



CONFIDENTIAL

Clearing the air on cloud computing

Discussion Document
March 2009

CONFIDENTIAL AND PROPRIETARY
Any use of this material without specific permission of McKinsey & Company is strictly prohibited.

McKinsey & Company

“ For I should not have rightly discovered things celestial if I had not suspended the intellect and the thought in a subtle form with its kindred air. But if, being on the ground, I should never have discovered them.

”

Aristophanes "The Clouds"

Context and objectives of this presentation



- **Using “clouds” for computing tasks promises a revolution in IT similar to the birth of the web and e-commerce**

- Much lower cost
- Faster time to market
- Great opportunities for creating new sources of value

- **While it has great potential, many of the claims being made about cloud computing have lead some to the point of “irrational exuberance” and unrealistic expectations**

- **The purpose of this report is to focus the nascent cloud industry and its consumers on setting realistic expectations by taking a “hype free” approach starting with the most basic question of what a “cloud” actually is**

- **Our specific objectives include**

- Proposing an industry-standard definition of cloud computing including what makes it unique and exciting
- Identifying the types of customers that should be early adopters—those that can benefit from existing and planned commercial offerings
- Understanding barriers that prevent large-scale adoption by corporations and government entities that represent the bulk of IT spending

Key findings

“Cloud computing” is approaching the top of the Gartner Hype-cycle



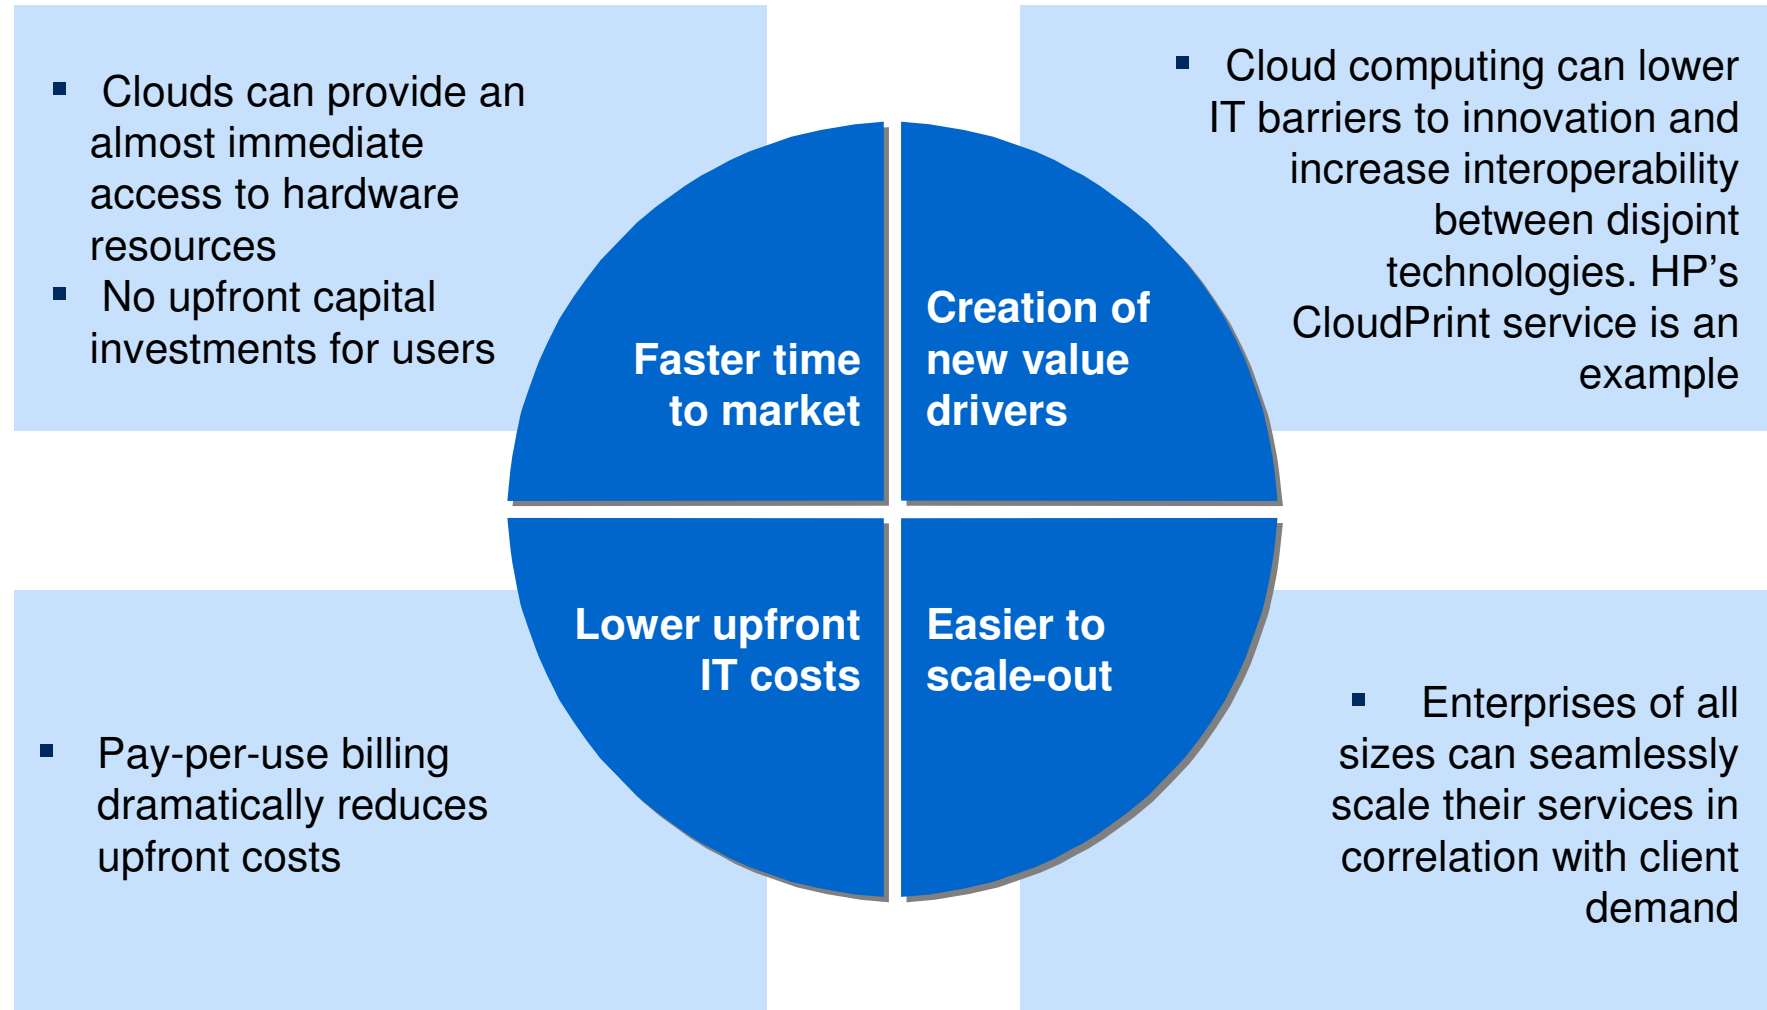
Getting an industry “fix” on the definition of what a cloud actually is would be a good first step and we propose one for adoption here

Clouds already make sense for many small and medium-size businesses, but technical, operational and financial hurdles will need to be overcome before clouds will be used extensively by large public and private enterprises

Rather than create unrealizable expectations for “internal clouds,” CIOs should focus now on the immediate benefits of virtualizing server storage, network operations, and other critical building blocks

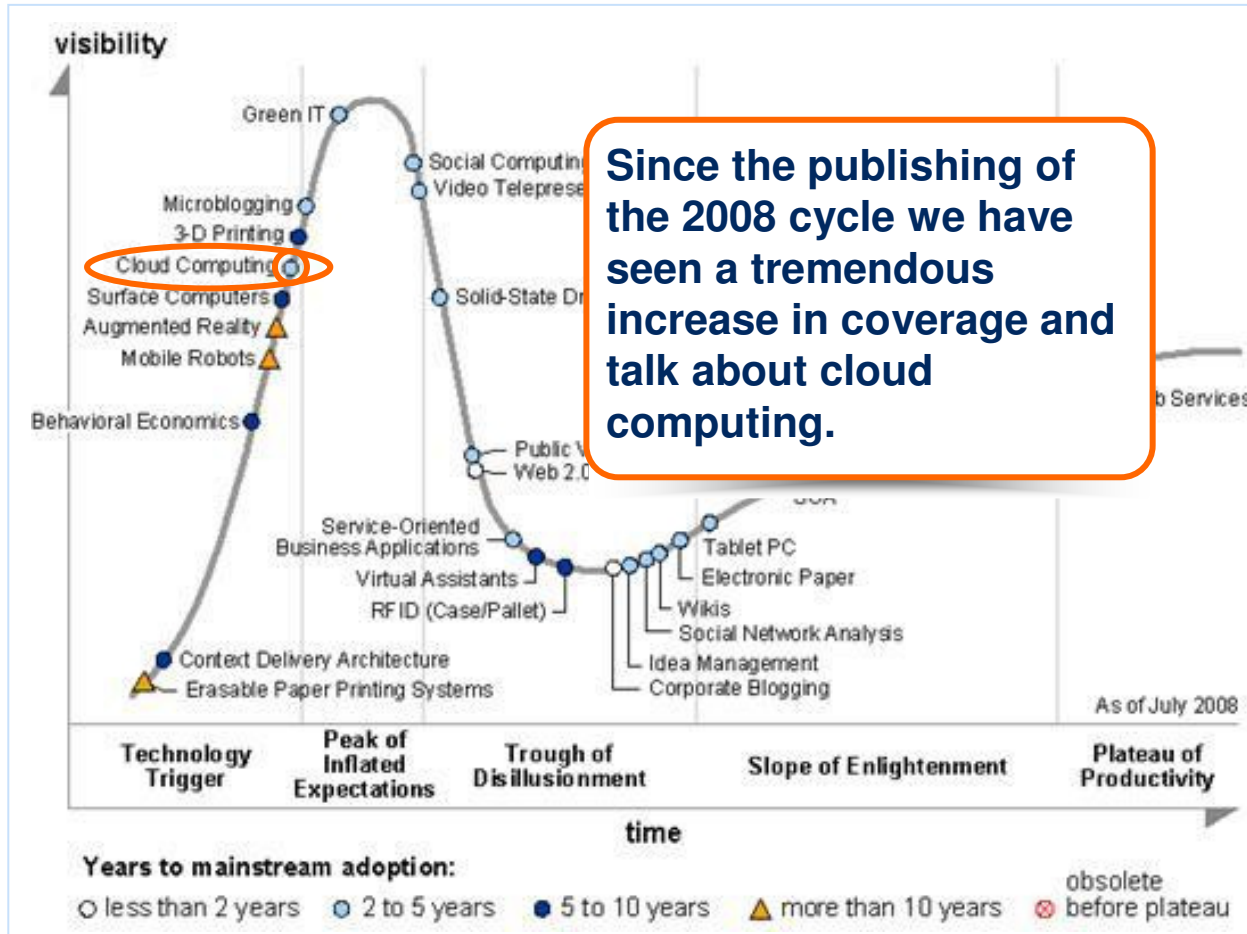
Users, hardware vendors and service suppliers can take specific steps to ensure the successful adoption of cloud technology—and prevent it from getting stuck in the “trough of disillusionment”

Cloud computing has tremendous promise



Cloud is quickly becoming top-of-mind in the IT world

2008 Gartner IT hype cycle



“The current U.S. economic woes will only drive more enterprises to consider and adopt cloud offerings. Spending on IT cloud services will hit \$42 billion by 2012”

IDC
Analyze the Future

“58 percent say cloud computing will cause a radical shift in IT and 47 percent say they're already using it or actively researching it”

CIO

The tremendous buzz around clouds is being further fueled by intense analyst and media speculation ...

“no less influential than e-business” (Gartner, 2008)

“ economic downturn, the appeal of that cost advantage will be greatly magnified” (IDC, 2008)

Not only is it **faster and more flexible**, it is **cheaper**. [...] the emergence of cloud models **radically alters the cost benefit decision**“ (FT Mar 6, 2009)

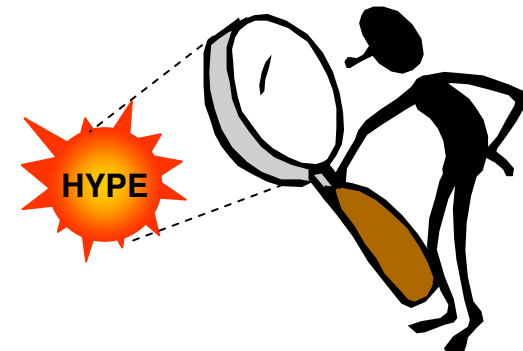


“Cloud computing achieves a **quicker return on investment** “
(Lindsay Armstrong of salesforce.com, Dec 2008)

“**revolution, the biggest upheaval since the invention of the PC in the 1970s [...]** IT departments will have little left to do once the bulk of business computing shifts [...] into the cloud” (Nicholas Carr, 2008)

The **economics are compelling**, with business applications made three to five times cheaper and **consumer applications five to 10 times cheaper** (Merrill Lynch, May, 2008)

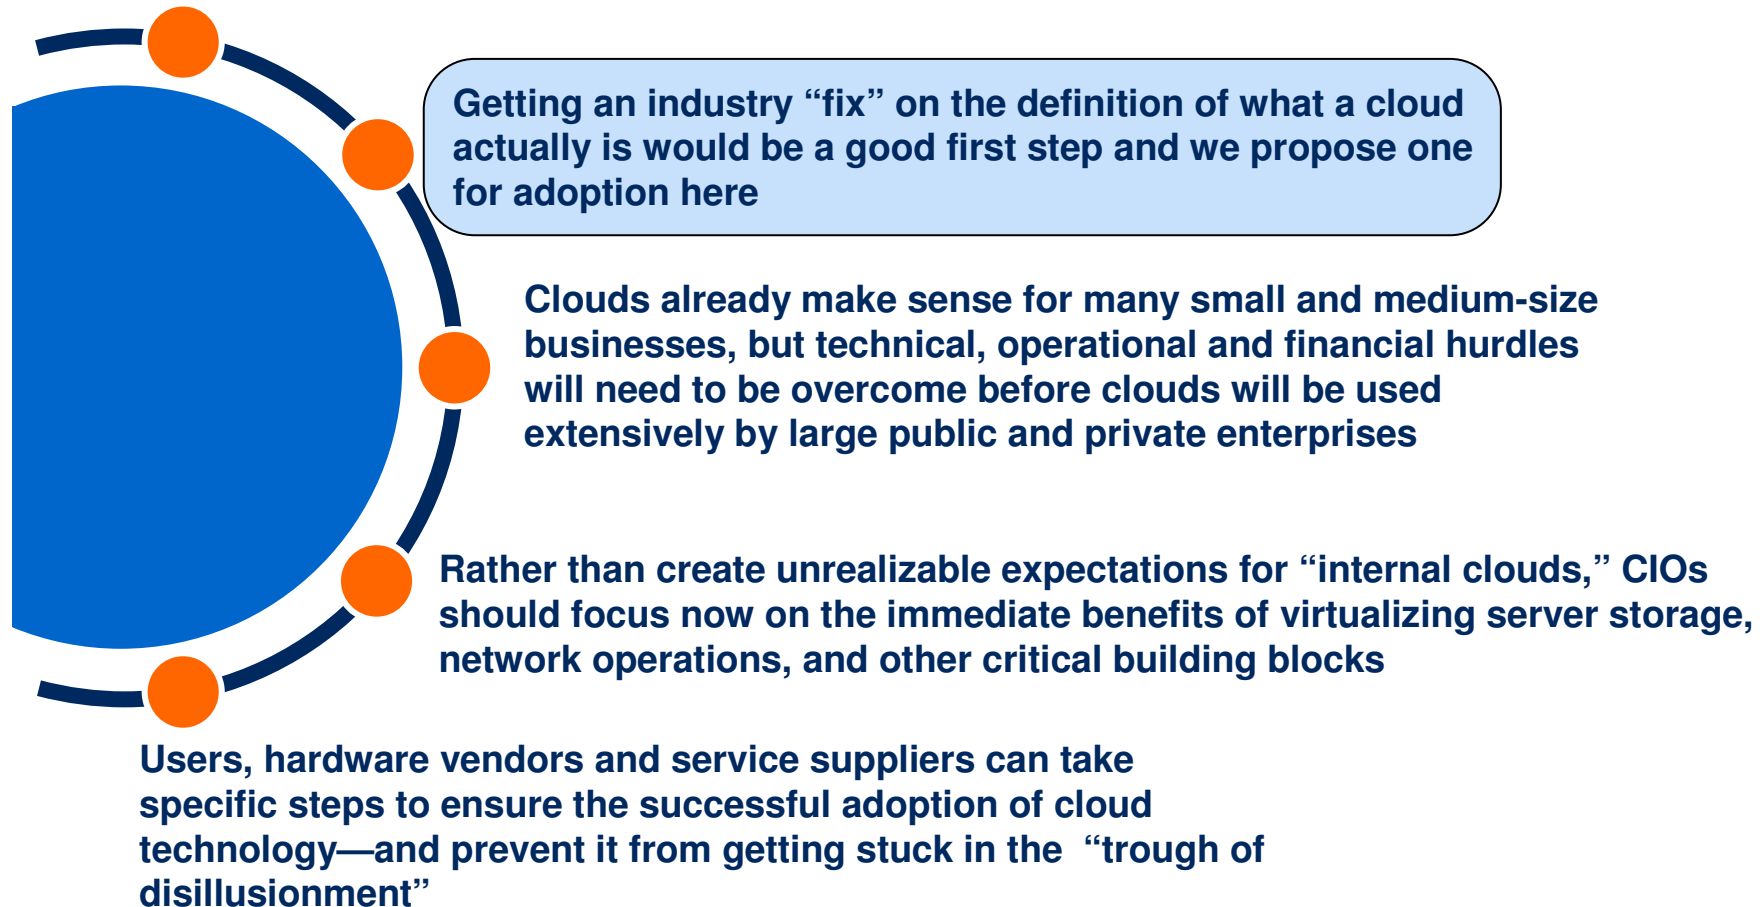
... and seems to be following the familiar IT hype pattern (for CIOs at least)

Problem	Works well in theory	The “Gold rush”	Lack of interest in technologies that work
	Cloud computing has shown great promise for start-ups and pet projects for large organizations. However, it is not ready to help with the big challenges of big companies	Initial excitement around a new technology grows rapidly into unrealistic expectations. Early on, the main beneficiaries of emerging technologies are usually the vendors, not the enterprise users	Cloud computing can divert IT departments’ attention from technologies that can actually deliver sizeable benefits; e.g., aggressive virtualization



Key findings


“Cloud computing” is approaching the top of the Gartner Hype-cycle



Why is it important to get a fix on clouds?

- Allows participants in cloud discussions to share a common understanding of what is meant by clouds
- A rigorous definition allows CIOs to be analytical in their decision making and to make more informed investment decisions
- Conversely, clarity also allows technology and service providers to build meaningful product, marketing, and sales strategies that translate into real value for their customers
- Allows an industry to begin moving forward with standards to promote interoperability amongst cloud products

A recent study has found 22+ definitions of clouds



Author/Reference	Year	Definition/Excerpt
M. Klems [11]	2008	you can scale your infrastructure on demand within minutes or even seconds, instead of days or weeks, thereby avoiding under-utilization (idle servers) and over-utilization (blue screen) of in-house resources...
P. Gaw [11]	2008	using the internet to allow people to access technology-enabled services. Those services must be 'massively scalable'...
R. Buyya [6]	2008	A Cloud is a type of parallel and distributed system consisting of a collection of interconnected and virtualized computers that are dynamically provisioned and presented as one or more unified computing resources based on service-level agreements established through negotiation between the service provider and consumers
R. Cohen [11]	2008	Cloud computing is one of those catch all buzz words that tries to encompass a variety of aspects ranging from deployment, load balancing, provisioning, business model and architecture (like Web2.0). It's the next logical step in software (software 10.0). For me the simplest explanation for Cloud Computing is describing it as, "internet centric software..."
J. Kaplan [11]	2008	a broad array of web-based services aimed at allowing users to obtain a wide range of functional capabilities on a 'pay-as-you-go' basis that previously required tremendous hardware/software investments and professional skills to acquire. Cloud computing is the realization of the earlier ideals of utility computing without the technical complexities or complicated deployment worries...
D. Gourlay [11]	2008	...the next hype-term...building off of the software models that virtualization enabled
D. Edwards [11]	2008	...what is possible when you leverage web-scale infrastructure (application and physical) in an on-demand way...
B. de Haff [11]	2008	...There really are only three types of services that are Cloud-based: SaaS, PaaS, and Cloud Computing Platforms. I am not sure being massively scalable is a requirement to fit into any one category
B. Kepes [11]	2008	...Put simply Clouds focused on capacity. This is a complete application and managed in the
O. Sultan [11]	2008	...In a fully implemented Data Center 3.0 environment, you can decide if an app is run locally (cook at home), in someone else's data center (take-out) and you can change your mind on the fly in case you are short on data center resources (pantry is empty) or you having environmental/facilities issues (too hot to cook). In fact, with automation, a lot of this can be done with policy and real-time triggers...
K. Hartig [11]	2008	...really is accessing resources and services needed to perform functions with dynamically changing needs...is a virtualization of resources that maintains and manages itself.
J. Pritzker [11]	2008	Clouds are vast resource pools with on-demand resource allocation...virtualized ...and priced like utilities
T. Doerksen [11]	2008	Cloud computing is ... the user-friendly version of Grid computing
T. von Eicken [11]	2008	outsourced, pay-as-you-go, on-demand, somewhere in the Internet, etc
M. Sheedan [11]	2008	... 'Cloud Pyramid' to help differentiate the various Cloud offerings out there...Top: SaaS; Middle: PaaS; Bottom: IaaS
A. Ricadela [11]	2008	...Cloud Computing projects are more powerful and crash-proof than Grid systems developed even in recent years
I. Wladawsky Berger [11]	2008	...the key thing we want to virtualize or hide from the user is complexity...all that software will be virtualized or hidden from us and taken care of by systems and/or professionals that are somewhere else - out there in The Cloud
B. Martin [11]	2008	Cloud computing encompasses any subscription-based or pay-per-use service that, in real time over the Internet, extends IT's existing capabilities
R. Bragg [5]	2008	The key concept behind the Cloud is Web application... a more developed and reliable Cloud. Many find it's now cheaper to migrate to the Web Cloud than invest in their own server farm ... it is a desktop for people without a computer
G. Gruman and E. Knorr [14]	2008	Cloud is all about: SaaS...utility computing...Web Services... PaaS...Internet integration...commerce platforms...
P. McFedries [22, 15]	2008	Cloud Computing, in which not just our data but even our software resides within the Cloud, and we access everything not only through our PCs but also Cloud-friendly devices, such as smart phones, PDAs... the megacomputer enabled by virtualization and software as a service...This is utility computing powered by massive utility data centers.

Two notable attempts at a rigorous definition that captures key characteristics

“Clouds are a large pool of easily usable and accessible **virtualized** resources (such as hardware, development platforms and/or services). These resources can be **dynamically reconfigured** to adjust to a variable load (scale), allowing also for an optimum resource utilization. This pool of resources is typically exploited by a **pay-per-use** model in which guarantees are offered by the Infrastructure Provider by means of customized SLAs.” - Vaquero et al

Does not distinguish cloud services from clouds and does not provide definitive economic implications

“Cloud Computing refers to both the applications delivered as services over the Internet and the hardware and systems software in the datacenters that provide those services...The datacenter hardware and software is what we will call a Cloud...Cloud computing has the following characteristics

1. The **illusion of infinite computing** resources...
2. The **elimination of an up-front commitment** by Cloud users...
3. The ability to **pay for use** ... as needed...” – UC Berkeley RAD Labs

Does not emphasize abstraction of infrastructure explicitly

SOURCE: “A Break in the Clouds: Towards a Cloud Definition”, 2009; Vaquero, Rodero-Merino, Caceres, Linder
“Above the Clouds: A Berkeley View of Cloud Computing”, 2009; UC Berkeley Reliable Adaptive Distributed Systems Laboratory

How we have defined a cloud

Definition: Clouds are hardware-based services offering compute, network and storage capacity where:

- 1 Hardware management is highly abstracted from the buyer
- 2 Buyers incur infrastructure costs as variable OPEX
- 3 Infrastructure capacity is highly elastic (up or down)

Characteristics of clouds

Characteristic: Enterprises incur no infrastructure capital costs, just operational costs and operational costs are incurred on a pay-per-use basis, with no contractual obligations

Definition: Clouds are hardware-based services offering compute, network and storage capacity where:

- 1 Hardware management is highly abstracted from the buyer
- 2 Buyers incur infrastructure costs as variable OPEX
- 3 Infrastructure capacity is highly elastic (up or down)

Characteristics of clouds

Characteristic: Enterprises incur no infrastructure capital costs, just operational costs and operational costs are incurred on a pay-per-use basis, with no contractual obligations

Definition: Clouds are hardware-based services offering compute, network and storage capacity where:

- 1 Hardware management is highly abstracted from the buyer
- 2 Buyers incur infrastructure costs as variable OPEX
- 3 Infrastructure capacity is highly elastic (up or down)

Characteristic: Capacity can be scaled up or down dynamically, and immediately, which differentiates from traditional hosting service providers

Characteristics of clouds

Characteristic: Enterprises incur no infrastructure capital costs, just operational costs and operational costs are incurred on a pay-per-use basis, with no contractual obligations

Definition: Clouds are hardware-based services offering compute, network and storage capacity where:

- 1 Hardware management is highly abstracted from the buyer
- 2 Buyers incur infrastructure costs as variable OPEX
- 3 Infrastructure capacity is highly elastic (up or down)

Characteristic: Capacity can be scaled up or down dynamically, and immediately, which differentiates from traditional hosting service providers

Characteristic: The underlying hardware can be anywhere geographically

Characteristics of clouds

Characteristic: Enterprises incur no infrastructure capital costs, just operational costs and operational costs are incurred on a pay-per-use basis, with no contractual obligations

Characteristic: Architecture specifics are abstracted. In addition, run in multi-tenancy mode with multiple users accessing the infrastructure simultaneously

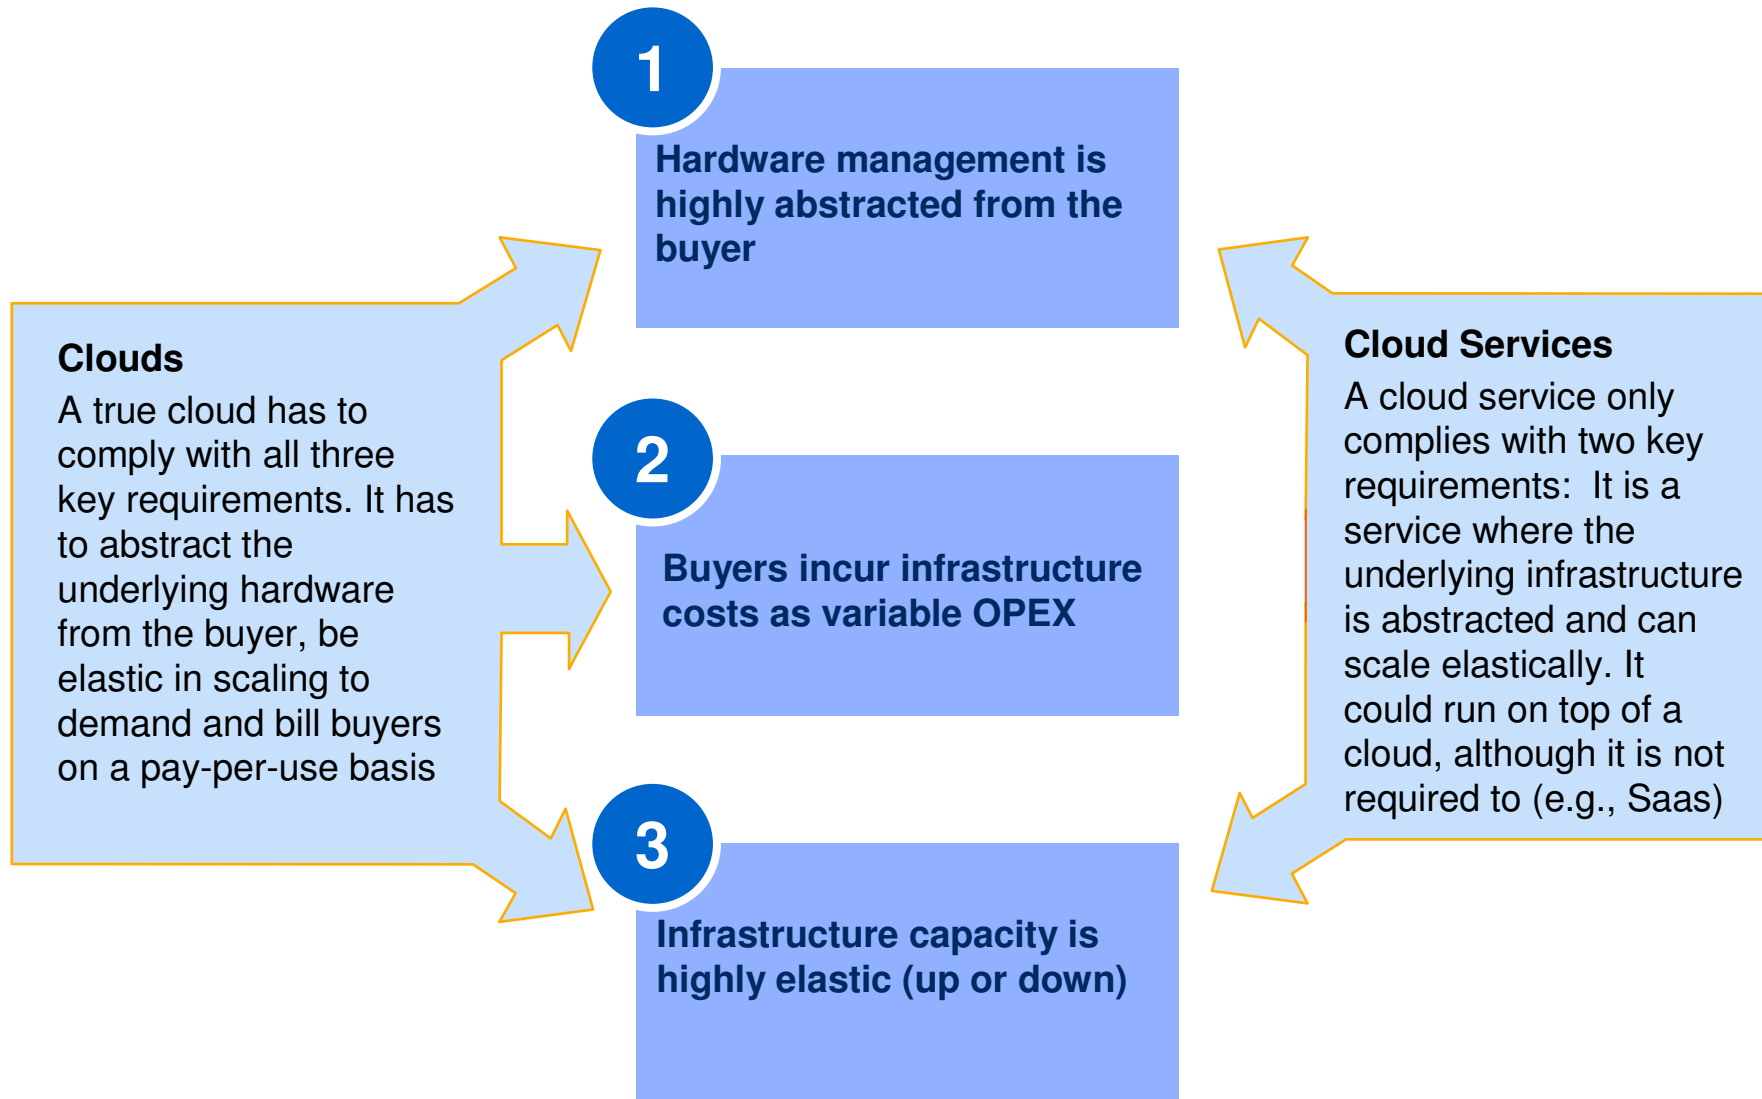
Definition: Clouds are hardware-based services offering compute, network and storage capacity where:

- 1 Hardware management is highly abstracted from the buyer
- 2 Buyers incur infrastructure costs as variable OPEX
- 3 Infrastructure capacity is highly elastic (up or down)

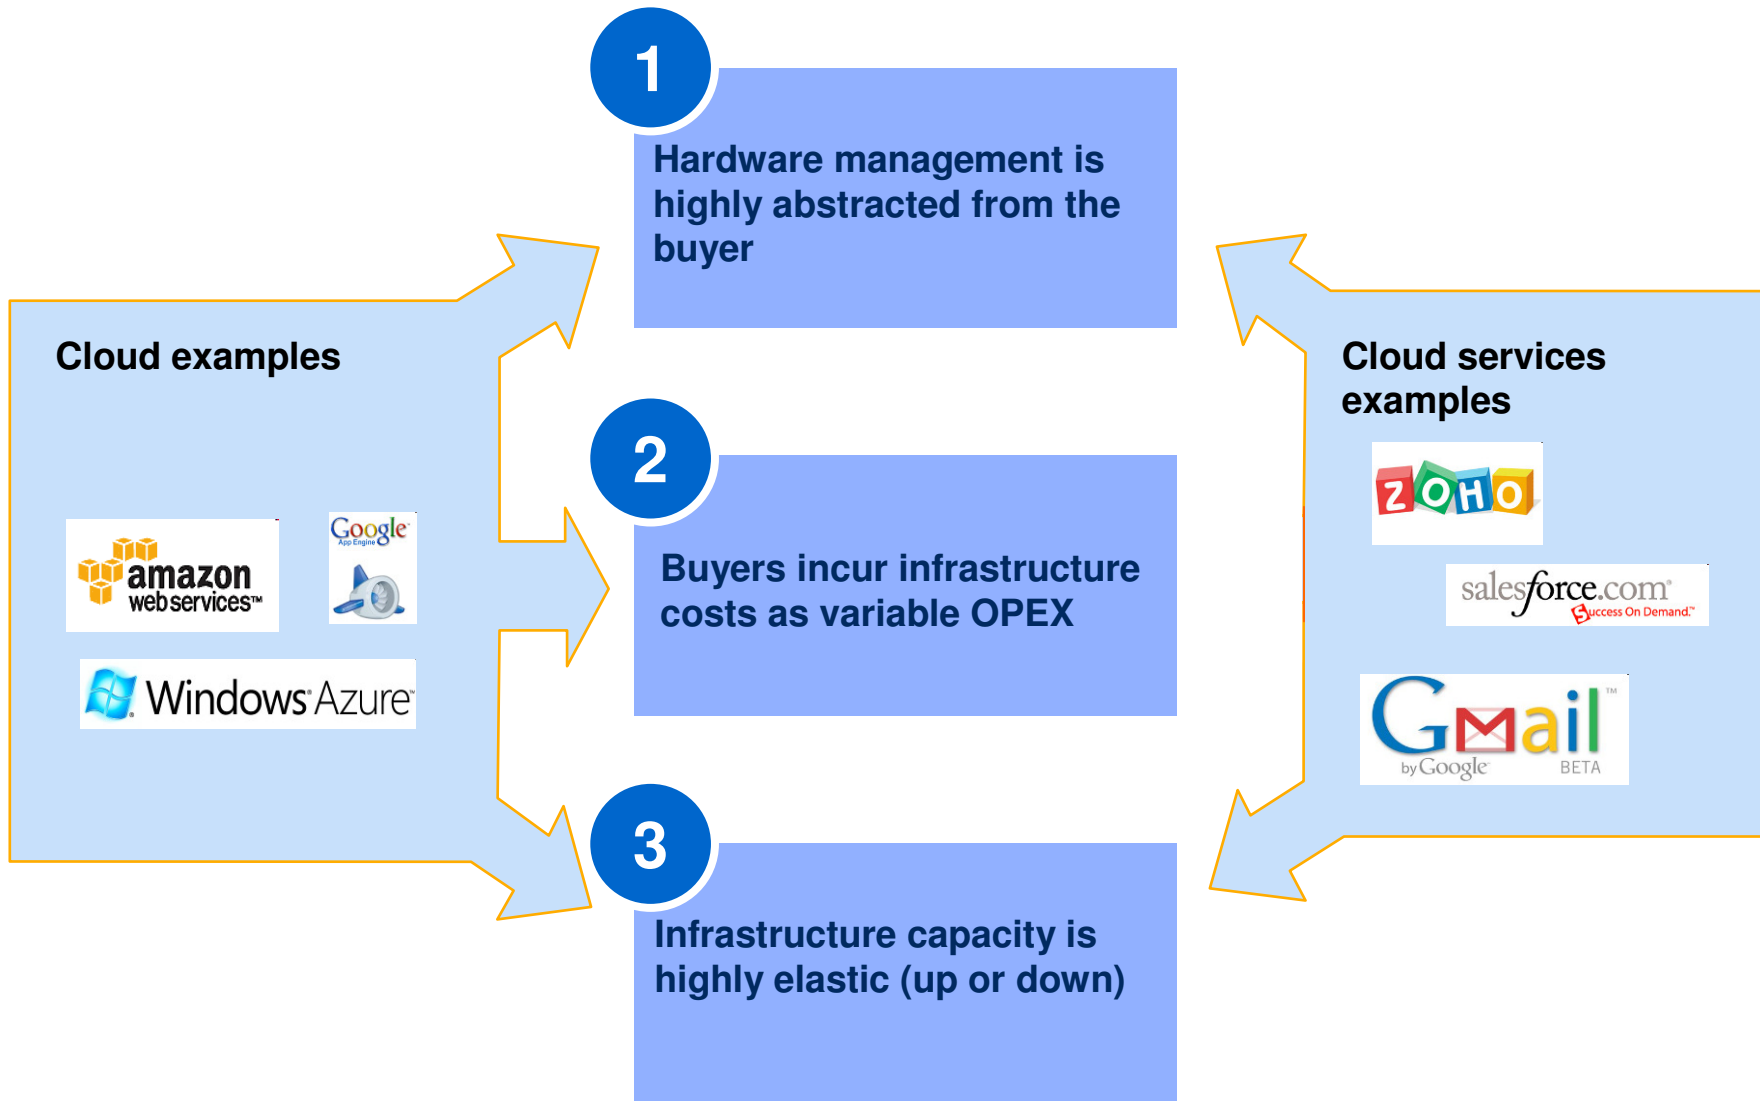
Characteristic: Capacity can be scaled up or down dynamically, and immediately, which differentiates from traditional hosting service providers

Characteristic: The underlying hardware can be anywhere geographically

Many cloud services are confused as clouds



Examples of clouds and cloud services




Key findings

“Cloud computing” is approaching the top of the Gartner Hype-cycle



Cloud offerings currently are most attractive for small and medium-sized enterprises

Clouds are very cost effective for SMEs ...



... and most customers of clouds are small businesses

unfuddle

SmugMug

JUNGLE DISK
Online storage powered by Amazon[®] and Rackspace[®]

ShareThis[®]

37signals

GIGAVOX
MEDIA

There are significant hurdles to the adoption of cloud services by large enterprises




Current cloud computing services are generally not cost effective for larger enterprises



EC2 monthly CPU equivalent price options (virtual cores)

\$



- Most EC2 options are more costly than TCO for a typical data center
- Enterprises could get lower TCO through pre-pay agreements—but only for Linux systems

The thin green line
TCA for typical data center¹ around \$45/month for CPU equivalent

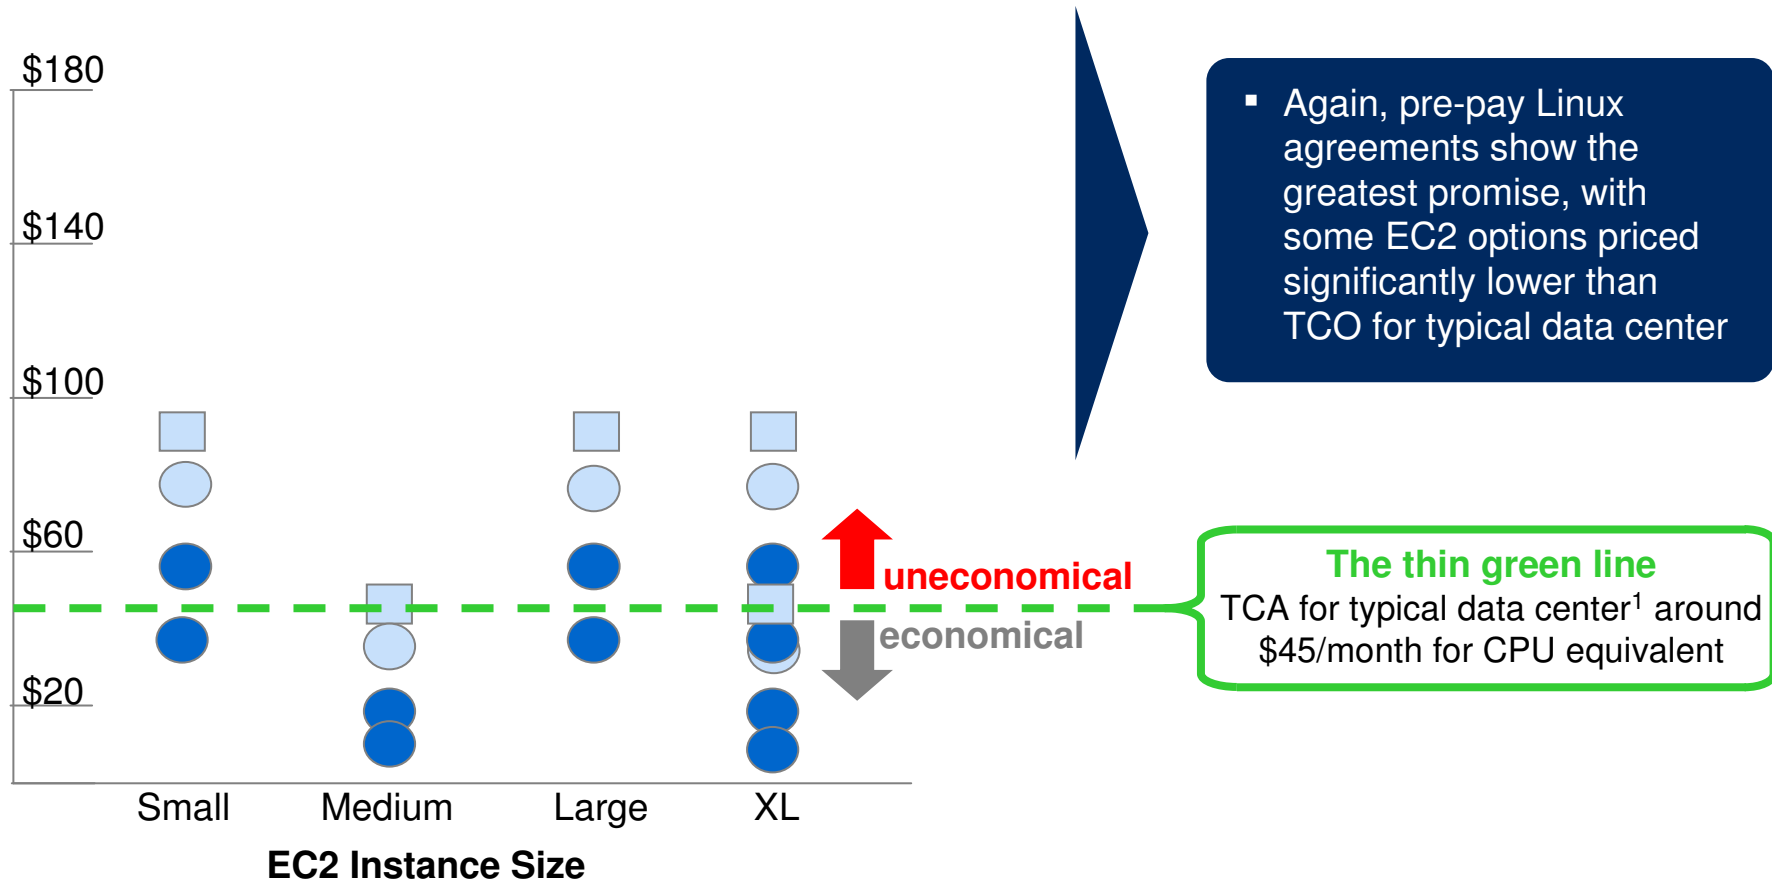
¹ Total Cost of Assets for “typical” data center: 10% utilization, \$20M/MW for facility, \$.1kW-hour, \$14K/Server (2 CPU, 4 core)

However, for smaller compute equivalents, today's clouds may be attractive for some computing work



EC2 monthly CPU equivalent price options (compute units²)

\$



Again, pre-pay Linux agreements show the greatest promise, with some EC2 options priced significantly lower than TCO for typical data center



The thin green line TCA for typical data center¹ around \$45/month for CPU equivalent

1 Total Cost of Assets for "typical" data center: 10% utilization, \$20M/MW for facility, \$.1kW-hour, \$14K/Server (2 CPU, 4 core)

2 One EC2 Compute Unit provides the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor

The cost of cloud must come down significantly for outsourcing a complete data center to make economic sense

DISGUISED CLIENT EXAMPLE



- Assumes migration of total Windows and non-console Linux capacity for entire data center
- Based on comparable hardware configurations, current pricing for cloud computing services is significantly higher than the CPU/month TCO achievable in data centers today.
- The key factor is that the majority of servers that can be migrated are Windows servers


1 Cost for comparable configuration – 75% of EC2 Large Standard Windows configuration
 2 Typical CPU/month cost for 3GHz dual-core Xeon Windows-based servers
 3 Estimated based on 10% labour savings from moving to a third-party cloud provider

There is a 10-15% total infrastructure labor base saving potential if moving whole data center to cloud

DISGUISED CLIENT EXAMPLE

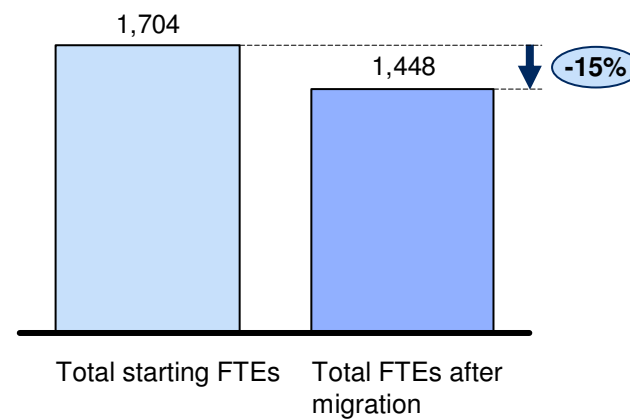
■ Total starting FTEs
■ Total FTEs after migration

FTEs by role group



- Approx 10-15% of headcount can be reduced by migrating capacity over to the cloud
- Most of this headcount comes from facilities and touch labor roles
- Consequently, labor savings are modest, but not insignificant...

Total FTEs



Many enterprise (necessarily or unnecessarily) set their SLAs uptimes at 99.99% or higher, which cloud providers have not yet been prepared to match

Amazon's cloud outages receive a lot of exposure ...

- July 20, 2008 Failure due to stranded zombies, lasts 5 hours
- Feb 15, 2008 Authentication overload leads to two-hour service outage
- October 2007 Service failure lasts two days
- October 2006 Security breach where users could see other users data

IT shops do not always deliver on their SLAs but their failures are less public and their customers can't switch easily

... and their current SLAs don't match those of enterprises*

Amazon EC2 **99.95%** Amazon S3 **99.9%**

It is not clear that all applications require such high services levels

* SLAs expressed in Monthly Uptime Percentages

Key findings

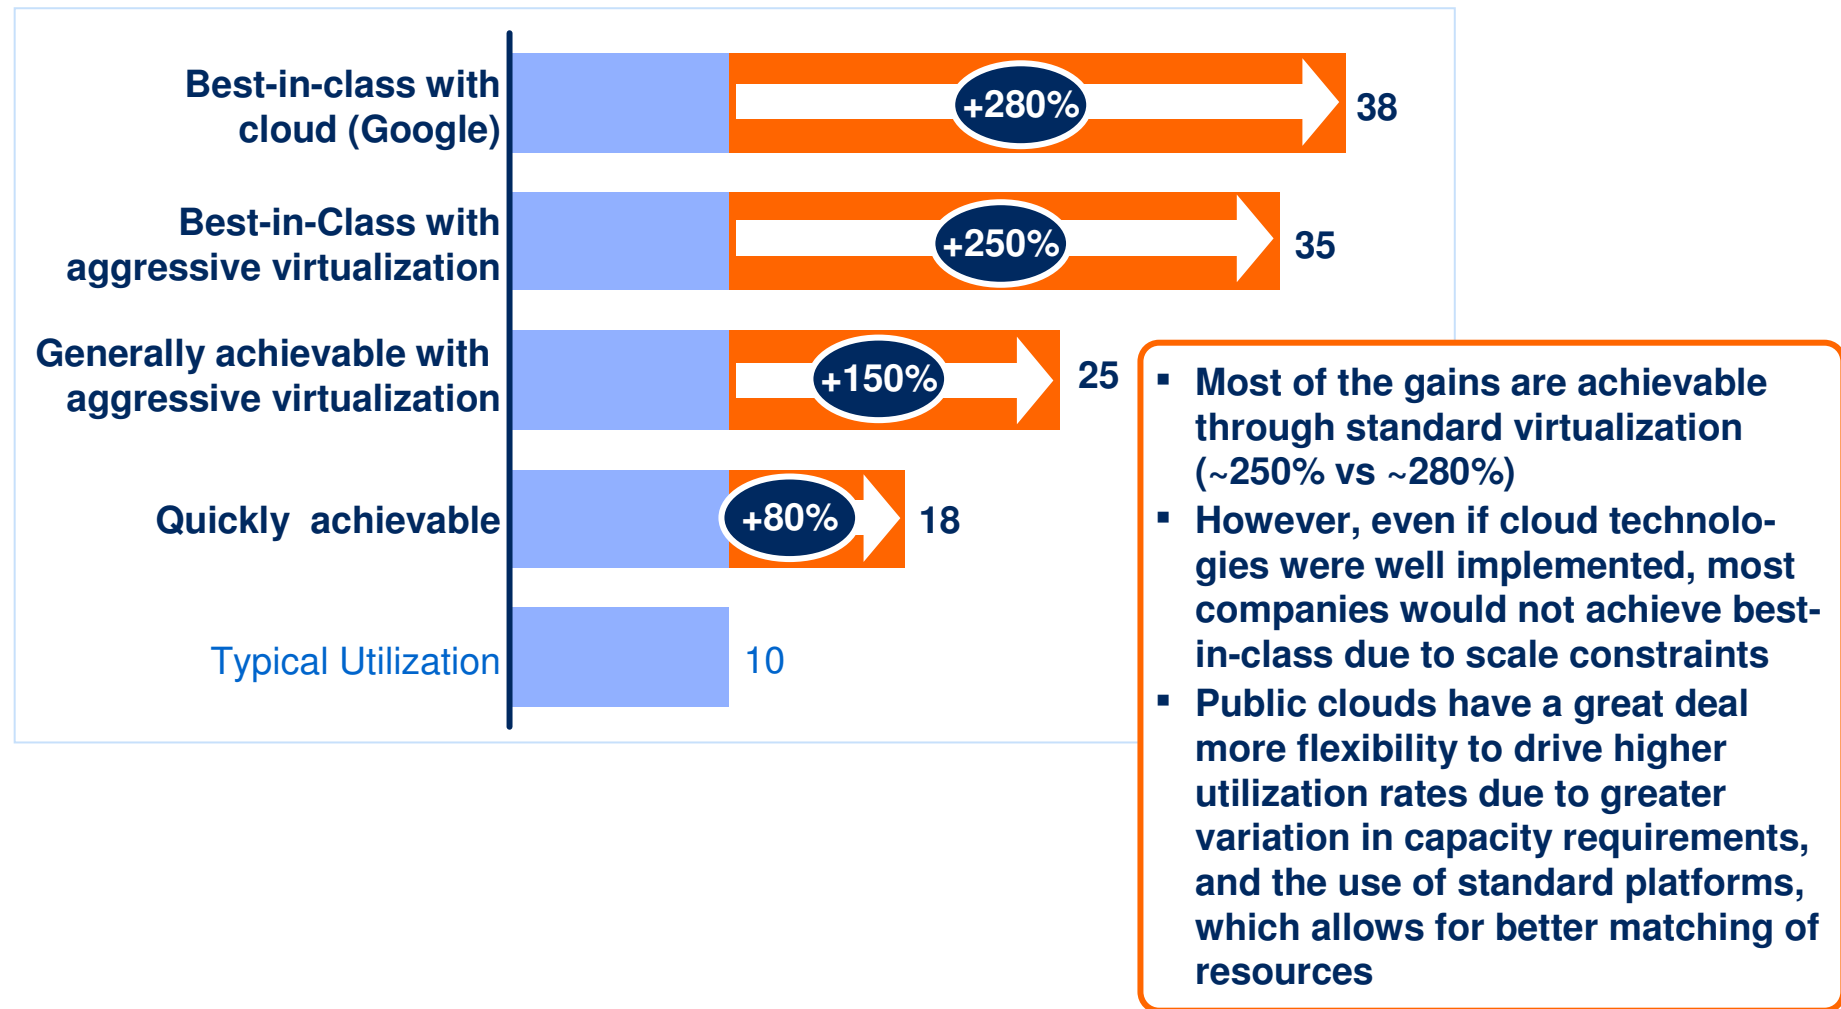
“Cloud computing” is approaching the top of the Gartner Hype-cycle



Users, hardware vendors and service suppliers can take specific steps to ensure the successful adoption of cloud technology—and prevent it from getting stuck in the “trough of disillusionment”

Large enterprises can achieve server utilization rates similar to those cloud providers are achieving from their platforms ...

Average server utilization rates

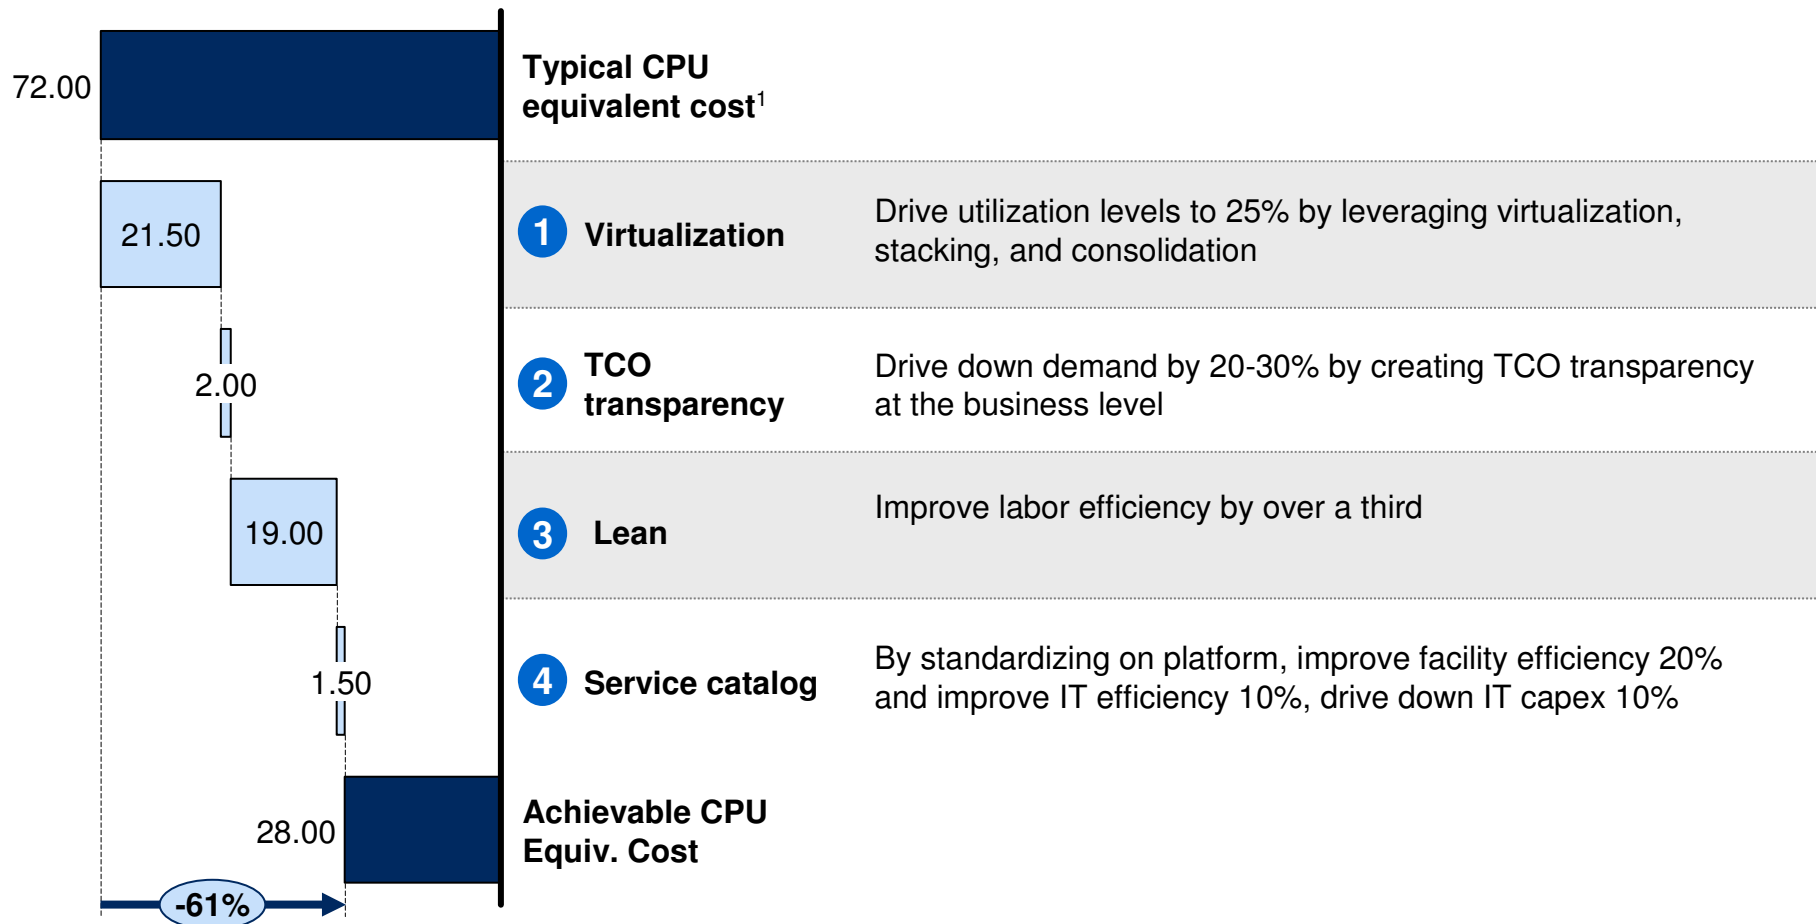


... and, by adopting data center best practices, can drive down server TCO by more than 50%

ILLUSTRATIVE

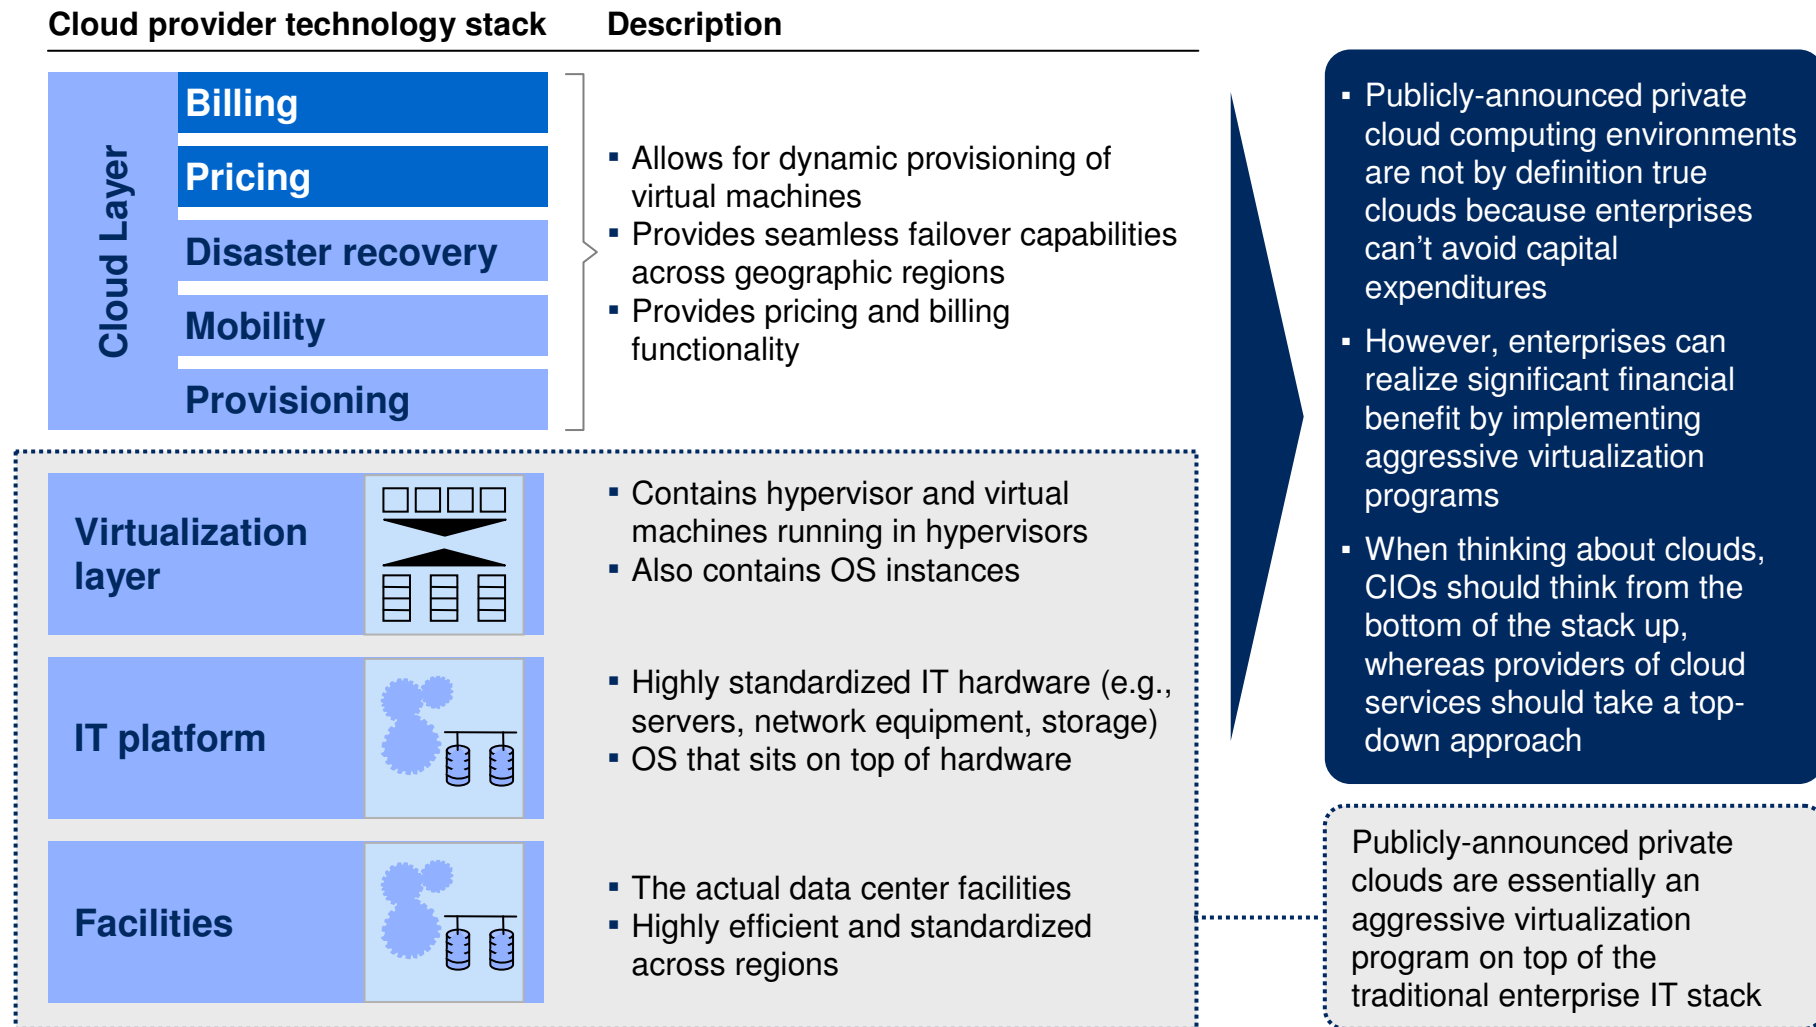
Key value drivers

Total monthly cost per CPU equivalent
\$, including labor



¹ TCO for "typical" data center: 10% utilization, \$20M/MW for facility, \$14K/Server (2 CPU, 4 core), PUE of 2.0, \$.1/kW-hour, facility utilization of 65%, 30-40 Server/FTE ration, \$100K/year/FTE, includes IT support and facility support labor

CIOs should be working the cloud technology stack bottom up; cloud providers top down



Key findings


“Cloud computing” is approaching the top of the Gartner Hype-cycle



Avoiding the trough of disillusionment will require appropriate action from all players in the cloud computing arena


Stakeholder groups	Cloud users (e.g. CIOs, CTOs)	<ul style="list-style-type: none">▪ Develop an overall strategy for XaaS based on solid business cases not “cloud for the sake of cloud”▪ Use modular design in all new and re-architected software to minimize costs when it comes time to migrate to cloud▪ Set up Cloud CIO Council to advise industry
	Cloud tools and infrastructure (e.g. HP, IBM)	<ul style="list-style-type: none">▪ Develop improved security standards to allay fears of client base▪ Implement technologies that will allow for fine grain billing and management across a cluster of compute devices
	Cloud providers (e.g. Amazon, Azure)	<ul style="list-style-type: none">▪ In the near term, focus on small and medium sized businesses▪ Work on improving uptime rates into the 99.99% range▪ Continue to drive down prices through scale/innovation to increase potential market

The 2000 dot-com bubble provides an extreme example of the dangers of investing in hype...



- Huge investments were made based largely on hype
- Eventually, the inability to generate profits led to the collapse of most of the boom-time dot-coms
- From peak to trough, the NASDAQ-composite lost ~ 80% of it's value

Further information:



Will Forrest

william_forrest@mckinsey.com

+1 (312) 551-3975



ayewill

#mckclouds



Media inquiries

Charlie Barthold

charles_barthold@mckinsey.com

+1 (203) 977-6915