

Key Debates Around Cloud Formation

A Look at Impact on VMW, RHT, MSFT, ORCL and others

- **Supply and demand side driving evolution to cloud infrastructure** – Economic and technology forces from both sides are driving enterprise customers to look to web-scale peers for efficiency and agility in underlying IT infrastructure. Transparent infrastructure pricing, commodity hardware components, the rise of the cloud management platform, lines of business driving the IT agenda, and an increasingly iterative pace of software development cycles are driving a pace of change in enterprise IT infrastructure not seen in a generation. We expect the result to be adoption and evolution of cloud architectures, which we label Cloud 1.0 (virtualized), 2.0 (automated) and 3.0 (native resource consumption).
- **Applications will move the cloud more slowly and less simply than the consensus view** – The demographics of enterprise apps are diverse and not all will benefit from natively consuming cloud resources. To achieve Cloud 3.0 “nirvana”, apps will have to be re-written and with enterprise application lifecycles of around 20 years, the shift in architecture will take longer than many investors expect.
- **Even if all apps ultimately end up in the public cloud, a hybrid architecture is required for some apps** – Several factors may conspire to keep applications from going quickly to the public cloud, including economic considerations, control / governance concerns and inability of public cloud to satisfy “long-tail” requirements. We expect the hybrid scenario will be the on-ramp for enterprise customers to the public cloud may be a long-term intermediate stage for many workloads.
- **Open standards may have less impact than “de-facto” standards** – Investor consensus is that the cloud must be “open.” However we believe true openness is difficult to achieve, something that is already showing up in the rise of OpenStack. Meanwhile Amazon is rising as the emerging de-facto standard, benefitting from significant network effects.
- **Software stock implications** – VMware is a winner if enterprise migration to cloud architectures happens in an evolutionary fashion. So far, early indications around hybrid cloud services (vCHS) are positive and if the company can capitalize on storage and network, it is well positioned to capture share in the cloud control plane. Microsoft has is one of three players with web-scale public cloud service. However, we believe with a Windows-only focus, it will be difficult for Microsoft to complete broadly in IaaS and instead must be successful reinvigorating its developers around the Azure platform as a service (PaaS) offerings. Red Hat has firmly made a commitment to OpenStack and is at the mercy of the timing of maturity of OpenStack, something that could still be 18-24 months away. We continue to believe service provider adoption is the harbinger of enterprise demand. Oracle is behind in public cloud and while its workloads are amongst the stickiest, the company must close the gap as enterprises become more comfortable with IaaS.

Walter H Pritchard, CFA

+1-415-951-1770

walter.h.pritchard@citi.com

Robert Chen

robert.chen@citi.com

See Appendix A-1 for Analyst Certification, Important Disclosures and non-US research analyst disclosures.

Citi Research is a division of Citigroup Global Markets Inc. (the “Firm”), which does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the Firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. Certain products (not inconsistent with the author’s published research) are available only on Citi’s portals.

Contents

Executive Summary	4	Implications of a de facto standard	33
Brief Cloud Primer	5	Network effects = winner take all if the tracks are laid fresh	34
Cutting through the fluff: What's different about Cloud Computing	5	What keeps AWS from becoming the next Wintel? (i.e. the de-facto standard)	34
Evolution in the way computing capacity is built and consumed	5	Stock implications	35
The Supply Side: Computing architectures are changing	6	Overall Cloud Formation Stock	
Evolutionary steps to cloud nirvana	7	Conclusions	37
First, some distinctions: What makes a cloud a cloud?	7	What you have to believe for each to be a winner	37
Cloud 1.0: Talks like a cloud but doesn't walk like a cloud	8	VMware – Control of Mission Critical Workloads and a differentiated IaaS offering are key	38
Cloud 2.0: Enterprises adopting cloud principles into the data center	9	VMware's entry into public cloud is early and behind on AWS definition of market	38
Cloud 3.0: Not all apps go to cloud heaven?	10	VMware is in a tough spot if enterprise moves quickly to cloud	38
The Demand Side: Trends in software development align with cloud	10	More likely case is slow migration to cloud, which favors VMware approach	38
Agile development and DevOps go hand-in-hand with Cloud Computing	11	Pay attention to mission-critical workloads and success with network and storage as harbingers of cloud success	39
Debate #1: How Fast Will Applications Move to the Cloud?	13	Success may never be in PaaS and new custom app building	39
Options for migrating to the cloud	14	What about VMware and OpenStack?	39
Catalog of applications is diverse	14	Stock view – We believe investors are too negative around VMware cloud position	39
Costs of moving to the cloud are high	15	Microsoft – Symmetry for Windows and web-scale public cloud offering is the Hope	40
Application lifecycles imply a 20-year migration cycle	15	Microsoft is locked into Windows for hybrid value prop	40
If some apps don't move quickly to the cloud, who benefits?	15	Microsoft needs to re-invigorate the developer base around cloud	40
Debate #2: Public vs. Private Cloud Prizefight	18	Azure will play to non-Windows developers, but that sounds like a tough battle to win as long as Windows is the strategic focus	41
Does cost matter?	19	Stock view – Best positioned mega cap for cloud, still with Windows baggage	41
Amazon is cheaper but more expensive	19	Red Hat – Is OpenStack the answer to RHEL cannibalization by the public cloud?	41
The public cloud has better economics	19	Big investment in OpenStack	41
Bridging the cost gap: Cost is not the same as price, and the price printed is not always the price you pay	20	Can RHEL make the transition?	41
It's a Control Thing	23	Stock view – Best way to play OpenStack but unlikely to be the medium-term growth driver	42
Is the Hybrid Cloud the answer, especially in the short run?	23	Oracle – Secular loser in cloud unless pace of innovation picks up	42
True Hybrid Clouds can be hard to build	23	No significant exposure to IaaS infrastructure software	42
The Long March to the Public Cloud	24	Application middleware and database in theory under threat from Cloud 2.0 / Cloud 3.0, but these may be sticky for some time	42
Stock implications	25		
Debate #3: Will the Cloud be Proprietary or Open?	28		
Why open vs. proprietary matters	29		
The multiple meanings of "open"	29		
Open core has largely superseded open source and is creating a Tower of Babel	30		
Open source code is a red herring, but could APIs be the answer?	30		
API compatibility is also a red herring	32		

Primary cloud strategy is a focus on appliances	43
Dark horse in Nimbula/OpenStack, which opens up Oracle appliances	43
Stock view – Mostly negative implications from Cloud with a dark horse	43
Citrix	43
Appendix A-1	44

Executive Summary

The content here was first published in a three-part series of “Cloud Formation” reports. We are republishing all three parts into one report for the sake of convenience. In this report, we address the key questions surrounding the market evolution of the infrastructure layer of cloud computing known as Infrastructure as a Service (IaaS). For our view on the Software as a Service (SaaS) market, see our report [State of SaaS: A Look at Adoption, Economics and Valuation Implications](#). We start with a quick, intermediate-level primer, then go in-depth on three key debates, and conclude:

- Brief Cloud Primer
- Debate #1: How Fast Will Apps Move to the Cloud?
- Debate #2: Will Public or Private or Hybrid Clouds prevail?
- Debate #3: Will the Cloud be proprietary or open?
- Overall Stock Conclusions

Brief Cloud Primer

Cutting through the fluff: What's different about Cloud Computing

The modern cloud was born with Amazon Web Services in 2006...

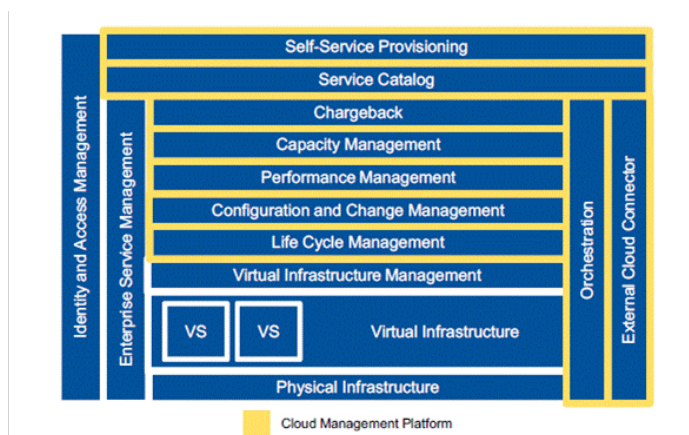
Variants of cloud-like computing have been around since the mainframe, a platform where server virtualization technology was introduced in 1972. In the late 1990s, the computing vendors such as H-P (HPQ.N; US\$29.58; 1) and Sun pushed "utility computing" initiatives, allowing customers to provision virtual private data centers through a graphical user interface. These later efforts, however, lacked the complete set of modern cloud attributes such as pooled infrastructure resources or high-levels automation. While Internet access, server virtualization, server hardware commoditization, and early efforts to build metering software were prerequisites for cloud computing, we trace the true launch of cloud computing to Amazon's (AMZN.O; US\$353.65; 1) introduction of Amazon Web Services (AWS) in 2006.

Evolution in the way computing capacity is built and consumed

...driven by shifts in both technology and technology consumption.

IT architectures are evolving, and so are the applications that are driving computing demand. New cloud architectures are building to a "scale-out" model instead of the traditional "scale-up" model, leveraging commodity components and copying innovations first invented by web companies to provide services at massive scale. New projects such as OpenStack (see our report [Is OpenStack for Real and How Does It Impact the Software Landscape?](#)) represent a new software layer called the Cloud Management Platform (CMP), which automates provisioning, configuration, and management of compute, storage and network fabrics.

Figure 1. The Cloud Management Platform offers new capabilities



Source: Gartner

On the demand side, the consumers of IT (mostly lines of business and others looking to use IT as a competitive differentiator) are demanding that IT be more responsive to their business needs. Lines of businesses want to move at a faster pace—for example, building new applications, driving innovations in these applications more rapidly, and reducing the amount of time it takes to provision new infrastructure to support these applications. Within the software world, there is a shift toward "agile" software development, resulting in the creation of a "DevOps" function (the merger of traditionally siloed development and operations), which enables shorter release cycles. These new applications are forcing traditional IT architectures to keep up with the high service levels at external public cloud service

providers (which many lines of business now often consider as an alternative for building new applications).

The Supply Side: Computing architectures are changing

While Amazon popularized cloud computing, modern cloud principles are driving a new architecture with very different design principles than traditional computing. Perhaps the best way to think about modern cloud computing is to understand an enterprise's treatment of servers. Bill Baker (formerly with Microsoft) and Randy Bias (CloudScaling) have popularized the analogy of servers being treated as “pets” in a scale-up model and as “cattle” in a scale-out model.

Figure 2. The pets vs. cattle analogy



	Pets	Cattle
Server name	BigDog.company.com	vm5802.company.com
When sick:	Nurse	Kill & Replace
Architectural model	Scale up	Scale out
Key software enabler	Hypervisor	CMP
Primary consumer	Sys admin	App developer
Fault tolerance resides in:	Infrastructure	Application

Source: Citi Research, adapted from Bill Baker and Randy Bias

In the pet world, a server is treated with extreme care and given the attention, resources, and customizations it needs to fulfill its duty. Each server is manually tuned to run at optimum utilization and management tools are deployed to ensure maximum uptime of a server. Fault tolerance is embedded in the infrastructure horizontally, many times with proprietary hardware features. In the cattle world, servers are largely anonymous and replaceable. For example, Yahoo (YHOO.O; US\$38.31; 1) walks through its data centers every two months to remove dead physical servers. Virtual servers are “stood up” instantly and are “put down” with ruthless regularity. In the cattle view of the server world, the emphasis is on automation of the compute fabric to avoid human intervention. Servers are built to scale out without human interaction. Implicit in this is “control plane” software (i.e. the CMP) to make the compute, storage, and network fabrics fully automated and fully elastic. Concurrently, applications need to be aware of the new control plane—they are designed with the necessary fault tolerance built-in.

In all, the “pets” model represents the vast majority of servers that are operating in data centers today (even encompassing virtualized servers), while the “cattle” population today primarily consists of servers operated by web 2.0 companies, who have built applications from the ground up within the last decade. In all, the virtualized (pets) vs. cloud (cattle) infrastructure worlds are very different. We don’t believe the market fully understands the implications of this difference and the speed with which enterprises can make the transition.

Evolutionary steps to cloud nirvana

The cattle scenario sounds appealing—especially in a world where hardware is moving toward standard commodity parts and data center components are being instantiated in software. In the web 2.0 world, companies are new and unencumbered by legacy infrastructure. These companies have been able to adopt the cloud paradigm more quickly than enterprises. However, our focus is on the enterprise, which is the primary driver of revenue today for publicly traded tech infrastructure companies. In the enterprise, many legacy applications and architectures coexist. For private and hybrid clouds, we see enterprises taking several intermediate steps on their ascendance into the cloud. In reality, early cloud architectures deliver some but not all the benefits of cloud computing. In the section below, we describe a common path that we see enterprises taking toward cloud nirvana.

First, some distinctions: What makes a cloud a cloud?

The key elements of cloud computing
require a new set of software capabilities.

The US Government’s National Institute of Standards and Technology (NIST) defines cloud computing as having five key characteristics (shown in the Figure below). Previous variants of cloud computing delivered on some but not all of the elements of cloud computing. Broad network access has been available since the Internet, and resource pooling was largely enabled by virtualization. Other elements of cloud computing such as self-provisioning and elastic release of resources require some additional capabilities of a CMP, where automation and management are core capabilities. For example, new automation tools such as Chef and Puppet enable developers to think about “infrastructure as code” – i.e. abstracting the underlying hardware elements into software recipes that can automate provisioning and configuration.

Figure 3. Cloud Computing Definition – Old World vs. New World cloud implementations

Cloud characteristics	NIST Description	Old world (e.g. HP UDC)	Cloud 3.0 (e.g. AWS, Azure)
Self Service	..Unilaterally provision computing...without requiring human interaction	Provision in days. Human intervention required to set-up.	Provision in minutes, fully automated
Broad Network Access	...accessed through standard mechanisms...(e.g. mobile phones, tables, laptops)	Yes	Yes
Resource Pooling	using a multi-tenant model...customer has no control over exact location...	Managed hosting model (no server virtualization)	Fully multitenant, high consolidation ratios
Rapid elasticity	elastically provisioned and released...capabilities appear to be unlimited...	Capabilities limited to racks pre-provisioned	Seemingly limitless resources
Measured service	Resource usage can be monitored, controlled, and reported	The "computon"	Automated metering and billing (\$/server-hour)

Source: NIST, Citi Research

Cloud 1.0: Talks like a cloud but doesn't walk like a cloud

Adding to the confusion is the choice of the industry to label data center consolidation using server virtualization technology (synonymous with VMware (VMW.N; US\$95.61; 1)) as "private cloud." We label these consolidated data centers as "Cloud 1.0", even though they largely pre-dated AWS and other modern clouds and were inferior to modern clouds in the five cloud characteristics above. Nevertheless, Cloud 1.0 and the data center consolidation it enabled helped drive savings from reducing physical footprints (hardware, real estate, and power consumption) as well as improving operational efficiency (shorter or no downtime and cheaper, software-based disaster recovery). Virtualization alone doesn't offer maximum elasticity, since the elasticity is largely limited to the Cloud 1.0 resource pool of a single server or a limited cluster of servers. Self-service, automated provisioning and delivering resources "as a service" were not typically part of the pre-cloud architecture, and Cloud 1.0 adopters largely built their own tools to enable these other capabilities. This Cloud 1.0 world drove better resource pooling through higher utilization levels, but it lacked the automation and elasticity that characterizes modern clouds.

Figure 4. Evolutionary (or revolutionary?) steps to Cloud 3.0

	Cloud 1.0	Cloud 2.0	Cloud 3.0
Investment required	Server virtualization	Re-platform, modest automation	Application re-write, CMP deployment
Economic gain	Data center consolidation (modest resource pooling)	More resource pooling (Greater density, more automation)	Full cloud benefits (true elasticity, full automation)
Key technologies	Server virtualization	Configuration management, re-platforming, exposed REST APIs, backward compatibility	Full next-gen CMP
Key products	vSphere, Hyper-V, KVM	vCloud, BMC CLM, IBM SDM	AWS, OpenStack, Chef, Puppet

Key discontinuity; not all apps will make this leap

Cloud Characteristics	On-demand self service			
	Broad network access			
	Resource pooling			
	Elasticity			
	Measured service			




Source: Citi Research

Cloud 2.0: Enterprises adopting cloud principles into the data center

Cloud 2.0 is about reducing heterogeneity and increasing automation.

A clear trend within enterprises is an ongoing evolution of their Cloud 1.0 architecture into one with more cloud-like characteristics. What differentiates Cloud 2.0 from Cloud 1.0 is a greater degree of homogeneity in applications, which in turn enables greater automation and therefore higher compute densities. No longer are enterprises forklifting legacy apps onto virtualized infrastructure, but instead they are minimizing heterogeneity and introducing some automation to internal cloud services. Cloud management platforms (CMPs) or middleware that are backward compatible with legacy applications such as VMware's vCloud, BMC Cloud Lifecycle Manager, or Apprenda form the basis for many of these Cloud 2.0 installations. For example, one large financial institution is re-platforming their Java applications onto a common middleware stack such that provisioning and configuration of the middleware down becomes highly automated. At another large financial institution, their Cloud 2.0 initiative is likely to move forward on a CloudStack implementation from Citrix (CTXS.O; US\$57.48; 2) which has CMP components to automate provisioning of commodity resources while having backwards compatibility with VMware infrastructure. This Cloud 2.0 model finds a common ground between disruptive application re-writes but that common ground prevents Cloud 2.0 from delivering on the full promise of Cloud 3.0.

Figure 5. While Cloud 2.0 solves many Cloud 1.0 problems, backward compatibility with legacy environments and code prevents it from seeing full Cloud 3.0 benefits

Cloud 1.0 problems	Cloud 2.0 approaches	Description
Manual intervention for provisioning/deprovisioning		Adopt lifecycle management capabilities while keeping <u>backward compatibility with legacy tooling</u>
VM sprawl, inconsistent policies leading to unique configurations		Additional tooling to manage ephemeral workloads and new IT policies <u>while integrating with ITSM frameworks</u>
Heterogeneous middleware requires matching infrastructure		Re-platform existing applications to prepare for Cloud 3.0 <u>while minimizing code rewrite</u>

Source: Citi Research

Cloud 3.0: Not all apps go to cloud heaven?

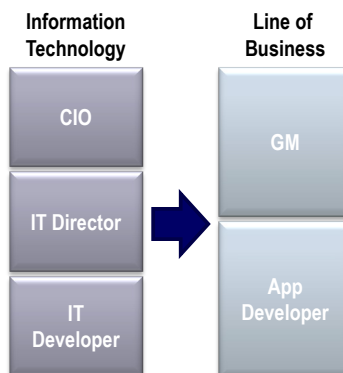
Cloud 3.0 goes hand in hand with the new software development model of “DevOps” (the underlying drivers of which are described below). Here applications themselves are programmed to directly consume infrastructure resources (compute, storage, and network) through application program interfaces (APIs) with certain service level requirements (resiliency, fault tolerance, latency tolerance) as part of the application code. Fundamental within the design of the app is a cloud environment where resources, response times, and undocumented behavior are not guaranteed. The majority of Web 2.0 apps (e.g. Netflix, Twitter, and Zynga) are virtually all written for Cloud 3.0 architectures. However, there is considerable debate about whether Cloud 3.0 architectures are the right fit for all applications and whether a traditional application like an Oracle financial accounting app is ever destined for the cloud. Our view (which we discuss further in Debate #1 below) is that this question is partly about the demographics of applications (some apps are more cloud ready than others) but more importantly it’s about the timeline to get to cloud nirvana. We offer a SaaS analogy for IaaS: How long will it take SaaS financials to displace Oracle or SAP (SAPG.DE; €58.50; 2)? Oracle and SAP application lifecycles are typically 30+ years and will likely take some time to displace. Not all applications will make it to the Cloud 3.0 architecture. For some applications, it will never make economic sense to re-write as this process is too expensive or the application lacks benefits of moving to the cloud. However, it’s useful to use the framework above to understand how the supply side (cloud builders) is offering various options to their customers.

The Demand Side: Trends in software development align with cloud

Increasingly, demand for applications is being driven by line of business (LOB) owners, who historically relied on their internal IT departments who have often failed to deliver on software development projects in a timely or satisfactory manner. These LOB owners were first to latch onto Software as a Service (SaaS). SaaS not only takes advantage of the ubiquitous web-based client, but is run as a service by the vendor, enabling customers to immediately have access to new capabilities with little up-front investment. These benefits, among others, have enabled SaaS as a category to take share from traditional on-premise software (see our report [State of SaaS: A Look at Adoption, Economics and Valuation Implications](#)). This ability of SaaS applications to keep up with the requirements of faster-moving LOBs has parallels to drivers of cloud infrastructure adoption, including agile development and “DevOps”. It is indeed LOBs that are increasingly trying to differentiate themselves


from competition with technology and driving IT to deploy cloud infrastructure (or going outside of IT by consuming public cloud capacity). As some commentators have noted, LOB application owners are multiplying: “every company is a software company” (recent Forbes article) and “software is eating the world” (Marc Andreessen).

Figure 6. They buyer of compute resources is shifting to the line of business



Source: Adapted from VMware analyst day presentation, August 2013

Figure 7. New cloud capabilities are putting pressure on traditional applications to transform

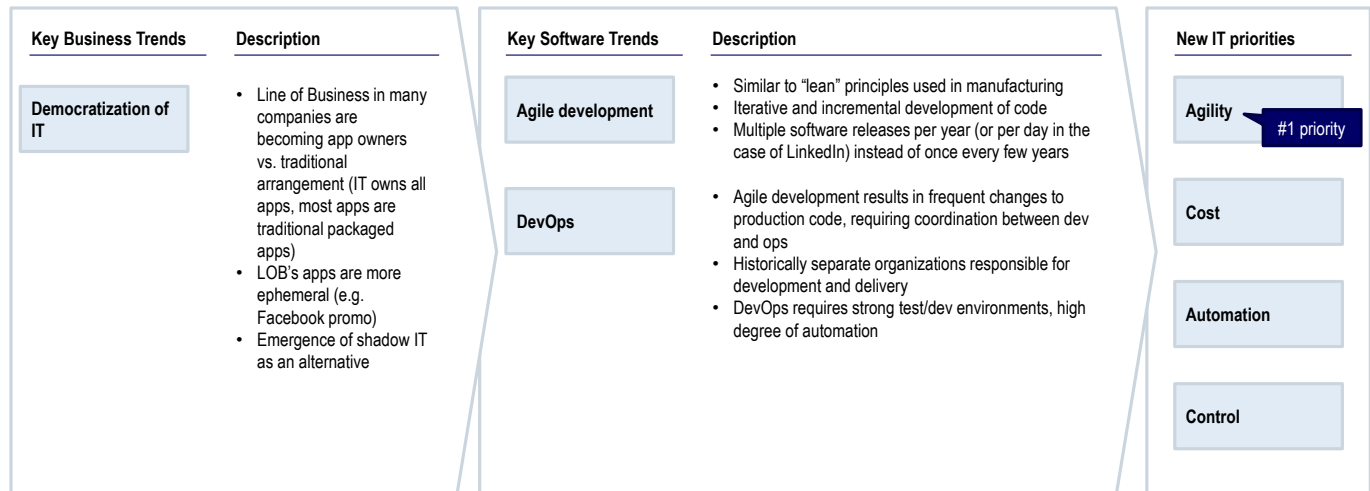


Source: Adapted from VMware analyst day presentation, August 2013

Agile development and DevOps go hand-in-hand with Cloud Computing

LOB ownership of software development is driving downstream changes in software development and deployment methodologies. SaaS companies have delivered more innovation in a shorter period of time than on-premise competitors and consequently have driven incremental adoption of software. However, SaaS vendors don't cover the multiplying range of applications that enterprise customers are beginning to demand. As a result, demand for Infrastructure as a Service (IaaS), raw cloud-based computing resources used as building blocks of modern applications, has exploded. This is the focus of our series of cloud reports.

Figure 8. Software trends are creating demand for cloud computing



Source: Citi Research

Debate #1: How Fast Will Applications Move to the Cloud?

Consensus view: Applications will quickly migrate to “cloud 3.0” architectures such as OpenStack and AWS

Citi view: Not that fast. Multiple decades.


Implication: If applications move at a relatively measured pace in enterprise, with IT budgets remaining under pressure, we expect VMware to be the biggest beneficiary. VMware’s technology enables enterprise customers to continue to reduce infrastructure CAPEX and operational costs.

- **“De-mistifying” cloud computing** – In this series of reports on cloud computing, we look at the potential of the Cloud to reshape the enterprise computing fabric. Enterprises now have at least two flavors to choose from when consuming computing infrastructure: vanilla “on prem” and new “as a Service” alternatives. There are even more flavors when considering hybrid clouds, and the many variants of IaaS being peddled by a hodgepodge of vendors including software companies, hardware vendors, telcos, and Amazon.com. We cut through the fluff and focus on what’s really changed: the emergence of the cloud management platform (CMP) delivering a high degree of automation and the resulting shift in the software development paradigm to take advantage of this new architecture.
- **The Not-So-Great Migration & The Great Transformation** – While it’s easy to imagine millions of workloads migrating across the firewall to reach the Cloud, the reality is that before any migration to the public or private cloud, applications need first to transform themselves to take full advantage of the Cloud. This Great Transformation is wrought with challenges, and we believe the current market view—that legacy vendors will be quickly displaced by the emergence of Cloud technologies such as OpenStack and AWS—is unfounded.
- **Watching ice melt: the application lifecycle and application demographics** – Current Cloud champions generally fall into two buckets: new companies that have exclusively developed on the Cloud or companies who consume computing resources in a way that are uniquely suited to the cloud (ephemeral workloads like genomic sequencing and “bursty” workloads like back-up). There is a large demographic of applications that don’t benefit much from the cloud, and this demo will likely live on given high switching costs in software. In addition, the application lifecycle of 20+ years promises a slow evolution toward the Cloud, except for app demos that skew cloud. Instead, we believe the majority of apps that will be built for the cloud are new apps, much the way that mainframe apps largely didn’t migrate to new architectures but new apps written went to new platforms.
- **Software stock implications** – Two of our names most exposed to The Great Migration are VMware and Red Hat (RHT.N; US\$58.08; 2). As highlighted in our [recent upgrade of VMware](#), the anxiety over the cloud threat may have peaked, and VMW should benefit from the market adopting the long view on the Cloud migration. Red Hat has invested wisely in Cloud 3.0 technology (OpenStack), but a slower migration relative to the consensus view would be less favorable for RHT. Microsoft provides a leading public/hybrid cloud, but its significant share in on-prem infrastructure could also be at risk. SaaS players with a Platform-as-a-Service like CRM (CRM.N; US\$63.73; 1) and NOW (NOW.N; US\$67.52; 1) could benefit from a slow migration as their PaaS could be an alternative to IaaS.

Options for migrating to the cloud

Enterprises face several options in dealing with their current set of applications, which today likely reside in a private data center in a Cloud 1.0-like environment (i.e. virtualized). We outlined the options that an enterprise faces for their existing applications in our OpenStack report, [Is OpenStack for Real and How Does It Impact the Software Landscape?](#). We summarize these options below.

Figure 9. Deployment options for new and existing apps



Source: Citi Research

Catalog of applications is diverse

Perhaps the more relevant question is not “how fast” but rather “which” applications will move to the cloud. The cloud has characteristics that are optimized for certain types of applications, such as “bursty” apps (e-commerce applications at Christmas peaks and other examples noted below). Some systems integrators perform a cloud diagnostic service for enterprises that buckets/triages their applications into varying degrees of cloud suitability. From there, enterprises might consider moving a small fraction of applications to the cloud.

Figure 10. Categories of Cloud-Friendly Apps

Application characteristics	Examples	Cloud advantage
Bursty	E-commerce (1-800-Flowers) TurboTax.com	Pay variable cost instead of capex for peak capacity
Ephemeral	Test Development	Quicker provisioning, cheap de-provisioning
Distributed access	Netflix Photo sharing	Existing cloud footprint
Stateless	Video transcoding Hadoop	No advantage. Cloud requirement in some cases.

Source: Citi Research

Costs of moving to the cloud are high

Traditional applications with none of these characteristics might ultimately never move to the cloud since the cost / benefit economics are less compelling. The cost of re-writing applications for anything other than Cloud 1.0 (which requires minimal if any work) can be very high and traditional enterprises have fewer resources trained to write applications on cloud platforms that are resilient and optimized for the environment of cattle-like servers. This developer expertise is scarce and mostly resides at web 2.0 companies that are using cloud 3.0 technology to build cutting-edge consumer apps (which require these capabilities). Some applications such as an ERP application might never move to the cloud in its current form, as the economics will not make sense.

In fact, further evidence of this is the fact that many SaaS companies remain in a Cloud 1.0 environment. Many of these early SaaS companies began their businesses with a hosting model first or before public cloud resources become commonly available. Products such as Taleo Enterprise Edition, ServiceNow, SAP's SuccessFactors, and many other SaaS companies were built in a Cloud 1.0 model utilizing traditional virtualization models, traditional relational databases, etc. The fact that these SaaS companies have not re-written for a next-generation cloud model underscores the difficulty of adopting a next-generation infrastructure.

Application lifecycles imply a 20-year migration cycle

If the cost of re-writing is greater than the benefit from cloud properties, then these applications are likely to remain on their existing Cloud 1.0 platforms, whether in a private cloud or hybrid cloud. However, there is a natural rate at which applications are replaced, and over time, we would expect that many of these applications will default into a new cloud-like environment as they are replaced, either with SaaS applications, applications built on PaaS, or (most likely) applications deployed on a cloud management platform like OpenStack. Ultimately some apps will never move, just as many enterprises rely on mainframe and Unix computing today.

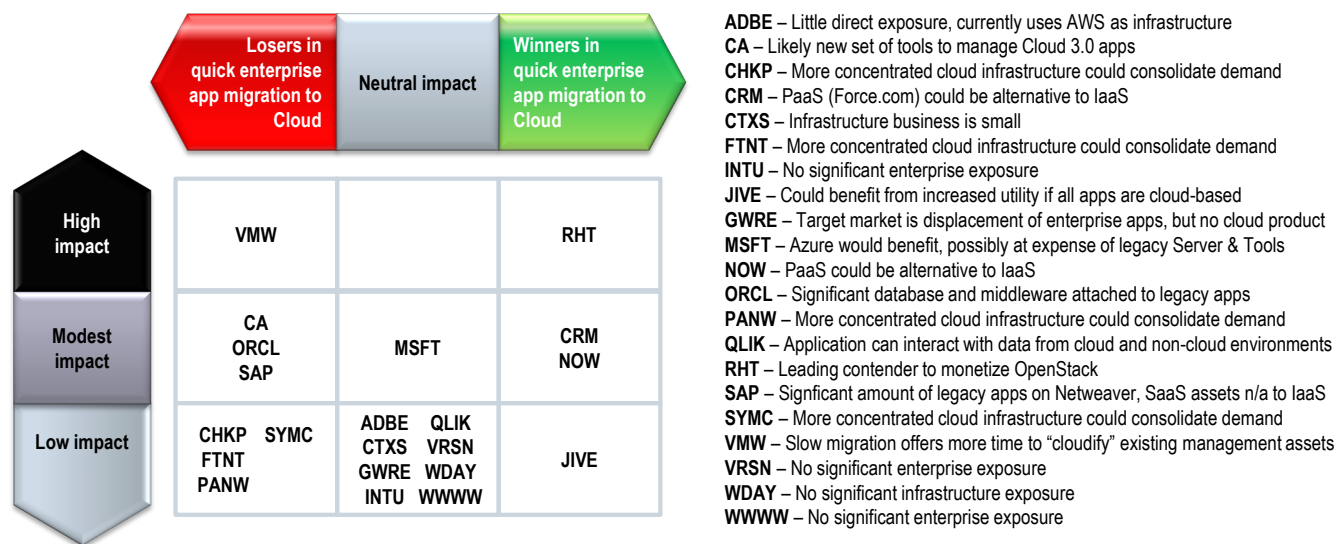
While we don't have comprehensive data here, we think that the move to the cloud is a 20-year exercise for enterprise applications. First, we think the move is similar to prior computing paradigm shifts (e.g. mainframe to client server). Second, in our research, we sampled some data from enterprises that shared their application inventories with us. Taking one example, a multinational company is retiring or replacing roughly 10% of its applications (both in-house and commercial) while adding 5% to its base for a net attrition of 5% per year as it rationalizes its application portfolio. We conclude that the natural replacement rate is somewhere in between 5-10% since in this example the attrition rate is not sustainable. Third, we note that on average, renewal rates of commercial enterprise applications (maintenance or subscription) are roughly 95% in the enterprise space. This implies an average customer lifecycle of 20 years (the inverse of the churn rate). However, there are some distortions, including reasons for churning that are unrelated to the lifespan of the application (e.g. bankruptcy) and the fact that the renewal rates are for commercial applications (excluding in-house apps).

If some apps don't move quickly to the cloud, who benefits?

We believe we have a non-consensus view on the lifecycle of applications. We believe the application replacement cycle is much longer. With the growth of AWS and the hype surrounding new cloud startups, cloud vendors are ready to give eulogies for Cloud 1.0 applications. We covered what the impact could be on the

software sector if Amazon ends up with a sizeable share of the enterprise infrastructure market (see [Software Sector Takeaways from AWS Re:invent - Headwinds for the Infrastructure Stocks](#)). While there are a category of applications that are moving quickly to the cloud (e.g. stateless/bursty apps such as media transcoding), we believe most applications within an enterprise will likely remain in Cloud 1.0 or Cloud 2.0 architecture.

Figure 11. Summary of the impact of app migration velocity on our coverage universe



Source: Citi Research

- VMware benefits from a long migration** – The friction for moving apps to the cloud is a benefit to VMware, which is building a toolset that allows enterprises to maintain Cloud 1.0 environments and evolve to Cloud 2.0 environments. Over the long-term, VMware will likely expose more of the raw infrastructure elements that are managed by vSphere, bring its own APIs and a hybrid private and public cloud deployment option to applications that are built in the Cloud 3.0 style. Key to a bullish view is 1) application migration to the cloud happening at a measured pace; 2) VMware is able to continue to add value to Cloud 1.0 application infrastructure and evolve to cloud 2.0 and 3) continue to develop a robust hybrid service (vCloud Hybrid Service – vCHS) which would enable customers to have seamless execution of applications, no matter where those applications “live” (even if this includes non-VMware infrastructure). We expect the company to continue to build out its vCloud suites and value around management and over the medium-term bring storage and networking virtualization to these applications in order to evolve to cloud 2.0. This combined infrastructure solution, which VMW has labeled the software defined data center, needs to remain the standard for legacy applications, hosting these apps as they remain on-prem in cloud 1.0 longer than some are expecting.
- Red Hat cuts both ways** – While Red Hat stands to gain from its efforts to control and monetize OpenStack, most of its current business is tied to enterprise Linux running in Cloud 1.0 environments. Unlike VMware, the company hasn’t had as much success adding additional value to its installed base. For example, we’ve seen minimal uptake of its RHEV virtualization tools (despite the fact KVM is doing well in early cloud deployments). As a result, Red Hat’s ability to monetize RHEV with its RHEV-M management tools has been lackluster. Similarly, JBoss and storage haven’t reached an inflection point in the way

VMware appears to be seeing with its vCloud Suites (although we do hear of improved recent JBoss traction). Similar to VMware though, if applications migrate relatively slowly to the cloud, the current Linux-based franchise is likely to have a long duration. If applications do migrate more quickly to the cloud, they might run on alternative distributions in public cloud providers (e.g. Amazon Linux) where Red Hat doesn't derive any support revenue

- **SaaS vendors with PaaS (CRM and NOW)** – Software-as-a-Service (SaaS) vendors with a Platform-as-a-Service (PaaS) offering such as salesforce.com(CRM.N; US\$63.73; 1) and ServiceNow would likely benefit from a quick evolution to Cloud 3.0. Specifically, they might benefit from 1) recognition by enterprises that building Cloud 3.0 is difficult, which might accelerate PaaS adoption over IaaS adoption as a way to minimize the effort in re-writing code, 2) broader uptake of cloud applications at large that would be complementary (and possibly integrated with) PaaS applications.
- **MSFT leading in enterprise public cloud, but with a lot to lose if not successful** –Microsoft (MSFT.O; US\$37.42; 2) was among the first to market with a set of enterprise cloud services (under Azure brand) built on Cloud 2.0 technology and has a degree of architectural symmetry with its on-prem family of Windows Server products and Hyper-V virtualization platform. Azure is both infrastructure as a service and platform as a service with early adoption being mostly IaaS based on our inputs. On one hand, this positions Microsoft well for the evolution towards the cloud and specifically, Microsoft's approach of enabling some cloud efficiencies (cloud 2.0-like) with minimal re-write of apps of change in developer technique is differentiated. This would mostly leverage the PaaS flavor of Azure. On the other hand, the majority of Microsoft's business today comes from on-prem infrastructure and over the last 10 years, VMware has replaced Microsoft as the hardware control plan under these applications. If customer look to consume IaaS, Azure is much less differentiated and likely has to contend with competition from AMZN and others who long-term will seek to become the platform on which developers build Cloud 2.0 / 3.0 application. We believe, like VMware, Microsoft benefits in a slow, evolutionary migration to the cloud, not a rapid one as much of the consensus view embeds. Under this evolutionary approach, Microsoft has time to complete architectural symmetry between the public cloud and on-prem, giving its customers a strong hybrid option.

Debate #2: Public vs. Private Cloud Prizefight

Consensus view: Public cloud and specifically AWS are huge winners

Citi view: No dominance in the medium term as enterprises experiment with Public and Hybrid clouds, but after some period where Hybrid Clouds bridge the transition, Public clouds win majority in the terminal state.

Stock impact: In software, VMW, RHT, and MSFT are most exposed with VMW a winner in the Private and Hybrid scenarios, MSFT well hedged, and RHT tied to the Private Cloud unless OpenStack is successful in the Public Cloud.


- **The Public vs. Private Cloud debate** – In Debate #1, we look at the debate about whether Public, Private, or Hybrid clouds will capture the majority of future cloud deployments. In one corner, we have the Public Cloud, most notably represented by current heavyweight champion of the skies, Amazon.com. In the other corner, we have traditional infrastructure companies such as VMware, which are rushing to cloudify their legacy solutions to create Private Clouds that pack the economic punch of Cloud 3.0. The Hybrid Cloud is singing kumbaya in front of the arena, asking why we can't all get along.
- **Give the fans what they want: Cloud 3.0 economics** – We look at Public vs. Private Cloud economics and conclude that for all of the scale that the Public Cloud enjoys, the economics are not always clear cut for the end customer. Clouding the cost equation are factors such as the margin earned by service providers, the mismatch between specialty workloads and commodity hardware, and forgotten costs such as bandwidth, which is independent of build economics. Some former public cloud sponsors such as Zynga are now favoring the Private Cloud.
- **Hybrid Cloud takes the early lead, Public Cloud goes the distance** – We expect this fight to go the distance in a drawn-out fight between the Public and Private Cloud. In early rounds, the Private Cloud has an advantage given application compatibility with all workloads. In middle rounds (10-year timeframe), the Hybrid Cloud will play an important role in bridging Cloud 2.0 workloads with the Public Cloud. In the end state, we expect the Public Cloud to take majority share given the benefits of scale, once the application lifecycle has run its course to take advantage of Cloud 3.0 architectures. Private Cloud will always have a place among customers, especially those who pay a premium for control over their infrastructure, but we believe the Public Cloud's cost and innovation edge will ultimately prevail.
- **Stock implications** – VMware, Red Hat, and Microsoft are the vendors most impacted by the public vs. private debate. The winners and losers here to some extent mirror the conclusions of Debate #1 (how fast do apps move to the cloud), as Private Cloud vendors are focused on bridging the gap between legacy applications and a fully evolved Cloud 3.0. VMware should benefit from the short-term focus on Private Cloud evolution, which we expect to evolve into a Hybrid Cloud, where VMware is again well positioned. Microsoft is well hedged, since it's built a cloud solution that is largely agnostic to the question of public vs. private. Red Hat's Linux business is at risk in the Public Cloud (as most public cloud providers don't use RHEL by default), so its future rests with OpenStack, which would help to drive enterprise OS demand if Public Cloud service providers are successful with OpenStack as the basis for their services.

Does cost matter?

As noted in Debate #1, the key driver of cloud adoption is agility, but a secondary driver is cost. We believe public clouds will likely always maintain an advantage on agility. With private clouds, agility is largely limited by the ability of the internal IT department to innovate, although it might benefit from outside innovation through the use of open-source technologies such as OpenStack. Public clouds have been pushed to innovate and have greater resources to innovate. Ultimately, we believe private clouds are subscale and will have difficulty replicating the economics of the public cloud.

Amazon is cheaper but more expensive


Figure 12. AWS is an order of magnitude larger than the field of IaaS competitors (2012 Cloud Revenue)



Source: Company Reports and Citi Research Estimates

Most private clouds are subscale. There is some minimum efficient scale (MES) in data centers as noted in the cost curve of data centers in the figure below. Costs continue to drop even after a data center has reached 100,000 nodes, suggesting that MES is very large and within the reach of webscale companies only such as Amazon.com. For example, Google (GOOG.O; US\$1,210.88; 1) – likely the largest data center operator – has approximately 1 million servers. The implications of the cost curve suggest that private clouds can never reach the scale of public clouds. So it is an open-and-shut case that public clouds will win both the agility and cost arguments. Or is it?

Figure 13. There are significant economies of scale in data center costs, and the shape of this curve is key to this debate



Source: Microsoft presentation

The public cloud has better economics

There is a persistent debate about the cost of public cloud services relative to internal IT. On the one hand, there are numerous examples of companies moving Amazon Web Services (AWS) workloads in-house once they reach scale, all under the banner of cost reduction. Zynga is a prominent example here. On the other hand, Amazon has enormous cost advantages ranging from sourcing servers and storage nodes directly from Taiwanese ODMs (bypassing Dell and others) and cheap power supply from nearby hydroelectric sources. In networking, AWS uses custom routers, custom protocol stacks, and custom switch firmware.


Figure 14. Sources of public clouds' cost advantages

Public Cloud Cost ADVANTAGE	Description
Procurement	Direct procurement of servers via ODMs. Buying power for power, networking, and other equipment.
Resource pooling	Smooth demand over a broader pool of applications (i.e. smooth over multiple peaks instead of separate peaks for Tax Day, Christmas, Valentine's Day).
Other economies of scale	Amortize fixed costs (e.g. custom networking protocol stacks, billing system, finance overhead, etc) over bigger revenue base.
Variable cost economics	Capacity build at web scale looks more like variable pricing than step-function costs.

Source: Citi Research

Public cloud providers such as Google and Amazon have been on the cutting edge of data center innovation. Mid-sized enterprises would appear hard-pressed to match the innovation and costs of these webscale datacenters. For example, AWS allows customers to sell their idle compute capacity in the spot market to further increase utilization. With public cloud vendors investing heavily in R&D, we presume they are seeing positive ROI and are getting further cost advantages beyond procurement. In the Figure below, we outline back-of-envelope math on the cost advantages enjoyed by a webscale public cloud. Layering in several basic cost advantages (primarily procurement), we come to the conclusion that public clouds enjoy a convincing cost advantage in economics – build economics, anyway.

Figure 15. Back-of-envelope math: Private clouds (even at scale) conservatively need to make up at least a 45% cost disadvantage vs. public clouds



Source: James Hamilton, VP at Amazon Web Services (mvdirona.com for Public cloud starting point), Citi Research analysis

Bridging the cost gap: Cost is not the same as price, and the price printed is not always the price you pay

On the face of it, public clouds should be cheaper, yet we find some public cloud customers are migrating to private or hybrid solutions in the name of cost or performance. Performance and cost are generally but not always the same thing in cloud computing. In the Figure below, we note several examples of customers migrating away from public cloud computing environments.

Figure 16. Examples of customers migrating workloads off the public cloud

Customer	From	To	Primary Rationale	Comment
Zynga	AWS	Hybrid Cloud (zCloud)	Cost & Performance	AWS cost advantage diminished as demand became more predictable. \$63M annual spend, moved 80%. Still on AWS for burst capacity. Says zCloud is 3x more efficient as zCloud is tuned to gaming workload.
Private SaaS company	AWS	Private Cloud (RackSpace OpenStack)	Performance & Cost	Zombie servers had a negative impact on performance and monthly bills.
MixPanel	RackSpace	Private Cloud (Softlayer)	Performance & Cost	Highly variable performance difficult to design around.
Sony	AWS	Public Cloud (RackSpace OpenStack)	Control	Wanted to reduce dependency on a single cloud provider.
Limelight	AWS	Private Cloud (Softlayer)	Control & Performance	Wanted control over downtime. Still using AWS for S3, DynamoDB, SimpleDB, Elastic Map Reduce, and EC2 for burst capacity. Multitenancy of AWS had negative impact on performance.

Source: Citi Research

The public cloud's cost advantages are relatively easy to understand (see left-hand side of the Figure below), but there other factors that favor the private cloud that can have an equalizing impact on cloud economics (see right-hand side of the Figure below). These offsetting cost disadvantages are the reason why the answer as to which is cheaper is not clear cut. If AWS cuts prices to achieve margins similar to its retail business, it's possible that the public cloud could be cheaper than almost any circumstance, except for the few application workloads that skew heavily toward a public cloud infrastructure (e.g. an irregular Hadoop job, test/dev workloads). However, Amazon does earn a margin, and for an enterprise with empty data centers and workloads with level compute characteristics (e.g. back-office transaction processing running 23 hours per day), public clouds could very well be more expensive.

Figure 17. Public clouds also have cost disadvantages

Public Cloud Cost DISADVANTAGE	Description
Margin requirement	Citi estimates AWS gross margins are in the 70% range. Fully burdened operating margins might be closer to 20-30%.
Peak capacity charging	For a workload with steady demand, there will be a public cloud premium for peak demand. This is mitigated by reserved instance pricing.
Sunk cost accounting	Enterprises with underutilized data centers have sunk, step-function fixed costs and could have lower incremental costs to build capacity in small quantity.
Lowest common denominator	Instances have generic performance characteristics that are not optimized for specialty workloads.
Variable performance	Since resources are opaque, customer can experience variable performance characteristics
Forgotten costs	While build costs might be similar when including the mark-up on services, bandwidth costs can be incremental in the public cloud


Source: Citi Research

The first three public cloud cost disadvantages are relatively easy to understand. Amazon needs to earn a margin, which could alone offset any procurement advantage. Our Internet analyst Mark May believes Amazon AWS's gross margins are well in excess of 70% while internal IT's transfer pricing likely embeds a slimmer margin if not zero margin. The next two factors are shaped by the utility nature of cloud computing and the public cloud's ability to best achieve cloud economics given the massive scale.

The last three factors likely require some more explanation:

- **Lowest common denominator and diversity in enterprise workloads** – As we noted in Debate #1, there is considerable diversity in enterprise workloads, and only a subset of applications are primed for the Cloud. Yet within this demographic, there is significant diversity. Many enterprises run on non-standard infrastructure configurations, including compute, storage, and network as well as custom software stacks. While companies have been pursuing a move toward Cloud 2.0, where they attempt to reduce diversity, few companies have fully made the transition. In the meantime, cloud vendors are broadening their SKUs to address this diversity in workloads. For example, AWS has introduced more specialized EC2 instances such as GPU instances. However, public cloud adoption will always be limited by this “long tail” requirement within the enterprise, and this “long tail” might be underestimated in terms of size. As we wrote in Debate #1, the refresh cycle is underway, and the “long tail” could migrate to the cloud, but this requires application re-write many times and with application lifecycles long (can be 10 years +), this “long tail” will take a long time to disappear.
- **Variable performance and shadow servers** – One of the key reasons is that public cloud performance can be variable. Cloud resources are always shared, and if you are sharing servers with a resource-intensive customer (i.e. the noisy neighbor problem), your performance will degrade. Sometimes. For example, heavy users know that instances provisioned in some data centers have better performance than those in other data centers. The end result is that customers might have to overprovision servers to get a minimum level of reliable performance and see some of the public cloud cost advantage negated by variable performance. Ultimately we think that the problem of variable performance is solvable. AWS has introduced Amazon EC2 Dedicated Instances, where hardware is dedicated to a single customer. Private clouds will face similar variable performance issues between internal “customers” when running at multitenant scale. We’ve heard from some service providers that VMware’s server virtualization technology continues to hold strong performance advantages, and it could be that some of Amazon’s performance disadvantage (equal to a cost disadvantage) is due to its use of the open source Xen hypervisor.
- **Forgotten costs such as bandwidth** – While costs to store a gigabyte might be lower in the public cloud, if the data is accessed frequently, bandwidth costs for transport of data to and from that public storage might eat up any public cloud cost savings. One infrastructure presents such data in the right-hand side of the Figure below and asserts that on-prem storage is cheaper than cloud storage, although clearly this non-whitebox vendor has a bias for showing on-prem to be cheaper. The hidden bandwidth costs especially matter in a hybrid cloud environment where an application transfers data across the firewall.

Figure 18. Network costs can eat up storage costs



Source: Infrastructure vendor presentation

It's a Control Thing

While we think that agility and cost are the two most important factors in infrastructure selection, some organizations value control most of all. For example, some industries such as financial services have a strong desire to control their infrastructure, whether to preserve competitive advantage (e.g. electronic trading) or to comply with regulatory requirements. In the section above, we listed several examples of customers who have moved off of the public cloud. While some moved off because of the variable performance issues discussed above, other such as Sony and Limelight moved off AWS to get greater control over their infrastructure. When service outages happen (a relatively frequent occurrence), IT departments are often powerless until service is restored. It is likely that AWS has better uptime than most enterprise data centers and this AWS downtime can be further minimized by optimally deploying workloads on AWS. However, during downtime, customers may feel like having more visibility into issues when workloads are in their own data center, rather than being without visibility and having to rely on a third-party to restore service. Amazon does sell premium support services at the price of 20% of your AWS bill, but even with that premium support, companies are still often at the mercy of Amazon fixing its own problems.

Is the Hybrid Cloud the answer, especially in the short run?

With all the pluses and minuses to public and private clouds, it seems logical to conclude that the hybrid cloud is the answer to the debate. After all, as we pointed out in Debate #1, there is significant diversity in applications, and some applications will likely stay in the private cloud until their demise. Hybrid clouds appear to bridge this gap, allowing enterprises to maintain applications in both worlds and maximize the benefit to their respective application demographics. In addition, the burst use case or the redundancy use case requires a hybrid cloud.

True Hybrid Clouds can be hard to build

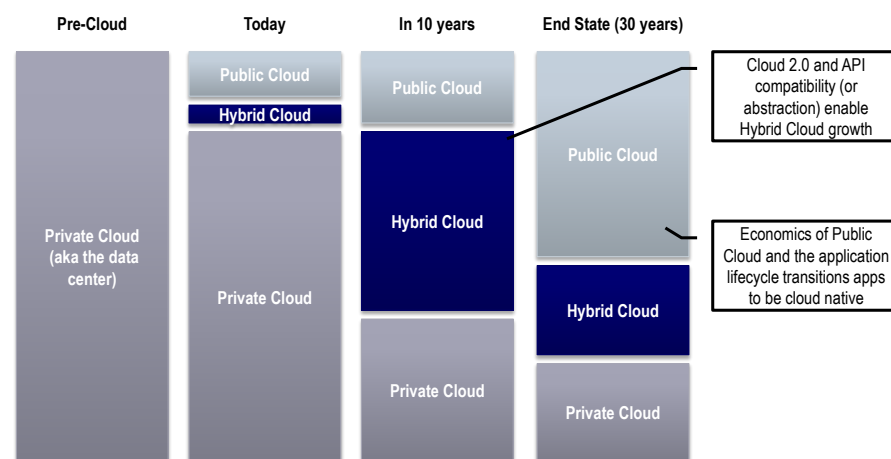
While it's likely that in the short run as applications begin their Great Migration to the cloud, companies utilize both public and private clouds. But using both public and

private doesn't mean that the enterprise is utilizing a hybrid cloud. A hybrid cloud has one of two technical features:

- **Symmetry:** For a workload to move seamlessly back and forth from public and private clouds there needs to be some form of symmetry – whether it is identical code or just application programming interface (API) compatibility. Microsoft promises this with its Windows Server / Azure combination and so does VMware with vSphere / vCHS. Also Eucalyptus brings Amazon APIs on prem. All of these efforts are early. We will discuss the issue of open APIs in Debate #3.
- **Abstraction layer:** Some tools such as RightScale and VMware's vCAC adds an abstraction layer beneath the workload to abstract the IaaS from the application. This allows the workload to move more easily to another cloud, whether private or public.

We believe few enterprises are operating true hybrid clouds today, as evidenced by 1) the lack of revenue for OpenStack-based service providers, 2) the majority of OpenStack customers using "native" APIs instead of Amazon APIs, and 3) Only fledgling demand for the public side of hybrid architectures from VMware (vCHS) and Microsoft (Azure). Over time, we expect hybrid adoption to accelerate as private clouds mature, abstraction layers mature, and, as we will discuss in Debate #3, API standards converge. To some degree, this will be helped by the maturation of public cloud services from Microsoft and VMware.

Figure 19. Citi view on Cloud evolution



Source: Citi Research

The Long March to the Public Cloud

Similar to our answer to Debate #1, we believe that the secular trend suggest workloads will migrate to the public cloud over time with hybrid clouds eventually occupying an important intermediate step. However, there are workloads and customers that will not move to public or hybrid clouds in the short run. The segment of customers/apps that will not move to the public cloud include:

- **Customers who require control.** For regulatory, strategic, or peace of mind reasons, some customers will prefer the higher degree of control afforded by the private cloud.

- **Workloads not suited for one size fits all.** Public utilities cater to the lowest common denominator in order to minimize SKUs and reduce complexity. Some customers (e.g. Zynga) will see benefits to moving workloads in-house. As enterprise reengineer their applications for Cloud 2.0, some of this diversity will go away over time.
- **Workloads that don't benefit from cloud economics.** As we discussed in Debate #1, some workloads benefit from cloud more than others. Workloads with predictable demand, especially those that are relatively mature, might never benefit from migration to the cloud.

The move to the public cloud mirrors the secular shift from Cloud 1.0 to Cloud 2.0 and eventually Cloud 3.0. While private clouds will likely get more advanced as projects such as OpenStack mature, we believe it is likely the enterprise market ends up with a healthy mix of public and private cloud deployments, with a secular movement toward public clouds due to their economic advantages and rapid speed of innovation.

Stock implications

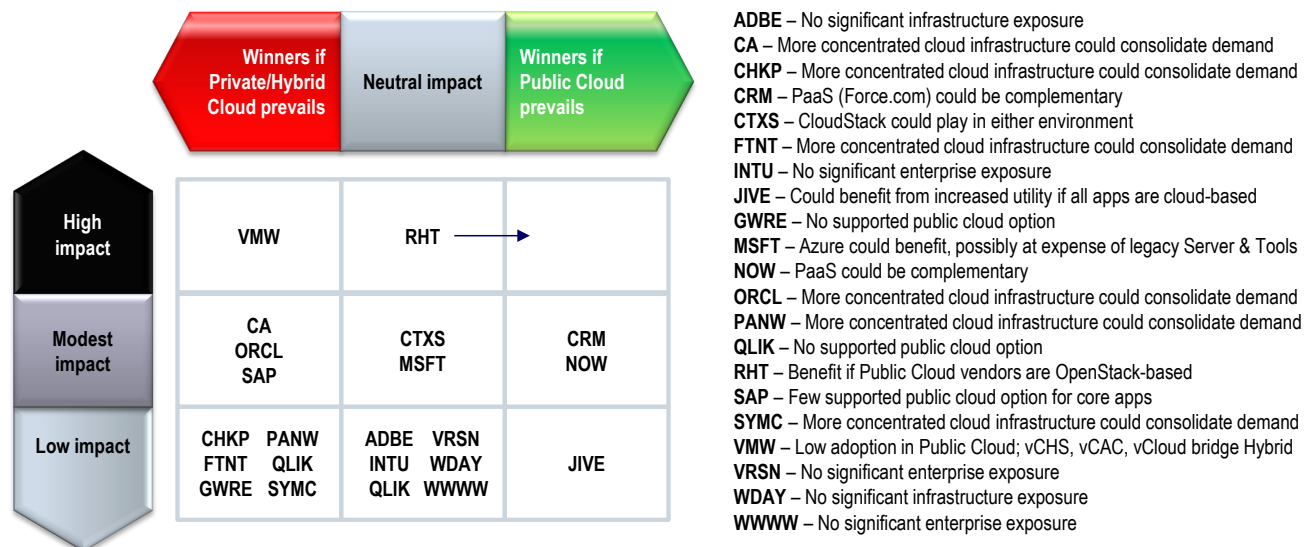
Companies most impacted by the transition from private to hybrid to public cloud are Microsoft, Red Hat, and VMware. Additionally, over time, we believe companies providing lower-level on-premise infrastructure elements such as security, back-up and certain IT management functions are likely to see some impact.

- **Microsoft** – In our view, Microsoft was the first large IT company to recognize the importance of having an architecture that had architectural symmetry between private and public deployment models. As a result, with the release of Windows Azure (Feb 2010), the company was early to market. The strategy has evolved significantly, but even four years later, the solution is still immature, especially in its key hybrid characteristic. Due to its early recognition of the hybrid scenario as important and its aggressive product roadmap here, we expect Microsoft to remain best positioned of the large IT vendors as it relates to “future proofing” its product line for the evolution from private to public cloud. Long-term, key for Microsoft is keeping developers committed to its various languages and overall framework. We will discuss this further in Debate #3. From an overall stock perspective, as we noted in our recent downgrade of MSFT shares from Buy to Neutral (see [Microsoft Corp. \(MSFT\) - The Decisions Ahead are Bigger than the Man; Downgrading to Neutral on High Expectations and Slow Expected Rate of Change](#)), we are concerned that a new CEO of Microsoft will continue to split investment between aggressively chasing the consumer devices market (where “success” has lower probability) and the enterprise cloud market (where “success” is more likely). By splitting investment and remaining un-focused, there is a greater chance that other players catch up with Microsoft in its early lead in enterprise public cloud.
- **Red Hat** – Red Hat's own public cloud initiatives are weak compared to others (Microsoft, even VMware). Red Hat's business is probably best-served long-term if enterprise workloads stay on-prem. Red Hat's public cloud strategy revolves around helping to establish OpenStack as the open standard for private and public clouds and then ride the wave of deployments (and especially the need for interoperability and support) by enterprise customers. As we noted in our report last year on OpenStack ([Is OpenStack for Real and How Does It Impact the Software Landscape?](#)), we continue to expect early traction will come from public cloud providers (Rackspace, HP's HP Cloud, telcos, etc.) as well as web companies (Yahoo, PayPal, Apple, etc.). These entities are less likely to pay Red

Hat for support for OpenStack, but if their adoption helps to add credibility to the initiative and bring interest from real enterprise customers (that are the mainstay of Red Hat's customer base), then this is progress for Red Hat. We expect these data points to build, but progress appears slower than a year ago. Meanwhile, Red Hat shares have recently appreciated significantly (up 2% YTD), with part of the thesis being that OpenStack (and other newer initiatives) will drive revenue re-acceleration at Red Hat in 2014. We believe the jury is still out on this. We'd rather play VMware instead on revenue re-acceleration from new products.

- **VMware** – VMware went from being the darling cloud play in 2011-2012 to a company investors worried significantly about secular position vis-à-vis public cloud in 2013. We believe the truth is somewhere between but over time, we believe VMware's public cloud position will improve and the company has the potential to become a very relevant players here. VMware's strategy looks very much like Microsoft's, except it was started later (vCHS was announced in 2013 vs. Azure in 2010) and also VMware targets lower-level cloud services (it is mainly IaaS – although that may evolve) while Windows Azure started off as Platform as a service (PaaS) and only released IaaS in the last 12-18 months. We believe vCHS will evolve rapidly over the next 12 months, as VMware has done a significant amount of hiring here and has an aggressive roadmap for new services. VMware's infrastructure is expensive relative to Windows and open source technology, and as a result, we believe the focus will need to be on driving efficiencies out of expensive, proprietary infrastructure (as happened with compute and is likely in storage and network). Similarly in cloud, we expect VMware to focus on more mission-critical workloads. This means VMware is less likely to have success with "born in the cloud" (i.e. "Cloud 3.0") applications and instead applications that can benefit from Cloud 1.0 and 2.0 architectures. If VMware finds itself competing against Amazon's AWS, it is likely going to find itself at a significant disadvantage. However, it is likely VMware can successfully provide a cloud service offering to rival that of IBM (IBM.N; US\$183.19; 1), Oracle (ORCL.K; US\$37.97; 1), HP and others that have traditionally focused on managed hosting and/or expensive, proprietary on-premise offerings whose value VMware can undercut with a software and automation-based solution.

Figure 20. Summary of the impact of public vs. private share on our coverage universe



Source: Citi Research

Debate #3: Will the Cloud be Proprietary or Open?

Consensus view: The Cloud will be based on open technologies as the technology industry coalesces around OpenStack and those not embracing open source technology will be disadvantaged.

Citi view: "Open-core" solutions might drive short-term innovation, but the mainstream CMP will be led by Amazon's reference model and pace of innovation. The gating item to AWS will be ability to embrace hybrid architecture as well as how quickly alternatives can emerge to eliminate AWS lock-in. We cover stock implications later on in this report.

■ **The Open vs. Proprietary debate** – In Debate #3, we look at whether tomorrow's Cloud will be built on open or proprietary technology. This is seemingly important as the OpenStack community is using lack of vendor lock-in as a key selling point. The reality is, key elements of the open source value prop are difficult to make happen and in the meantime, proprietary solutions like Amazon Web Services are building scale and driving network effects. Microsoft and Google are also following Amazon's lead. Given powerful commercial interests seeking to influence and potentially fracture open-source projects into open-core offspring (a Tower of Babel, so to speak), open source may become more about marketing. Inter-cloud compatibility among hardware, source code and APIs is difficult to achieve, making the open vs. proprietary debate less relevant in our view.

■ **Amazon is creating a proprietary chokepoint and benefitting from network effects, can it be challenged?** – Computing architecture shifts have had extremely valuable proprietary chokepoints (operating system, hypervisor, etc), and in Cloud, we see a similar potential for a chokepoint in the form of a "reference model" – a particular software configuration and behavior embedded in a cloud fabric. Given early lead in public cloud, we see AWS as the leading contender to have its reference model adopted as a de-facto standard. We are already seeing some start to emulate this model (OpenStack for example), reinforcing AWS' position. Challengers will likely exploit the following drivers: 1) AWS lacks a private cloud solution to complete hybrid scenario, 2) others are building critical mass, although significantly lag that of AWS, 3) AWS reference model appears able to be emulated, 4) there are migration and abstraction layers that enable some cloud portability.

There are multiple dimensions to this question, making open vs. proprietary a complex question to answer.

As we noted in the introduction to Debate #1, one of the main selling points of new cloud technologies is that enterprises can rid themselves of expensive, proprietary hardware. Cloud 3.0 technologies are built for the "cattle" model of computing, where it's assumed that inexpensive commodity components will fail and the cloud management platform (CMP) contains all the intelligence to keep applications humming. So it might be natural to assume that the Cloud will be based on open technologies and that proprietary platforms will face extinction.

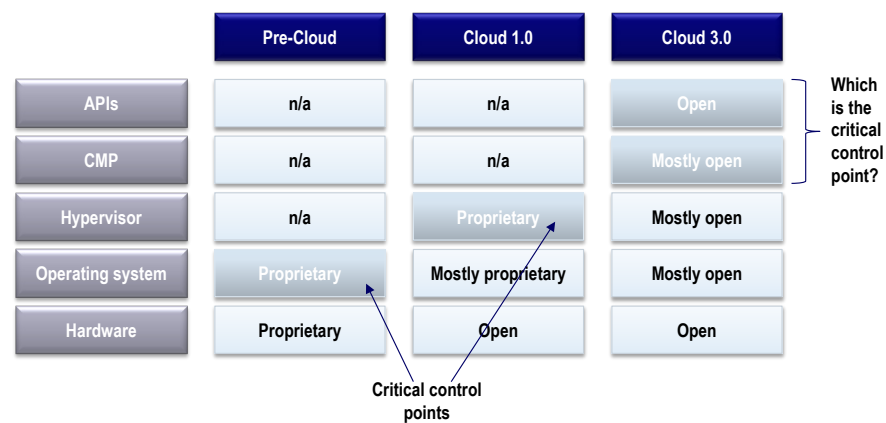
While that might be true of hardware, there is intense debate about whether the CMP will be proprietary or open. There are multiple dimensions to this question, making this a complex question to answer. Formerly proprietary CMPs such as Citrix's CloudStack have gone open source (at least a token open core) in an effort to accelerate development and adoption. The OpenStack community is debating whether to more broadly adopt Amazon APIs and the Amazon reference model. We break down the dimensions of this question in the report below.

Why open vs. proprietary matters

Control points shifting up the stack, layers below commoditizing and opening up.

In previous generations of enterprise computing, massive fortunes were created by owning a proprietary control point – whether it was Microsoft in server operating systems or VMware in Cloud 1.0 (i.e. server virtualization). Investors looking at the paradigm shift in enterprise computing are hoping to own a similar control point in the Cloud and to reap gains from proprietary software (see Figure 21). Similar battles in consumer tech have been fought (Mac vs. PC) and are being fought today (Android vs. iOS), and a common analogy is that Amazon Web Services is Apple, and OpenStack is Android/PCs.

Figure 21. Previous computing paradigms have had proprietary chokepoints



Source: Citi Research

What makes Cloud 3.0 an interesting puzzle is 1) the variety of players (web 2.0 companies, tech vendors, telcos, and software vendors), 2) the extent to which incumbent players have recognized the threat from Cloud 3.0 and responded (e.g. even “legacy” HP is building an OpenStack-based cloud), and 3) the uncertainty regarding where the critical control points reside. But whether we will see the next Microsoft or the next VMware, which have both established well defended “choke-points” in modern IT architectures, is less certain.

The multiple meanings of “open”

“Open” in software typically means “open source” software – all the underlying source code (which in software is all the intellectual property) is freely available for download to anyone for any use. Linux is the prototypical example in computer operating systems. Customers benefit from less vendor lock-in (they can modify the code themselves), while vendors benefit from potentially faster and less costly development if the open source project has sufficient momentum in the community. OpenStack (if you can’t tell by its name) is a CMP that hopes to follow this open-source model and has attracted significant critical mass.

Figure 22. The various meanings of “open” in the CMP

	Description	Customer benefits	Vendor benefits	CMP examples
Open Source	Complete source code freely available for any use	Lowest cost, control over code	Lowest cost of development	OpenNebula, Eucalyptus
Open Core	Core code is open but proprietary extensions atop core	Likely lower cost, faster innovation	Highest monetization potential	OpenStack, CloudStack
Open Interfaces	APIs documented & open communication with other platforms	Total infrastructure abstraction (perfect commodity)	Broader adoption	AWS (de facto standard)

Source: Citi Research

Open core has largely superseded open source and is creating a Tower of Babel

Speed is of the essence, so CMPs rushing to get critical mass will likely be open core

In reality, however, open-source software rarely remains the altruistic movement that it started as. Once commercial interests recognize the value of open source, the model usually evolves into “open core.” While Linux was once a completely open-source project, commercial vendors such as Red Hat have perfected the “open core” model, where a vendor develops proprietary code on top of a “core” of open-source code. True open-source models have become scarce (with some notable exceptions such as Apache), but there are examples of open source in the CMP layer – OpenNebula is the primary example, and Eucalyptus started as open source, went to open core, and is now reverting to open source. Citrix’s CloudStack has a token open-source core. But these rival cloud CMPs lack the momentum of OpenStack, where vendors with commercial interests are investing millions in developing the open-source core (plus proprietary extensions) in an attack that is trumping the ability of others to keep pace.

The myriad open-core OpenStack packages are creating a “Tower of Babel”-like environment, where even though these distributions share the same core, they are largely incompatible with one another. While a single-node application with basic storage and networking requirements might easily bridge this compatibility divide, a web 2.0 application running on an OpenStack would likely tap many proprietary extensions so as to render it incompatible with another OpenStack cloud.

Open source and open core questions are increasingly irrelevant in the long run.


While over the long run, open source might eventually subsume proprietary components built on an immature open core, the key debate is not over open source vs. open core but really whether open source (core or in full) really matters if the infrastructure runs “as a service” and the underlying guts are less important than how workloads are managed once up and running. While open source promises customers won’t get locked in, the reality is that as long as clouds have open and compatible interfaces (for example, conforming to Amazon APIs), in many instances the underlying code bases don’t need to have symmetry. The question then becomes whether open core/source still matters if they all support the same open interfaces (APIs) and therefore are compatible with all the major platforms anyway. Could it be that API compatibility is the lingua franca that essentially makes the underlying code irrelevant as applications embed a standard API?

Open source code is a red herring, but could APIs be the answer?

We conclude that the open source and open core questions are increasingly irrelevant in the cloud, and we focus on the interface layer. For enterprise use cases, we believe key ways to add value to the open stack (lower case) will rest in

1) the ability to bridge the public-private clouds, 2) the ability to help with Cloud 1.0 to Cloud 3.0 migrations, and 3) ability to achieve a minimum efficient scale to gain significant cost efficiencies.


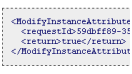

Figure 23. Enterprises want Multi-Cloud/Hybrid Cloud



Source: RightScale, State of the Cloud Report 2013

In Debate #2, we asserted that hybrid clouds will be an important driver in the transition to the public cloud. Given that CMPs are generally going open core, then compatibility is what matters. In a hybrid cloud environment, workloads can migrate between public and private clouds either because of symmetry (same code on both sides) or an abstraction layer (code that abstracts the application from differences in code or behavior, such as tools from RightScale). APIs are one kind of abstraction layer, where symmetry can happen. But there are other layers to target for symmetry, in increasing order of abstraction:

Figure 24. Potential areas for symmetry in the cloud stack

		Example	Software implications
Symmetry of ...	Hardware 	Open Compute Project	CMPs should be able to handle most of the heterogeneity of cloud compute fabrics, which are moving toward similar designs
	Source code <pre>from nova import utils from nova import volume LOG = logging.getLogger(__name__) get_notifier = functools.partial(wrap_exception = functools.partial(</pre>	OpenStack	Open core models destroy code symmetry as projects like OpenStack look more like a kernel than a full-blown open source stack
	APIs 	Eucalyptus	APIs are the first step in reference model integrity, but APIs alone do not capture full engineering specs
	Reference Model 	CloudScaling	Reference models are the more complete representation of symmetry that allows workloads to see symmetric environments

Source: Citi Research

- **Hardware symmetry** – Hardware specifications can drive operating characteristics, and initiatives like the Open Compute Project attempt to standardize reference designs to enable cloud-like commodity parts. While we acknowledge that symmetry here can drive efficiencies, we believe software can easily abstract many of the differences.
- **Source code symmetry** – The same exact source code exists in multiple clouds, and therefore all behavior is the same. This is the illusion that many OpenStack vendors are attempting perpetuate, but the reality is that nearly everyone is building an open core model and will do everything they can to avoid total

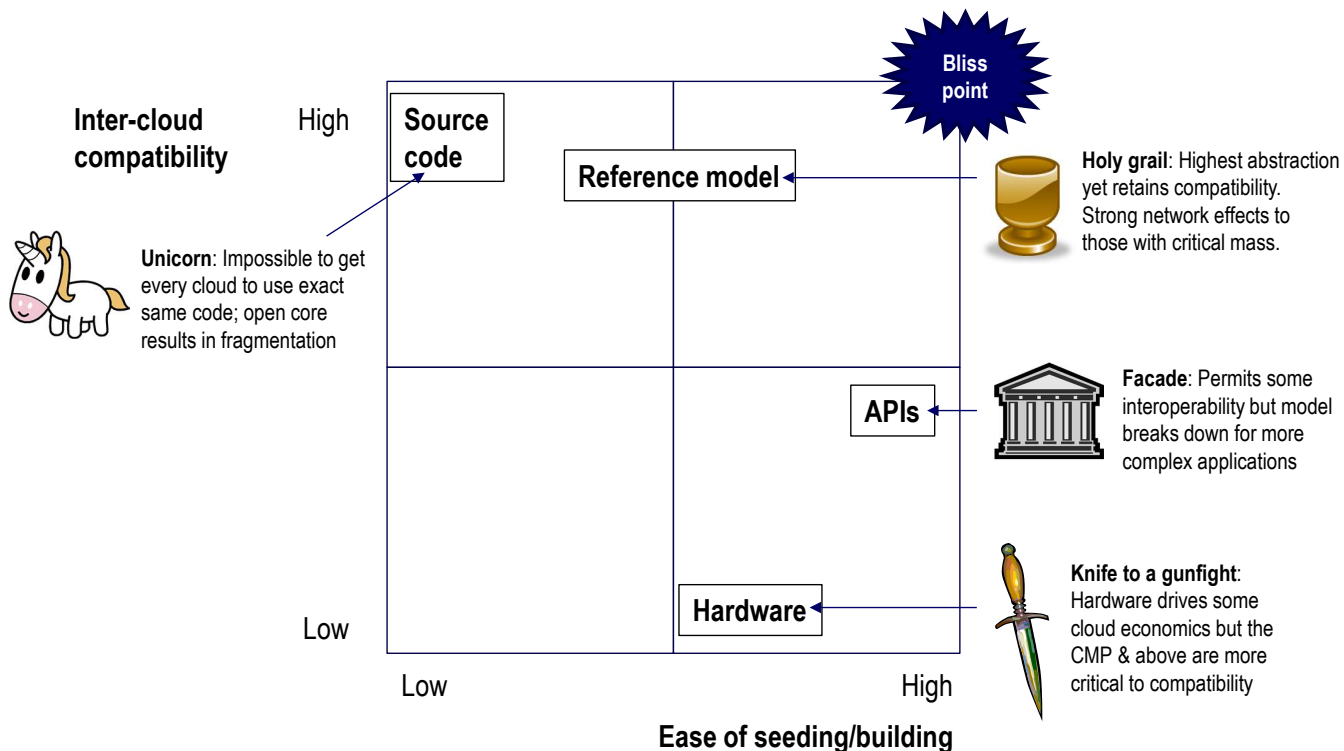
commoditization of the CMP layer. We note that even if it were possible to get to source code symmetry, hardware asymmetry could cloud compatibility.

- **API symmetry** – Symmetry of APIs is a necessary but not sufficient component to guarantee compatibility. While APIs are a higher level of abstraction that can potentially moot the open source, there are additional behaviors of a cloud environment (discussed below) that can break compatibility even if there is API symmetry.
- **Reference model symmetry** – In a higher level of abstraction, a reference model takes into account additional compatibility requirements that go beyond APIs including engineering specs and behavior (e.g. what happens when a hypervisor fails, how long should an app wait for an answer) that is undocumented in the API set.

API compatibility is also a red herring

API compatibility is necessary but not sufficient for compatibility. As mentioned before, an example of a Cloud 1.0 app such as, say, vanilla Wordpress running for a single tenant on a single node with very low traffic can likely survive migrations across API-compatible clouds in most cases. But for more complex requirements, this sort of compatibility breaks. In an anecdotal parallel to the hypervisor layer, a large enterprise that we spoke to would like to expand its use of KVM finds that it breaks many of its VMware workloads. Similarly APIs capture some but not all of the interoperability. Some vendors recognize the importance of compatibility and the insufficiency of APIs to get to compatibility. For example, Eucalyptus, a private CMP vendor, licensed Amazon's APIs to ensure better compatibility with Eucalyptus. Eucalyptus gets access to additional engineering specifications behind the API, which should in theory make Eucalyptus clouds work more seamlessly with AWS in a hybrid environment. Cloudscaling, an OpenStack integrator and packager of an open-core distribution, argues that this broader set of behaviors or "reference model" is a higher level of compatibility beyond the APIs. Their OpenStack-based technology is compatible with AWS at this level.

Figure 25. The Reference Model is the likely symmetry point to offer cloud compatibility



Source: Citi Research

Implications of a de facto standard

As we noted in Debate #1, applications will need to undergo significant transformation before they migrate to a Cloud 3.0-like environment, and we believe this transformation will take decades. In the meantime, customers are likely to adopt a variety of CMPs, consumed in a variety of delivery models (public, private, and hybrid). In this Tower of Babel of cloud deployments, it's unclear if/when/how these islands of clouds will come back together.

Amazon AWS has established itself as the largest provider of cloud services, with everyone from ISVs to web 2.0 companies to enterprises building on Amazon's infrastructure (although we are very early in enterprise adoption). AWS is becoming the de facto standard in the Cloud, and some rival CMPs are replicating the AWS APIs and reference model in an effort to facilitate hybrid cloud deployments. OpenStack, which has its own native APIs is enmeshed in an existential debate within its community about whether to fully adopt Amazon APIs, and even which API set – either the current Nova API developed by Rackspace (RAX.N; US\$34.97; Not Rated) or a new API set that is controlled by the community – to adopt as the “native” API for Nova, the compute fabric in OpenStack. Currently one-third of OpenStack deployments use Amazon's EC2 APIs.

There are rivals to AWS, including Google Compute Engine and Microsoft Azure, both of which take advantage of their webscale/hyperscale consumer services – following AWS's path. Additionally, like Microsoft, recent entrants including VMware (vCloud Hybrid Service) and IBM bring significant enterprise scale and a large customer base to bear on the cloud services market. While these rivals to AWS

matter because of their webscale consumer base, we believe there are two relevant questions that determine whether AWS will become the de-facto standard in public cloud services.

- 1) Is AWS lead is so significant in the market that strong network effects make it difficult for any competition to rival the de-facto standard position that its APIs and reference model have become? An affirmative answer here is most likely to drive AWS technology as the defacto open standard.
- 2) Do the public cloud needs of the enterprise customer evolve in a different way than have the requirements of customers for public cloud services have so far? Harkening back to Debate #1 and #2 of this report, if apps enterprise apps move more slowly to Cloud 3.0 architectures and the long-term intermediate state is more hybrid than public (even if eventually public), then some of the proprietary control points of some of today's on-prem architectures remain highly relevant and could give time for other reference models to catch up to AWS and displace it. However Cloud 3.0 apps are tied tightly to their underlying reference models, so we expect the reference model to become increasingly important and dominant.

Network effects = winner take all if the tracks are laid fresh

In addressing the first question, it is clear there are some network effects at play in the open interfaces. A common example of demand-side network effects is railroad gauges, where gauges became standardized in certain geographies to maximize the utility of rail cars on all rail networks. Rail cars could travel across previously incompatible networks, enabling transcontinental travel. The analogue in the cloud is that the application (rail car) is designed for a reference model (gauge). If the reference model is common among various clouds (rail networks), then the application can move seamlessly across clouds, and more companies will build more applications that conform to the most widely adopted reference model. Furthering the analogy, those clouds that don't conform to the Amazon reference model will be incompatible islands of computing. There are other benefits to Amazon's current scale – for example, there are increasingly pervasive software libraries that make it easy to store objects on S3, drawing incremental development.

What keeps AWS from becoming the next Wintel? (i.e. the de-facto standard)

As we note above, AWS has the early lead in cloud 3.0 architecture and is already being emulated by challenges that have yet to achieve this standard status. This lead, reinforced by network effects suggests that the Amazon reference model may have the run-away lead in a market the we have already concluded will inevitably move to public clouds for new applications. However, before we can conclude that the AWS reference model will be the Cloud proprietary control point, we need to believe the following:

1. **The AWS reference model must work in a private deployment.** For many use cases symmetry between public and private clouds will be critical. This will be necessary for workloads that “burst” as well as for many enterprise workloads that will likely exist in a permanent state of transition between private cloud and public cloud. We argue in Debate #2, the hybrid cloud is an important bridge for enterprises to get some cloud economics while they rewrite for Cloud 3.0. The AWS reference architecture assumes only public cloud infrastructure and many inherent properties of this architecture. Therefore, we believe the market (or Amazon) must prove the adaptability of the AWS

reference architecture for private clouds. The early steps here have begun with AWS partner Eucalyptus, and AWS continues to build some momentum. We note that Microsoft is pushing the state of the art on deployment symmetry between public and private with Azure and Windows Server 2013 and VMware is earlier but attempting the same thing between vSphere and vCHS.

2. **A more pure “open” cloud architecture do not rival scale of AWS in medium-term.** Amazon has an early lead and if this lead is not challenged by a rival “open” architecture, Amazon will be able to undercut its competition on price and also disproportionate mindshare with developers. OpenStack is the clearest challenger but does not have reference models and there is no single OpenStack player that has the scale to build critical mass for one today, but this could change very quickly. In theory, if an alternative to AWS emerged that was “more open” (i.e. a “unicorn” from Figure 25), multiple constituents in the cloud ecosystem, including developers and enterprise customers would likely prefer this to AWS’s de-facto standard position (and proprietary control point).
3. **Other vendors cannot replicate the AWS reference model.** While OpenStack’s original APIs were the same as Amazon’s EC2 and numerous CMPs support AWS APIs, APIs are necessary but not sufficient for compatibility. AWS APIs are publicly documented, but the engineering specs are not all documented and even the ones that are undocumented and known (e.g. the number of network interface cards per VM) are not uniform across OpenStack implementations.
4. **Abstraction layers can’t bridge the interoperability gap.** Products such as RightScale and VMware’s vCAC (formerly DynamicOps) can abstract some differences in reference models, especially Cloud 1.0 apps where the apps are not tied as much to the cloud environment, albeit at some cost to performance. Webscale companies are unlikely to take this performance hit, but enterprises who are more interested in price and negotiating leverage might find it worth their while. Also, there are migration platforms such as Rivermeadow that enable some portability between cloud architectures, although at a very basic level today. If differences in APIs and reference models can be abstracted away or applications easily moved, then there is no proprietary control point. At the same time, this lack of proprietary control point would then apply to others looking for a control point (like Microsoft and VMware).

Stock implications

In each of the first two debates, we outlined implications from a stock perspective from the dynamics discussed in the report. In this debate, we are doing the same, based on how “open” cloud architecture evolves to be. While each vendor has multiple points of exposure to cloud symmetry, we note that essentially the key battle is between AWS, OpenStack, and VMware. Red Hat and VMware are the two key vendors in software, each having considerable control over a number open and proprietary control points. In the reference model area, there is AWS and a plethora of vendors hoping to either mimic AWS or make difference irrelevant by abstraction or targeting workloads that might not be reliant on reference models (particularly Cloud 1.0 and Cloud 2.0 workloads).

Figure 26. Winners in the Open vs. Proprietary points of symmetry

	Hardware	Source code	APIs	Reference model
Open	OCP supporters & whitebox vendors	RHT	RHT, Eucalyptus	Model replicators & Webscale cloud vendors, incl. MSFT
Proprietary	ORCL	VMW, CTXS	Abstraction layers (e.g. VMW vCAC, RightScale)	VMW

Source: Citi Research

Overall Cloud Formation Stock Conclusions

■ **Grand finale stock implications** – Following the three debate sections, we take the following view on Infrastructure as a Service (IaaS) impact on stocks in our universe. VMware is a winner if enterprise migration to cloud architectures happens in an evolutionary fashion, with VMware increasingly becoming the data center control plane for mission critical applications. So far, early indications around public cloud services (vCHS) are positive and if the company can capitalize on storage and network, it is in a strong position into the cloud transition. Microsoft has is one of three players with web-scale public cloud service. However, we believe with a Windows-only focus, it will be difficult for Microsoft to complete broadly in IaaS and instead must be successful reinvigorating its developers around the Azure platform as a service (PaaS) offerings. Red Hat has firmly made a commitment to OpenStack (OS) and is at the mercy of the timing of maturity of OS, something that could still be 18-24 months away. We continue to believe service provider adoption is the harbinger of enterprise demand. Oracle is behind in public cloud and while its workloads are amongst the stickiest, the company must close the gap as enterprises become more comfortable with IaaS.

What you have to believe for each to be a winner

We believe the move to cloud architectures is a more complex change than most investors believe. Where we don't disagree is that this move will have a profound impact on the software landscape and on the stocks of companies involved. The fact that so many questions around the move to the cloud are unanswered at this early stage in the market transition makes it difficult to make sweeping and certain stock conclusions. As a result, we have outlined where each company participates in the Cloud and where each company is aligned vis à vis the 3 key questions posed in these Cloud Formation reports.

Figure 27. Software universe exposure to IaaS

Company	Participation in cloud layers				Position in Cloud			Position in			Alignment with			Key things to watch
	infra.	IaaS	PaaS	SaaS	1.0	2.0	3.0	Private	Hybrid	Public	AWS	Open-Stack	Own prop.	
VMW	●	●	○	○	●	●	●	●	●	●	●	○	○	Software defined NW / storage and vCHS traction
MSFT	●	●	●	●	●	●	●	●	●	●	○	○	●	Developer traction with Azure
CTXS	●	○	○	○	○	○	○	○	○	○	○	○	○	Cloud Stack adoption
RHT	●	○	○	○	○	○	○	○	○	○	○	○	○	Service provider traction with Open Stack
ORCL	○	○	○	○	○	○	○	○	○	○	○	○	○	Evolution of IaaS and PaaS, Nebula traction
SAP	○	○	○	○	○	○	○	○	○	○	○	○	○	Development of platform services on Hana
CRM	○	○	○	○										As apps get replatformed, PaaS is an alternative
NOW	○	○	○	○										
WDAY	○	○	○	○										Need partnerships on public cloud to extend reach
WWW	○	○	○	○										
INTU	○	○	○	○										
JIVE	○	○	○	○										
ADBE	○	○	○	○										
ADSK	○	○	○	○										
GWRE	○	○	○	○										
CA	○	○	○	○										
SYMC	○	○	○	○										
CHKP	○	○	○	○										
FTNT	○	○	○	○										
PANW	○	○	○	○										
VRSN	○	○	○	○										
QLIK	○	○	○	○										

Source: Citi Research

VMware – Control of Mission Critical Workloads and a differentiated IaaS offering are key

VMware's entry into public cloud is early and behind on AWS definition of market

VMware is well positioned in Cloud 1.0, Cloud 2.0, and the Hybrid cloud.

VMware is behind in the hybrid/public cloud market, having debated a strategy internally for years based on our research, but only releasing its vCloud Hybrid Service in May 2013. We understand that in the last 6 months, vCHS has seen an up-tick in traction, although customer counts are still low (one of the more successful vCHS channel partners has “dozens” of customers). We believe traction so far is in test/dev and disaster recovery.

The promise of vCHS is architectural symmetry with on-prem vSphere deployments and based on what we hear from early deployments, the simple ability to use vCenter and other management tools to control vCHS in the same way is driving the early success. We expect this symmetry will continue to expand in value if VMware is successful in expanding its footprint beyond compute and into network and storage with the broader “software defined data center” strategy. We don't expect that vCHS will ever be the cheapest or most open IaaS offering on the market.

At the same time, with the majority of cloud adoption being from web 2.0 and similar companies, we believe VMware will appear to be behind for some time in this market based on revenue scale. We expect VMware will be tempted to release AWS-copycat features such as hourly billing and self-service access, although these are not likely to be the features that its enterprise customers value. Instead, we believe that success for VMware will be to bring a hybrid cloud solution to high-value on-premise workloads. In bridging to public, VMware needs to initially bring the best “pets” (from our earlier “pets vs. cattle” analogy) architecture, something that Amazon doesn't do, even though the long-term future of the cloud will be a “cattle” model.

VMware is in a tough spot if enterprise moves quickly to cloud

If enterprise customers re-architect their applications quickly to Cloud 3.0 architecture (i.e. “cattle”), VMware's window of opportunity will be minimized as the world will move through from private to hybrid and into public quickly. In this scenario, the customer would not need the value provided by VMware's management of underlying hardware resources and the benefits of architectural symmetry between private and public.

More likely case is slow migration to cloud, which favors VMware approach

In our mind, the more likely case is that there are going to be enterprise workloads that stay in the current architecture (Cloud 1.0 – server virtualization) for some time and will need optimization of the entire infrastructure stack to lower costs as they transition to Cloud 2.0 (automated infrastructure). What VMware is doing around virtualization of compute, storage and network is ahead of the competition, in some cases by a substantial margin. This should continue to drive share gains in mission critical workloads onto vSphere (VMware compute platform) and vCloud (VMware automation platform). With a slower migration to cloud, VMware can expose APIs that enable data center operators to have a greater degree of control over underlying infrastructure elements. Ultimately these APIs will be much like what Amazon exposes, but done in a way with continuous backward compatibility to prior enterprise process and governance.

Pay attention to mission-critical workloads and success with network and storage as harbingers of cloud success

We believe what is important to watch for VMware is whether vSphere continues to onboard mission critical x86 workloads (IBM Websphere / Oracle WebLogic), Oracle database, SAP applications for example) and how VMware is able to impact the way storage and networks are managed. We believe VMware has a stronger presence than most investors believe in mission critical workloads already (we've seen some stats suggesting well over half of middleware instance run on VMware). If VMware hosts mission critical x86 workloads (where much of the IT budget is) and becomes the control plan for storage and network within this infrastructure, then the company is in a strong position going forward, almost no matter what the pace is of adoption of public cloud (where VMware is behind).

Success may never be in PaaS and new custom app building

Before the Pivotal spinout, VMware was assembling an entire stack, much like what Microsoft has built to enable developers to build natively on the platform and then run on infrastructure optimized for this native development. While this vision was elegant, it defocused VMware from focus on core infrastructure and arguable set VMware back competitively during 2010 and 2011. With the Pivotal assets no longer tightly integrated with VMware, we expect VMware's success in cloud is likely to come more at the infrastructure level (IaaS) and in this vein, with a differentiate infrastructure (distinct from AWS commodity infrastructure). We wouldn't expect VMware to come to market with a generic PaaS or generally application level services (open source database for example). Instead, we expect to see VMware work on building the broadest cloud 2.0 platform for hosting common, mission critical workloads. We expect there might be a flavor of the service that optimized for middleware and can host Websphere (IBM) and WebLogic (ORCL) on VMware, both on premise with vSphere and in the cloud with vCHS. Thus we look for the roll-out of these types of services in 2H14.

What about VMware and OpenStack?

We noted in [our report on OpenStack](#) about a year ago that OpenStack's success was VMware's loss. We still believe this in a pure sense (a built from scratch "new cloud" that is 100% OpenStack may be no place for VMware), but over the last year, it is clear to us that OpenStack is likely to be a looser "standard" the original envisioned and there is likely to be a significant amount of mix and match with componentry (as always is the case in enterprise IT). For enterprise customers that are already very comfortable with vSphere as the underlying compute in the data center, it is possible that the OpenStack (or CloudStack) cloud management platform (CMP) could spin up, spin-down, schedule and control computing. VMware continues to have a level of maturity in its hypervisor that brings better performance, broader hardware and tooling support and more institutional knowledge within customers on how to manage it. The more fragmented and bogged down OpenStack becomes, the more likely it is something that plays a looser role within enterprise IT and the more strategic the position of VMware in these accounts.

Stock view – We believe investors are too negative around VMware cloud position

A driver of multiple contraction for VMware in 2013 was the notion that AWS and OpenStack would quickly emerge as significant competitive threats to VMware. We believe this threat is less imminent than the consensus view. Furthermore, with a VMware focused around infrastructure, we believe the company is in a good

position to leave the transition from cloud 1.0 to 2.0 with on-prem infrastructure and provide a hybrid alternative with vCHS. We believe the 1.0 to 2.0 transition, where VMware will sell vCloud management tools as well as infrastructure software around networking and storage will drive 15%+ revenue growth through FY16. This is ahead of the street and we believe will result in multiple expansion, driving upside to shares.

Microsoft – Symmetry for Windows and web-scale public cloud offering is the Hope

Microsoft is locked into Windows for hybrid value prop

We noted above that VMware is likely to bring symmetry at the infrastructure level while providing infrastructure to run mission critical applications (whether they be Windows or Linux apps). We expect Microsoft will provide infrastructure that has symmetry of deployment between on prem (Windows Server) and the public cloud (Azure). We expect this to be limited to Windows applications as there is no reason enterprise customers will run Linux workloads on top of Windows servers. This limitation to Windows is likely to be an advantage and disadvantage. Amazon currently offers Windows instances on AWS and VMware is likely to do the same. We expect Microsoft to be the only one that will offer native Windows APIs on the public cloud (Cloud 3.0-like), in contrast to VMware and AWS offering Windows virtual machines (Cloud 1.0 and 2.0). Developers that design for Windows on Azure will likely be able to benefit from Windows properties in Azure that don't exist in Windows on AWS and VMware vCHS as Microsoft's capabilities will be embedded in developer tools (Visual Studio) that these other vendors don't have. Microsoft's hybrid capabilities will be superior to any one else for the Windows environment and since the 5-6M Windows developers aren't going anywhere, any time soon, this will be a good business for some time to come.

Microsoft needs to re-invigorate the developer base around cloud

We believe the 5-6M developers that Microsoft has are a significant asset, in fact after the hundreds of millions of information worker using Office products, this is Microsoft's biggest asset. These developers grew up building Windows apps, based on the windows common language runtime (CLR) which is common across both the client and server operating systems. Arguably, this is the largest, most loyal base of developers and one that has evolved over time with changes in Microsoft development frameworks. As a result, this developer base should carry Microsoft into the cloud era, with these developers picking up the skill set to build natively on Azure. This rich set of developer tools and developer loyalty, combined with the above-noted limitations to Microsoft infrastructure make Microsoft's most likely differentiated play as the Azure Platform as a Service (Windows Azure, SQL Azure, etc.). We expect here, Microsoft will effectively transition its existing developer base.

The long-term question will be whether or not Microsoft can attract new developers to the platform for "starting from scratch" apps. Here, developers will compare the merits of the Windows APIs vs. other cloud frameworks. If these apps are at all hybrid (communicating with on-prem Microsoft apps) or leverage Microsoft services such as Skype, Skydrive, Bing search, etc, Microsoft will have a leg up. If not, Microsoft will have to compete on equal footing with AWS.

Azure will play to non-Windows developers, but that sounds like a tough battle to win as long as Windows is the strategic focus

If native Windows coding is not part of the spec, Microsoft's Azure offering will have to compete on an even playing field with AWS and other IaaS. Also, AWS has built early forms of PaaS such as Elastic Beanstalk and RDS. AWS and Azure will overlap in both IaaS as well as PaaS around non-Microsoft languages and frameworks. A good example is Node.js (server-side Javascript), which both support natively. We believe here, AWS has the scale and the lead with developers, something that Microsoft will have to work hard to rival. We therefore believe that Microsoft is most likely to win where it can leverage native Windows development technologies and in customers that are looking for symmetry of development and deployment of these Windows apps.

Stock view – Best positioned mega cap for cloud, still with Windows baggage

With a new CEO at Microsoft and investors awaiting an understanding of how Satya Nadella may approach running the company differently than Steve Ballmer, we believe the ins and outs of Azure don't matter that much right now to the stock. Long-term, Microsoft is the best mega cap in terms of positioning for trends towards cloud computing. However, we believe that Microsoft's split personality, between wanting to continue to compete for market share in all relevant device categories and also provide applications and services across all devices is conflicts. We don't believe Microsoft can maximize success while trying to do both. As a result, we believe investors stand to be disappointed when the go-forward plan at Microsoft looks very similar to the current plan. We expect cloud to continue to get airtime amongst investors but at the same time, we believe there will be bigger investment debates around Microsoft shares in 2014.

Red Hat – Is OpenStack the answer to RHEL cannibalization by the public cloud?

Big investment in OpenStack

Success for Red Hat rests on success for OpenStack.

Red Hat has decided to invest heavily in OpenStack in a bid to take a similar position in OpenStack as it has with Linux – becoming the vendor of choice for enterprises or service providers wanting commercial support. While we believe this is a timely and strategically savvy bet, success for Red Hat in the CMP depends on OpenStack's success and there is also the question of timing. As we highlight in Debate #3, source code is unlikely to become the point of symmetry and on paper, this creates less potential for value creation for Red Hat. However, Red Hat has earned a position of trust with its enterprise customers and the packaging up and supporting of OpenStack (even if source code symmetry is not critical) can be a significant business. Also, with so many service providers and vendors adopting OpenStack, there is some likelihood that one of the winners (be it HP Cloud, RackSpace, Telefonica, etc.) will run OpenStack and pay Red Hat for support and proprietary extensions. We therefore see Red Hat has having covered its bases and putting itself in a position to succeed in public, hybrid, and private clouds (assuming OpenStack is successful) and in Cloud 2.0 and 3.0 environments.

Can RHEL make the transition?

While most of our discussion relates to the CMP layer and who can grab the lead there, trends in IaaS can have a significant impact on Red Hat's current core business, Red Hat Enterprise Linux (RHEL). RHEL wasn't widely available in the public cloud a couple of years ago due to licensing practices, but now public cloud vendors do offer RHEL although at a premium price that has dissuaded early

adopter customers such as Web 2.0 start-ups from deploying their applications on RHEL. Meanwhile, Red Hat's KVM virtualization is being adopted quite broadly by public and hybrid cloud vendors primarily due to zero licensing costs and KVM's efficient architecture. We are uncertain whether Red Hat can monetize either of its OS or virtualization assets in a Cloud 3.0 environment, given the price pressures involved with commodity hardware and the fact that the Linux operating system has less value than underlying cloud APIs. To the extent enterprises embrace a hybrid cloud and/or a slow move to Cloud 3.0, we expect Red Hat to be cushioned from an immediate impact to the core business.

Stock view – Best way to play OpenStack but unlikely to be the medium-term growth driver

Red Hat has been seeking new growth businesses to drive the next phase of growth after the tapering of RHEL following past catalysts such as accelerated Unix migrations. While on its most recent earnings call, Red Hat announced some six-figure deals for OpenStack, we believe it is still early days for RHT to drive meaningful growth from supporting OpenStack. We have not found significant traction for enterprise or service provider deployments in any of our conversations with cloud vendors outside of a handful of early adopters. In addition, given our longer-term view that the core OpenStack project could struggle with competing reference platforms and the splintering that follows an open core model, we suspect that OpenStack is not the catalyst that Red Hat needs to reinvigorate growth in the next 2 years.

Oracle – Secular loser in cloud unless pace of innovation picks up

No significant exposure to IaaS infrastructure software

Historically, the bulk of Oracle's infrastructure software resides above the CMP layer in the form of application middleware and databases. There are some small exceptions – with the acquisition of Sun, Oracle added the Solaris operating system and added to its server virtualization capabilities. In addition, Oracle has identity management software and management software (Oracle Enterprise Manager). The small size and low (or lack of) growth of these assets suggest de minimis impact from the shift to cloud.

Application middleware and database in theory under threat from Cloud 2.0 / Cloud 3.0, but these may be sticky for some time

Oracle database and middleware are, in theory, at risk from a move to Cloud 2.0 and especially Cloud 3.0 architectures. In Cloud 3.0, applications are re-written to natively take advantage of properties inherent in the application architecture. Many or even most applications running on an Oracle infrastructure are platformed there to take advantage of scale up resiliency and properties. While these applications could ultimately be re-platformed, this process may take some time. There is significant debate around what percentage of the apps running on Oracle infrastructure are "stuck" there, with replatforming very difficult and how many can / will move. Oracle also has strict licensing practices that are high barriers for customers looking to move to the Cloud. Oracle requires that customers pay Oracle based on the number of processors – whether physical or virtual – a licensing practice that could offset savings from a cloud migration. While licensing increases migration costs and helps to keep customers loyal, at the same time it increases the incentives for customers to replatform to non-Oracle infrastructure or alternatively just bypass Oracle to adopt new SaaS applications or build their own Cloud 3.0-style applications.

Primary cloud strategy is a focus on appliances

Oracle is laser-focused on selling Exadata and other appliances to cloud vendors. While this strategy is largely in contrast to the embrace of commodity hardware by virtually every CMP, we believe Oracle can carve out a significant niche for itself in its large base of legacy applications written on Java, Solaris, and WebLogic technologies. We have no doubt that enterprises can achieve high consolidation ratios with Oracle iron and there drive strong efficiencies with what are effectively Cloud 1.0 technologies.

Dark horse in Nimbula/OpenStack, which opens up Oracle appliances

In March 2013, Oracle acquired start-up Nimbula, whose flagship product Nimbula Director was a CMP designed for the hybrid cloud. Nimbula's founders boast the best pedigree – they created Amazon Web Services. Nimbula started as a proprietary CMP that was building OpenStack compatibility. In September 2013, Oracle revealed the role of Nimbula, which will be a software layer on top of Exalogic that supports OpenStack APIs to create an OpenStack cloud in a box. In December 2013, Oracle joined OpenStack and announced that Solaris 11.2 will be compatible with OpenStack, which it now views as very strategic. While we believe the Nimbula/OpenStack solution is unlikely to present any cost savings in comparison to open solutions and therefore unlikely to capture new workloads, there might be a niche for customers seeking high-performance OLTP applications in an elastic cloud environment (e.g. test/dev for Oracle applications).

Stock view – Mostly negative implications from Cloud with a dark horse

As with SaaS, implications for Oracle are largely negative. Oracle has benefitted from selling expensive proprietary solutions, and the cloud is the antithesis of that. Oracle's large base of legacy applications is exposed to the multiple threats of SaaS, Cloud 2.0 replatforming, and Cloud 3.0 rewriting. The debate will be around time and if this timeframe for disruption from Cloud 3.0 is beyond that of SaaS (which it likely is), we expect SaaS threat to remain the focus. While Oracle is creating an IaaS and PaaS strategy with Exadata/Exalogic/Nimbula/OpenStack, we do not believe that these solutions cater to the core cloud audience, which is looking to drive significant cost savings.

Citrix

Cloud is a relatively small part of Citrix business, with its CloudStack based offering a business that we estimate in the range of 1% of total revenue. This is primarily from service provider partners (like Korea Telecom) and to a lesser degree enterprise customers, which are using this as the basis for Cloud 2.0 deployments.

We expect the greater impact on Citrix revenue from cloud is likely to be the web-scale cloud companies such as Microsoft, Google and Amazon buying Netscaler appliances (20-25% of revenue) to accelerate and optimize the delivery of web-based services. We model Netscaler continued growth in the 15% range, a level that is below the growth rate over the last two years.

We expect that Netscaler and CloudPlatform businesses will help raise Citrix overall growth, however we don't expect that these business will drive a re-rating in the stock. This can only happen with acceleration in the growth rate of the desktop business.

Appendix A-1

Analyst Certification

The research analyst(s) primarily responsible for the preparation and content of this research report are named in bold text in the author block at the front of the product except for those sections where an analyst's name appears in bold alongside content which is attributable to that analyst. Each of these analyst(s) certify, with respect to the section(s) of the report for which they are responsible, that the views expressed therein accurately reflect their personal views about each issuer and security referenced and were prepared in an independent manner, including with respect to Citigroup Global Markets Inc and its affiliates. No part of the research analyst's compensation was, is, or will be, directly or indirectly, related to the specific recommendation(s) or view(s) expressed by that research analyst in this report.

IMPORTANT DISCLOSURES

Citigroup Global Markets Inc. owns a position of 1 million USD or more in the debt securities of VMware Inc

Citigroup Global Markets Inc. owns a position of 1 million USD or more in the debt securities of Amazon.Com Inc

A director of Citi serves on the board of Hewlett-Packard Co.

Citigroup Global Markets Inc. owns a position of 1 million USD or more in the debt securities of Citrix Systems Inc

Citigroup Global Markets Inc. owns a position of 1 million USD or more in the debt securities of Salesforce.com Inc

Citigroup Global Markets Inc. owns a position of 1 million USD or more in the debt securities of Microsoft Corp

Citigroup Global Markets Inc. owns a position of 1 million USD or more in the debt securities of Google Inc

A director of Citi serves on the board of IBM Corporation. Citigroup Global Markets Inc. owns a position of 1 million USD or more in the debt securities of International Business Machines Corp

Citigroup Global Markets Inc. owns a position of 1 million USD or more in the debt securities of Yahoo! Inc

Citigroup Global Markets Inc. owns a position of 1 million USD or more in the debt securities of Oracle Corp

Within the past 12 months, Citigroup Global Markets Inc. or its affiliates has acted as manager or co-manager of an offering of securities of VMware, Inc., Microsoft Corp., International Business Machines Corp, Yahoo! Inc., Oracle Corporation.

Citigroup Global Markets Inc. or its affiliates has received compensation for investment banking services provided within the past 12 months from VMware, Inc., SAP AG, ServiceNow, Inc., Hewlett-Packard Co, Microsoft Corp., Google Inc., International Business Machines Corp, Yahoo! Inc., Oracle Corporation.

Citigroup Global Markets Inc. or an affiliate received compensation for products and services other than investment banking services from VMware, Inc., SAP AG, Amazon.com, Inc., Hewlett-Packard Co, Citrix Systems, Inc., salesforce.com, inc., Microsoft Corp., Google Inc., International Business Machines Corp, Red Hat, Inc., Yahoo! Inc., Oracle Corporation in the past 12 months.

Citigroup Global Markets Inc. currently has, or had within the past 12 months, the following as investment banking client(s): Microsoft Corp., ServiceNow, Inc., Yahoo! Inc., VMware, Inc., SAP AG, Hewlett-Packard Co, Google Inc., International Business Machines Corp, Oracle Corporation.

Citigroup Global Markets Inc. currently has, or had within the past 12 months, the following as clients, and the services provided were non-investment-banking, securities-related: Oracle Corporation, VMware, Inc., SAP AG, Amazon.com, Inc., Hewlett-Packard Co, Citrix Systems, Inc., salesforce.com, inc., Microsoft Corp., Google Inc., International Business Machines Corp, Red Hat, Inc., Yahoo! Inc..

Citigroup Global Markets Inc. currently has, or had within the past 12 months, the following as clients, and the services provided were non-investment-banking, non-securities-related: VMware, Inc., SAP AG, Amazon.com, Inc., Hewlett-Packard Co, salesforce.com, inc., Microsoft Corp., Google Inc., International Business Machines Corp, Red Hat, Inc., Yahoo! Inc., Oracle Corporation.

Citigroup Global Markets Inc. or an affiliate received compensation in the past 12 months from Yahoo! Inc..

Analysts' compensation is determined based upon activities and services intended to benefit the investor clients of Citigroup Global Markets Inc. and its affiliates ("the Firm"). Like all Firm employees, analysts receive compensation that is impacted by overall firm profitability which includes investment banking revenues.

The Firm is a market maker in the publicly traded equity securities of ServiceNow, Inc., Amazon.com, Inc., Citrix Systems, Inc., Microsoft Corp., Google Inc., Yahoo! Inc., Oracle Corporation.

For important disclosures (including copies of historical disclosures) regarding the companies that are the subject of this Citi Research product ("the Product"), please contact Citi Research, 388 Greenwich Street, 28th Floor, New York, NY, 10013, Attention: Legal/Compliance [E6WYB6412478]. In addition, the same important disclosures, with the exception of the Valuation and Risk assessments and historical disclosures, are contained on the Firm's disclosure website at https://www.citivelocity.com/cvr/eppublic/citi_research_disclosures. Valuation and Risk assessments can be found in the text of the most recent research note/report regarding the subject company. Historical disclosures (for up to the past three years) will be provided upon request.

Citi Research Equity Ratings Distribution

Data current as of 31 Dec 2013	12 Month Rating			Relative Rating		
	Buy	Hold	Sell	Buy	Hold	Sell
Citi Research Global Fundamental Coverage	49%	40%	12%	6%	88%	6%
% of companies in each rating category that are investment banking clients	55%	52%	44%	62%	52%	49%

Guide to Citi Research Fundamental Research Investment Ratings:

Citi Research stock recommendations include an investment rating and an optional risk rating to highlight high risk stocks.

Risk rating takes into account both price volatility and fundamental criteria. Stocks will either have no risk rating or a High risk rating assigned.

Investment Ratings: Citi Research investment ratings are Buy, Neutral and Sell. Our ratings are a function of analyst expectations of expected total return ("ETR") and risk. ETR is the sum of the forecast price appreciation (or depreciation) plus the dividend yield for a stock within the next 12 months. The Investment rating definitions are: Buy (1) ETR of 15% or more or 25% or more for High risk stocks; and Sell (3) for negative ETR. Any covered stock not assigned a Buy or a Sell is a Neutral (2). For stocks rated Neutral (2), if an analyst believes that there are insufficient valuation drivers and/or investment catalysts to derive a positive or negative investment view, they may elect with the approval of Citi Research management not to assign a target price and, thus, not derive an ETR. Analysts may place covered stocks "Under Review" in response to exceptional circumstances (e.g. lack of information critical to the analyst's thesis) affecting the company and / or trading in the company's securities (e.g. trading suspension). As soon as practically possible, the analyst will publish a note re-establishing a rating and investment thesis. To satisfy regulatory requirements, we correspond Under Review and Neutral to Hold in our ratings distribution table for our 12-month fundamental rating system. However, we reiterate that we do not consider Under Review to be a recommendation.

Relative three-month ratings: Citi Research may also assign a three-month relative call (or rating) to a stock to highlight expected out-performance (most preferred) or under-performance (least preferred) versus the geographic and industry sector over a 3 month period. The relative call may highlight a specific near-term catalyst or event impacting the company or the market that is anticipated to have a short-term price impact on the equity securities of the company. Absent any specific catalyst the analyst(s) will indicate the most and least preferred stocks in the universe of stocks under consideration, explaining the basis for this short-term view. This three-month view may be different from and does not affect a stock's fundamental equity rating, which reflects a longer-term total absolute return expectation. For purposes of NASD/NYSE ratings-distribution-disclosure rules, most preferred calls correspond to a buy recommendation and least preferred calls correspond to a sell recommendation. Any stock not assigned to a most preferred or least preferred call is considered non-relative-rated (NRR). For purposes of NASD/NYSE ratings-distribution-disclosure rules we correspond NRR to Hold in our ratings distribution table for our 3-month relative rating system. However, we reiterate that we do not consider NRR to be a recommendation.

Prior to October 8, 2011, the firm's stock recommendation system included a risk rating and an investment rating. **Risk ratings**, which took into account both price volatility and fundamental criteria, were: Low (L), Medium (M), High (H), and Speculative (S). **Investment Ratings** of Buy, Hold and Sell were a function of the Citi Research expectation of total return (forecast price appreciation and dividend yield within the next 12 months) and risk rating. Additionally, analysts could have placed covered stocks "Under Review" in response to exceptional circumstances (e.g. lack of information critical to the analyst's thesis) affecting the company and/or trading in the company's securities (e.g. trading suspension). Stocks placed "Under Review" were monitored daily by management and as practically possible, the analyst published a note re-establishing a rating and investment thesis. For securities in developed markets (US, UK, Europe, Japan, and Australia/New Zealand), investment ratings were: Buy (1) (expected total return of 10% or more for Low-Risk stocks, 15% or more for Medium-Risk stocks, 20% or more for High-Risk stocks, and 35% or more for Speculative stocks); Hold (2) (0%-10% for Low-Risk stocks, 0%-15% for Medium-Risk stocks, 0%-20% for High-Risk stocks, and 0%-35% for Speculative stocks); and Sell (3) (negative total return). For securities in emerging markets (Asia Pacific, Emerging Europe/Middle East/Africa, and Latin America), investment ratings were: Buy (1) (expected total return of 15% or more for Low-Risk stocks, 20% or more for Medium-Risk stocks, 30% or more for High-Risk stocks, and 40% or more for Speculative stocks); Hold (2) (5%-15% for Low-Risk stocks, 10%-20% for Medium-Risk stocks, 15%-30% for High-Risk stocks, and 20%-40% for Speculative stocks); and Sell (3) (5% or less for Low-Risk stocks, 10% or less for Medium-Risk stocks, 15% or less for High-Risk stocks, and 20% or less for Speculative stocks).

Investment ratings are determined by the ranges described above at the time of initiation of coverage, a change in investment and/or risk rating, or a change in target price (subject to limited management discretion). At other times, the expected total returns may fall outside of these ranges because of market price movements and/or other short-term volatility or trading patterns. Such interim deviations from specified ranges will be permitted but will become subject to review by Research Management. Your decision to buy or sell a security should be based upon your personal investment objectives and should be made only after evaluating the stock's expected performance and risk.

NON-US RESEARCH ANALYST DISCLOSURES

Non-US research analysts who have prepared this report (i.e., all research analysts listed below other than those identified as employed by Citigroup Global Markets Inc.) are not registered/qualified as research analysts with FINRA. Such research analysts may not be associated persons of the member organization and therefore may not be subject to the NYSE Rule 472 and NASD Rule 2711 restrictions on communications with a subject company, public appearances and trading securities held by a research analyst account. The legal entities employing the authors of this report are listed below:

Citigroup Global Markets Inc

Walter H Pritchard, CFA; Robert Chen; Jim Suva, CPA; Mark May; Michael Rollins, CFA

OTHER DISCLOSURES

Citigroup Global Markets Inc. and/or its affiliates has a significant financial interest in relation to SAP AG, Hewlett-Packard Co, Microsoft Corp., Google Inc., International Business Machines Corp, Yahoo! Inc., Oracle Corporation. (For an explanation of the determination of significant financial interest, please refer to the policy for managing conflicts of interest which can be found at www.citiVelocity.com.)

For securities recommended in the Product in which the Firm is not a market maker, the Firm is a liquidity provider in the issuers' financial instruments and may act as principal in connection with such transactions. The Firm is a regular issuer of traded financial instruments linked to securities that may have been recommended in the Product. The Firm regularly trades in the securities of the issuer(s) discussed in the Product. The Firm may engage in securities transactions in a manner inconsistent with the Product and, with respect to securities covered by the Product, will buy or sell from customers on a principal basis.

Securities recommended, offered, or sold by the Firm: (i) are not insured by the Federal Deposit Insurance Corporation; (ii) are not deposits or other obligations of any insured depository institution (including Citibank); and (iii) are subject to investment risks, including the possible loss of the principal amount invested. Although information has been obtained from and is based upon sources that the Firm believes to be reliable, we do not guarantee its accuracy and it may be incomplete and condensed. Note, however, that the Firm has taken all reasonable steps to determine the accuracy and completeness of the disclosures made in the Important Disclosures section of the Product. The Firm's research department has received assistance from the subject company(ies) referred to in this Product including, but not limited to, discussions with management of the subject company(ies). Firm policy prohibits research analysts from sending draft research to subject companies. However, it should be presumed that the author of the Product has had

discussions with the subject company to ensure factual accuracy prior to publication. All opinions, projections and estimates constitute the judgment of the author as of the date of the Product and these, plus any other information contained in the Product, are subject to change without notice. Prices and availability of financial instruments also are subject to change without notice. Notwithstanding other departments within the Firm advising the companies discussed in this Product, information obtained in such role is not used in the preparation of the Product. Although Citi Research does not set a predetermined frequency for publication, if the Product is a fundamental research report, it is the intention of Citi Research to provide research coverage of the/those issuer(s) mentioned therein, including in response to news affecting this issuer, subject to applicable quiet periods and capacity constraints. The Product is for informational purposes only and is not intended as an offer or solicitation for the purchase or sale of a security. Any decision to purchase securities mentioned in the Product must take into account existing public information on such security or any registered prospectus.

Investing in non-U.S. securities, including ADRs, may entail certain risks. The securities of non-U.S. issuers may not be registered with, nor be subject to the reporting requirements of the U.S. Securities and Exchange Commission. There may be limited information available on foreign securities. Foreign companies are generally not subject to uniform audit and reporting standards, practices and requirements comparable to those in the U.S. Securities of some foreign companies may be less liquid and their prices more volatile than securities of comparable U.S. companies. In addition, exchange rate movements may have an adverse effect on the value of an investment in a foreign stock and its corresponding dividend payment for U.S. investors. Net dividends to ADR investors are estimated, using withholding tax rates conventions, deemed accurate, but investors are urged to consult their tax advisor for exact dividend computations. Investors who have received the Product from the Firm may be prohibited in certain states or other jurisdictions from purchasing securities mentioned in the Product from the Firm. Please ask your Financial Consultant for additional details. Citigroup Global Markets Inc. takes responsibility for the Product in the United States. Any orders by US investors resulting from the information contained in the Product may be placed only through Citigroup Global Markets Inc.

The Citigroup legal entity that takes responsibility for the production of the Product is the legal entity which the first named author is employed by. The Product is made available in **Australia** through Citi Global Markets Australia Pty Ltd. (ABN 64 003 114 832 and AFSL No. 240992), participant of the ASX Group and regulated by the Australian Securities & Investments Commission. Citigroup Centre, 2 Park Street, Sydney, NSW 2000. The Product is made available in Australia to Private Banking wholesale clients through Citigroup Pty Limited (ABN 88 004 325 080 and AFSL 238098). Citigroup Pty Limited provides all financial product advice to Australian Private Banking wholesale clients through bankers and relationship managers. If there is any doubt about the suitability of investments held in Citigroup Private Bank accounts, investors should contact the Citigroup Private Bank in Australia. Citigroup companies may compensate affiliates and their representatives for providing products and services to clients. The Product is made available in **Brazil** by Citigroup Global Markets Brasil - CCTVM SA, which is regulated by CVM - Comissão de Valores Mobiliários, BACEN - Brazilian Central Bank, APIMEC - Associação dos Analistas e Profissionais de Investimento do Mercado de Capitais and ANBID - Associação Nacional dos Bancos de Investimento. Av. Paulista, 1111 - 11º andar - CEP. 01311920 - São Paulo - SP. If the Product is being made available in certain provinces of **Canada** by Citigroup Global Markets (Canada) Inc. ("CGM Canada"), CGM Canada has approved the Product. Citigroup Place, 123 Front Street West, Suite 1100, Toronto, Ontario M5J 2M3. This product is available in **Chile** through Banchile Corredores de Bolsa S.A., an indirect subsidiary of Citigroup Inc., which is regulated by the Superintendencia de Valores y Seguros. Agustinas 975, piso 2, Santiago, Chile. The Product is distributed in **Germany** by Citigroup Global Markets Deutschland AG ("CGMD"), which is regulated by Bundesanstalt fuer Finanzdienstleistungsaufsicht (BaFin). CGMD, Reuterweg 16, 60323 Frankfurt am Main. Research which relates to "securities" (as defined in the Securities and Futures Ordinance (Cap. 571 of the Laws of Hong Kong)) is issued in **Hong Kong** by, or on behalf of, Citigroup Global Markets Asia Limited which takes full responsibility for its content. Citigroup Global Markets Asia Ltd. is regulated by Hong Kong Securities and Futures Commission. If the Research is made available through Citibank, N.A., Hong Kong Branch, for its clients in Citi Private Bank, it is made available by Citibank N.A., Citibank Tower, Citibank Plaza, 3 Garden Road, Hong Kong. Citibank N.A. is regulated by the Hong Kong Monetary Authority. Please contact your Private Banker in Citibank N.A., Hong Kong, Branch if you have any queries on or any matters arising from or in connection with this document. The Product is made available in **India** by Citigroup Global Markets India Private Limited, which is regulated by Securities and Exchange Board of India. Bakhtawar, Nariman Point, Mumbai 400-021. The Product is made available in **Indonesia** through PT Citigroup Securities Indonesia. 5/F, Citibank Tower, Bapindo Plaza, Jl. Jend. Sudirman Kav. 54-55, Jakarta 12190. Neither this Product nor any copy hereof may be distributed in Indonesia or to any Indonesian citizens wherever they are domiciled or to Indonesian residents except in compliance with applicable capital market laws and regulations. This Product is not an offer of securities in Indonesia. The securities referred to in this Product have not been registered with the Capital Market and Financial Institutions Supervisory Agency (BAPEPAM-LK) pursuant to relevant capital market laws and regulations, and may not be offered or sold within the territory of the Republic of Indonesia or to Indonesian citizens through a public offering or in circumstances which constitute an offer within the meaning of the Indonesian capital market laws and regulations. The Product is made available in **Israel** through Citibank NA, regulated by the Bank of Israel and the Israeli Securities Authority. Citibank, N.A., Platinum Building, 21 Ha'arba'ah St, Tel Aviv, Israel. The Product is made available in **Italy** by Citigroup Global Markets Limited, which is authorised by the PRA and regulated by the FCA and the PRA. Via dei Mercanti, 12, Milan, 20121, Italy. The Product is made available in **Japan** by Citigroup Global Markets Japan Inc. ("CGMJ"), which is regulated by Financial Services Agency, Securities and Exchange Surveillance Commission, Japan Securities Dealers Association, Tokyo Stock Exchange and Osaka Securities Exchange. Shin-Marunouchi Building, 1-5-1 Marunouchi, Chiyoda-ku, Tokyo 100-6520 Japan. If the Product was distributed by SMBC Nikko Securities Inc. it is being so distributed under license. In the event that an error is found in an CGMJ research report, a revised version will be posted on the Firm's Citi Velocity website. If you have questions regarding Citi Velocity, please call (81 3) 6270-3019 for help. The Product is made available in **Korea** by Citigroup Global Markets Korea Securities Ltd., which is regulated by the Financial Services Commission, the Financial Supervisory Service and the Korea Financial Investment Association (KOFIA). Citibank Building, 39 Da-dong, Jung-gu, Seoul 100-180, Korea. KOFIA makes available registration information of research analysts on its website. Please visit the following website if you wish to find KOFIA registration information on research analysts of Citigroup Global Markets Korea Securities Ltd. <http://dis.kofia.or.kr/fs/dis2/fundMgr/DISFundMgrAnalystPop.jsp?companyCd2=A03030&pageDiv=02>. The Product is made available in Korea by Citibank Korea Inc., which is regulated by the Financial Services Commission and the Financial Supervisory Service. Address is Citibank Building, 39 Da-dong, Jung-gu, Seoul 100-180, Korea. The Product is made available in **Malaysia** by Citigroup Global Markets Malaysia Sdn Bhd (Company No. 460819-D) ("CGMM") to its clients and CGMM takes responsibility for its contents. CGMM is regulated by the Securities Commission of Malaysia. Please contact CGMM at Level 43 Menara Citibank, 165 Jalan Ampang, 50450 Kuala Lumpur, Malaysia in respect of any matters arising from, or in connection with, the Product. The Product is made available in **Mexico** by Acciones y Valores Banamex, S.A. De C. V., Casa de Bolsa, Integrante del Grupo Financiero Banamex ("Accival") which is a wholly owned subsidiary of Citigroup Inc. and is regulated by Comision Nacional Bancaria y de Valores. Reforma 398, Col. Juarez, 06600 Mexico, D.F. In **New Zealand** the Product is made available to 'wholesale clients' only as defined by s5C(1) of the Financial Advisers Act

2008 ('FAA') through Citigroup Global Markets Australia Pty Ltd (ABN 64 003 114 832 and AFSL No. 240992), an overseas financial adviser as defined by the FAA, participant of the ASX Group and regulated by the Australian Securities & Investments Commission. Citigroup Centre, 2 Park Street, Sydney, NSW 2000. The Product is made available in **Pakistan** by Citibank N.A. Pakistan branch, which is regulated by the State Bank of Pakistan and Securities Exchange Commission, Pakistan. AWT Plaza, 1.1. Chundrigar Road, P.O. Box 4889, Karachi-74200. The Product is made available in the **Philippines** through Citicorp Financial Services and Insurance Brokerage Philippines, Inc., which is regulated by the Philippines Securities and Exchange Commission. 20th Floor Citibank Square Bldg. The Product is made available in the Philippines through Citibank NA Philippines branch, Citibank Tower, 8741 Paseo De Roxas, Makati City, Manila. Citibank NA Philippines NA is regulated by The Bangko Sentral ng Pilipinas. The Product is made available in **Poland** by Dom Maklerski Banku Handlowego SA an indirect subsidiary of Citigroup Inc., which is regulated by Komisja Nadzoru Finansowego. Dom Maklerski Banku Handlowego S.A. ul.Senatorska 16, 00-923 Warszawa. The Product is made available in the **Russian Federation** through ZAO Citibank, which is licensed to carry out banking activities in the Russian Federation in accordance with the general banking license issued by the Central Bank of the Russian Federation and brokerage activities in accordance with the license issued by the Federal Service for Financial Markets. Neither the Product nor any information contained in the Product shall be considered as advertising the securities mentioned in this report within the territory of the Russian Federation or outside the Russian Federation. The Product does not constitute an appraisal within the meaning of the Federal Law of the Russian Federation of 29 July 1998 No. 135-FZ (as amended) On Appraisal Activities in the Russian Federation. 8-10 Gasheka Street, 125047 Moscow. The Product is made available in **Singapore** through Citigroup Global Markets Singapore Pte. Ltd. ("CGMSPL"), a capital markets services license holder, and regulated by Monetary Authority of Singapore. Please contact CGMSPL at 8 Marina View, 21st Floor Asia Square Tower 1, Singapore 018960, in respect of any matters arising from, or in connection with, the analysis of this document. This report is intended for recipients who are accredited, expert and institutional investors as defined under the Securities and Futures Act (Cap. 289). The Product is made available by The Citigroup Private Bank in Singapore through Citibank, N.A., Singapore Branch, a licensed bank in Singapore that is regulated by Monetary Authority of Singapore. Please contact your Private Banker in Citibank N.A., Singapore Branch if you have any queries on or any matters arising from or in connection with this document. This report is intended for recipients who are accredited, expert and institutional investors as defined under the Securities and Futures Act (Cap. 289). This report is distributed in Singapore by Citibank Singapore Ltd ("CSL") to selected Citigold/Citigold Private Clients. CSL provides no independent research or analysis of the substance or in preparation of this report. Please contact your Citigold/Citigold Private Client Relationship Manager in CSL if you have any queries on or any matters arising from or in connection with this report. This report is intended for recipients who are accredited investors as defined under the Securities and Futures Act (Cap. 289). Citigroup Global Markets (Pty) Ltd. is incorporated in the **Republic of South Africa** (company registration number 2000/025866/07) and its registered office is at 145 West Street, Sandton, 2196, Saxonwold. Citigroup Global Markets (Pty) Ltd. is regulated by JSE Securities Exchange South Africa, South African Reserve Bank and the Financial Services Board. The investments and services contained herein are not available to private customers in South Africa. The Product is made available in the **Republic of China** through Citigroup Global Markets Taiwan Securities Company Ltd. ("CGMTS"), 14 and 15F, No. 1, Songzhi Road, Taipei 110, Taiwan and/or through Citibank Securities (Taiwan) Company Limited ("CSTL"), 14 and 15F, No. 1, Songzhi Road, Taipei 110, Taiwan, subject to the respective license scope of each entity and the applicable laws and regulations in the Republic of China. CGMTS and CSTL are both regulated by the Securities and Futures Bureau of the Financial Supervisory Commission of Taiwan, the Republic of China. No portion of the Product may be reproduced or quoted in the Republic of China by the press or any third parties [without the written authorization of CGMTS and CSTL]. If the Product covers securities which are not allowed to be offered or traded in the Republic of China, neither the Product nor any information contained in the Product shall be considered as advertising the securities or making recommendation of the securities in the Republic of China. The Product is for informational purposes only and is not intended as an offer or solicitation for the purchase or sale of a security or financial products. Any decision to purchase securities or financial products mentioned in the Product must take into account existing public information on such security or the financial products or any registered prospectus. The Product is made available in **Thailand** through Citicorp Securities (Thailand) Ltd., which is regulated by the Securities and Exchange Commission of Thailand. 18/F, 22/F and 29/F, 82 North Sathorn Road, Silom, Bangrak, Bangkok 10500, Thailand. The Product is made available in **Turkey** through Citibank AS which is regulated by Capital Markets Board. Tekfen Tower, Eski Buyukdere Caddesi # 209 Kat 2B, 23294 Levent, Istanbul, Turkey. In the **U.A.E.**, these materials (the "Materials") are communicated by Citigroup Global Markets Limited, DIFC branch ("CGML"), an entity registered in the Dubai International Financial Center ("DIFC") and licensed and regulated by the Dubai Financial Services Authority ("DFS") to Professional Clients and Market Counterparties only and should not be relied upon or distributed to Retail Clients. A distribution of the different Citi Research ratings distribution, in percentage terms for Investments in each sector covered is made available on request. Financial products and/or services to which the Materials relate will only be made available to Professional Clients and Market Counterparties. The Product is made available in **United Kingdom** by Citigroup Global Markets Limited, which is authorised by the Prudential Regulation Authority ("PRA") and regulated by the Financial Conduct Authority ("FCA") and the PRA. This material may relate to investments or services of a person outside of the UK or to other matters which are not authorised by the PRA nor regulated by the FCA and the PRA and further details as to where this may be the case are available upon request in respect of this material. Citigroup Centre, Canada Square, Canary Wharf, London, E14 5LB. The Product is made available in **United States** by Citigroup Global Markets Inc, which is a member of FINRA and registered with the US Securities and Exchange Commission. 388 Greenwich Street, New York, NY 10013. Unless specified to the contrary, within EU Member States, the Product is made available by Citigroup Global Markets Limited, which is authorised by the PRA and regulated by the FCA and the PRA.

Pursuant to Comissão de Valores Mobiliários Rule 483, Citi is required to disclose whether a Citi related company or business has a commercial relationship with the subject company. Considering that Citi operates multiple businesses in more than 100 countries around the world, it is likely that Citi has a commercial relationship with the subject company.

Many European regulators require that a firm must establish, implement and make available a policy for managing conflicts of interest arising as a result of publication or distribution of investment research. The policy applicable to Citi Research's Products can be found at https://www.citivelocity.com/cvr/epublic/citi_research_disclosures.

Compensation of equity research analysts is determined by equity research management and Citigroup's senior management and is not linked to specific transactions or recommendations.

The Product is not to be construed as providing investment services in any jurisdiction where the provision of such services would not be permitted.

Subject to the nature and contents of the Product, the investments described therein are subject to fluctuations in price and/or value and investors may get back less than originally invested. Certain high-volatility investments can be subject to sudden and large falls in value that could equal or exceed the amount invested. Certain investments contained in the Product may have tax implications for private customers whereby levels and basis of taxation may be subject to change. If in doubt, investors should seek advice from a tax adviser. The Product does not purport to identify the nature of the specific market or other

risks associated with a particular transaction. Advice in the Product is general and should not be construed as personal advice given it has been prepared without taking account of the objectives, financial situation or needs of any particular investor. Accordingly, investors should, before acting on the advice, consider the appropriateness of the advice, having regard to their objectives, financial situation and needs. Prior to acquiring any financial product, it is the client's responsibility to obtain the relevant offer document for the product and consider it before making a decision as to whether to purchase the product. Citi Research generally disseminates its research to the Firm's global institutional and retail clients via both proprietary (e.g., Citi Velocity and Citi Personal Wealth Management) and non-proprietary electronic distribution platforms. Certain research may be disseminated only via Citi's proprietary distribution platforms; however such research will not contain changes to earnings forecasts, target price, investment or risk rating or investment thesis or be otherwise inconsistent with the author's previously published research. Certain research is made available only to institutional investors to satisfy regulatory requirements. Individual Citi Research analysts may also opt to circulate published research to one or more clients by email; such email distribution is discretionary and is done only after the research has been disseminated.

The level and types of services provided by Citi Research analysts to clients may vary depending on various factors such as the client's individual preferences as to the frequency and manner of receiving communications from analysts, the client's risk profile and investment focus and perspective (e.g. market-wide, sector specific, long term, short-term etc.), the size and scope of the overall client relationship with Citi and legal and regulatory constraints. Citi Research product may source data from dataCentral. dataCentral is a Citi Research proprietary database, which includes Citi estimates, data from company reports and feeds from Thomson Reuters.

© 2014 Citigroup Global Markets Inc. Citi Research is a division of Citigroup Global Markets Inc. Citi and Citi with Arc Design are trademarks and service marks of Citigroup Inc. and its affiliates and are used and registered throughout the world. All rights reserved. Any unauthorized use, duplication, redistribution or disclosure of this report (the "Product"), including, but not limited to, redistribution of the Product by electronic mail, posting of the Product on a website or page, and/or providing to a third party a link to the Product, is prohibited by law and will result in prosecution. The information contained in the Product is intended solely for the recipient and may not be further distributed by the recipient to any third party. Where included in this report, MSCI sourced information is the exclusive property of Morgan Stanley Capital International Inc. (MSCI). Without prior written permission of MSCI, this information and any other MSCI intellectual property may not be reproduced, redisseminated or used to create any financial products, including any indices. This information is provided on an "as is" basis. The user assumes the entire risk of any use made of this information. MSCI, its affiliates and any third party involved in, or related to, computing or compiling the information hereby expressly disclaim all warranties of originality, accuracy, completeness, merchantability or fitness for a particular purpose with respect to any of this information. Without limiting any of the foregoing, in no event shall MSCI, any of its affiliates or any third party involved in, or related to, computing or compiling the information have any liability for any damages of any kind. MSCI, Morgan Stanley Capital International and the MSCI indexes are services marks of MSCI and its affiliates. The Firm accepts no liability whatsoever for the actions of third parties. The Product may provide the addresses of, or contain hyperlinks to, websites. Except to the extent to which the Product refers to website material of the Firm, the Firm has not reviewed the linked site. Equally, except to the extent to which the Product refers to website material of the Firm, the Firm takes no responsibility for, and makes no representations or warranties whatsoever as to, the data and information contained therein. Such address or hyperlink (including addresses or hyperlinks to website material of the Firm) is provided solely for your convenience and information and the content of the linked site does not in anyway form part of this document. Accessing such website or following such link through the Product or the website of the Firm shall be at your own risk and the Firm shall have no liability arising out of, or in connection with, any such referenced website.

ADDITIONAL INFORMATION IS AVAILABLE UPON REQUEST
