

Business white paper

Achieving the speed and agility of cloud giants

Composable infrastructure can help you succeed



Managing traditional and new infrastructure for speed

Table of contents

- 2 Managing traditional and new infrastructure for speed
- 3 Challenges of operating like a cloud giant
- 6 A new class of infrastructure that gives you cloud giant efficiencies
- 6 The path to the future of IT

It is easier than ever to disrupt industries and markets with new innovations. Cloud and mobile computing combined with the rise in social media has given business unprecedented opportunities to reach new customers with innovative products and services. This is called the Idea Economy, and it allows new businesses to impact their industries quickly while often leaving their more established competitors scrambling to keep up.

Succeeding in the Idea Economy requires that businesses have the ability to turn their ideas into value faster than the competition. The Idea Economy puts tremendous pressure on IT to maintain traditional operations while implementing new applications and services for mobile, social, and the cloud. The resulting shift towards on-demand, high-performance applications has left IT organizations scrambling.

“Between 2013 and 2020, organizations’ investment in mobile, social, and Big Data technologies (3rd platform) will grow over 20 times faster than organizations’ investment in client/server technologies (2nd platform).”¹

– IDC prediction

Traditionally, IT operations focused on maintaining existing infrastructure that was largely designed to address core business needs and running prepackaged applications to automate internal business functions. For these non-revenue generating applications, the main priorities for IT staff were minimizing cost, risk, and inefficiency using standard methodologies, conventional infrastructure, and strong governance.

Today, however, customer demands and expectations are forcing businesses to shift the focus of IT operations. IT must now support two very different modes of operation for delivering services:

1. Balance resources and investments needed to keep core legacy systems up and running.
2. Create new value by integrating new infrastructure and services to keep the business competitive and growing.

Managing both successfully is key to the sustainability of any business today.

Currently, enterprises are attempting to support these disparate models with separate teams, hardware, and processes. In an effort to reduce complexity and consolidate operations, some IT departments are attempting to implement public cloud operations as a solution. However, the significant differences between the enterprise environment and the service provider environment make it difficult for enterprises to reach a fully optimized state.

¹The Third Platform Enabling Digital Transformation, IDC, November 2013

Challenges of operating like a cloud giant

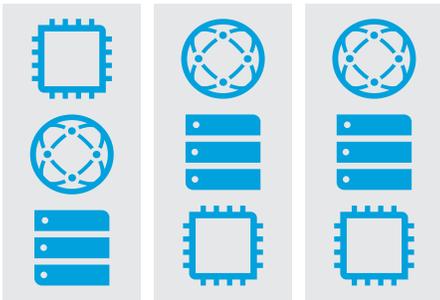
The biggest cloud-computing providers in the world—Tier-1 providers who are known as cloud giants—employ huge collections of servers, sometimes numbering in millions, to provide access to a finite set of applications for users around the world. Because applications are their primary source of revenue, cloud providers need the ability to build on flexible infrastructure that can adjust to the speed of their business.

Procurement and provisioning: Traditional enterprise IT simply cannot match the efficiency of cloud giants' custom server environments while maintaining traditional infrastructure. Cloud giants have the luxury of buying thousands of servers, and the associated networking and storage, at one time and provisioning everything in advance to meet the specifications of a handful of unique applications. Then they simply turn on pre-configured, pre-defined, and pre-wired equipment and provisioning for the application they want to run. Comparatively, the traditional enterprise hardware implementation process usually takes much longer since a large majority of applications carry unique requirements. Unlike cloud giants, enterprise IT cannot choose the infrastructure it needs until the applications are selected. Also, server, network, and storage equipment are often purchased by separate IT teams and then brought together for provisioning—a process which can take weeks or even months to coordinate and complete. Enterprise IT needs the ability to deploy, all at one time, the compute, storage and fabric needed to run their applications in sizes appropriate for their scale so new applications can be quickly provisioned.



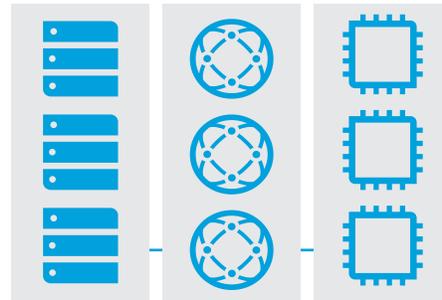
Cloud

Fully provisioned infrastructure purchased in massive increment



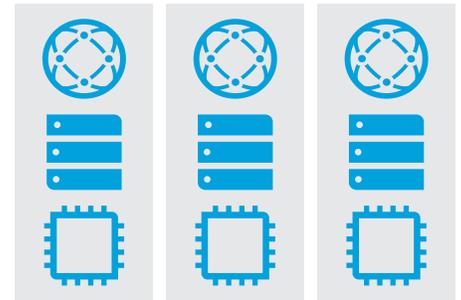
Traditional enterprise

Purchase silos in bulk and connected by long processes



New enterprise model

Right-sized, fully provisioned infrastructure



Hardware and applications: The applications and data for a cloud giant are typically designed to run across thousands of servers, with little variation in hardware due to their operational models. Comparatively, most enterprises operate only a fraction of the servers with a much greater degree of application and hardware diversity.

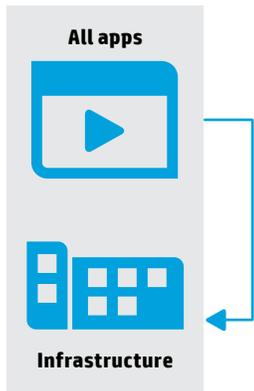
Cloud giants are able to operate efficiently because their software applications are written with the hardware in mind, making management of this vast infrastructure nimble and efficient. Enterprises, on the other hand, run mostly packaged applications, which often require optimized hardware configurations to meet performance requirements, adding more complexity across the infrastructure to manage.

In addition, depending on the application, cloud giants may be able to leverage inexpensive, shared-nothing storage architectures where data is stored in object format across multiple servers, creating redundancy through the application itself. This allows cloud providers to build inexpensive and low-resilience data centers with high-resilience software that runs smoothly with minimal downtime. Alternatively, most enterprises have “off-the-shelf” applications, which store data in block or file formats and are dependent on the SAN to enable shared access. These applications are not built with resilience, and therefore, they need resilience at the server- and storage-hardware level in order to run with high availability.

Like cloud giants, enterprises would benefit greatly if they could run all of their applications on a homogeneous infrastructure with the flexibility and resiliency necessary for both traditional and cloud-native applications. However, since packaged applications cannot be counted on to adapt to infrastructure, infrastructure must adapt to the needs of the applications.

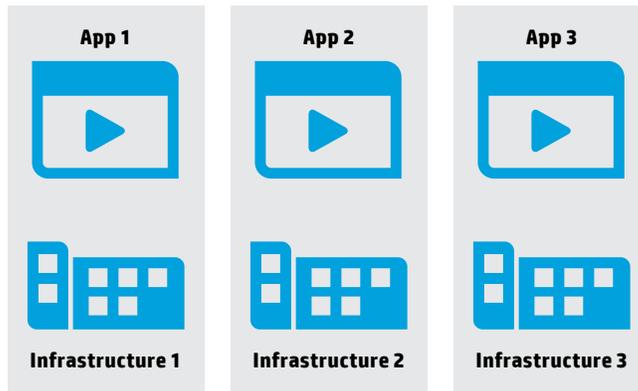
Cloud

Application flexes to the infrastructure



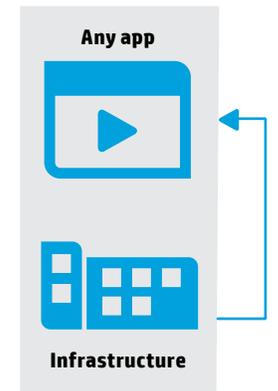
Traditional enterprise

Custom infrastructure per application



New enterprise model

Infrastructure flexes to the application



Operations and maintenance: The biggest IT expense for most businesses is labor.² Most enterprises spend about three times as much on labor as equipment. In contrast, when configuring hardware, cloud giants rely heavily on automation to provision resources. The homogeneity of their infrastructure makes this automation process simple. For an enterprise, hardware configuration is typically done through a management tool manually operated by an administrator. Some customers attempt to automate their processes. However, applications may have unique hardware size and configuration requirements, leading to a lengthy script writing process. Since there are a variety of devices and each device manager typically has a different scriptable interface or API, different data formats, and different data models, the complexity of this process significantly drives cost up and slows time-to-value.

“In one large cloud provider, they host tens of thousands of servers but only have 11 administrative staff to service the equipment. In the same provider, it takes around 90 people to run hundreds of servers in its managed business.”³

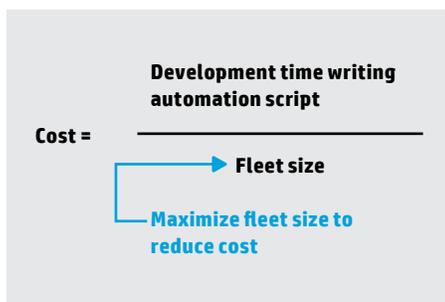
– Forrester Research

The scale and homogeneity of cloud giants’ operations leads to a lower cost-per-server maintenance spend than enterprises. Though the manpower required to write an automated script for implementing a configuration change is about the same, cloud giants spread the cost of writing one automated script across a vast number of servers, so it can be done at a fraction of the cost-per-server because of their scale. For an enterprise, that same scripting cost may only apply to a small number of servers and the hardware diversity makes it more difficult.

Since enterprises cannot reduce their cost-per-server through server fleet growth, they need to look at reducing costs by utilizing a substantially simpler automation interface for administrators. To optimize efficiency, enterprises need an easier way to automate maintenance by eliminating the complexity inherent in standard equipment.

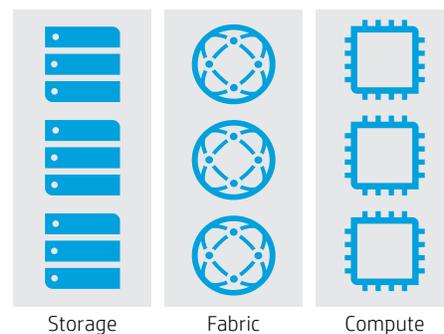
Cloud

Fully automated over massive homogenous fleet



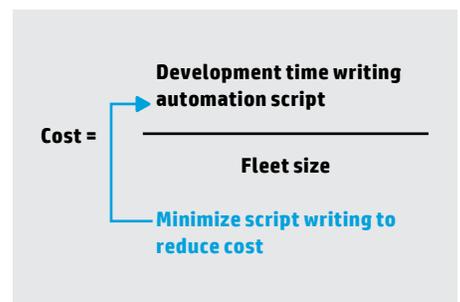
Traditional enterprise

Too many operations to configure each manually



New enterprise model

Fully automated over mostly homogenous fleet



²Business Technographics Global Budgets Survey, Forrester Research, August 2014 (In traditional IT, labor is 39 percent of operational costs while hardware is only 14 percent.)

³Five Data Center and IT Infrastructure Lessons From the Cloud Giants, Forrester Research, August 2013



Run anything, store everything



Deliver apps and services dramatically faster



Simplify lifecycle operations and integration



Unlock value creation

A new class of infrastructure that gives you cloud giant efficiencies

How can IT manage both traditional and cloud-native applications while gaining the speed, agility, and efficiency of cloud giants? A powerful new architecture called composable infrastructure integrates compute, storage, and fabric all-in-one solution so enterprises can easily implement the right-size infrastructure to run their applications. This enables transformational capabilities:

Bringing compute, storage and fabric together: Pre-configured, more flexible, and ready to run, a composable infrastructure can be provisioned on-demand as applications require. This dramatically cuts the time needed for setup and deployment of new infrastructure to run a new application or IT service. IT can optimize operations and meet business demands while still effectively managing efficiency and scale with a composable infrastructure.

Flexible and composable hardware runs all applications: A composable infrastructure has both resiliency and flexibility at the hardware level so it can adapt to the needs of virtually any application. It has the unique ability to run without being limited to any one computing paradigm because it can run virtual machines, bare-metal deployments, containers, and cloud-native applications.

Fluid pools of resources that can be allocated to meet the needs of applications give composable infrastructure the ability to run and store virtually anything. It can disaggregate and re-aggregate underlying blocks of resources with physical, virtual, and container resource pools. To meet application requirements, storage can directly attached, network-attached, or software-defined and be configured to present data services in block, file, or object storage formats. This provides the flexibility to respond effortlessly to unpredictable demands, regardless of data type, connectivity protocol, or service level requirement.

Ease of maintenance: Composable infrastructure simplifies maintenance by using software-defined intelligence with template-driven, single-touch provisioning and operations management. With a self-discovering, self-securing, self-orchestrating, and self-diagnosing nature, it minimizes manual intervention. This flexible architecture uses a single-pane-of-glass, as well as a unified API, to simplify automation of operations and maintenance tasks.

The path to the future of IT

Composable infrastructure lets enterprise IT operate like a cloud provider to lines of business and the extended enterprise. It maximizes the speed, agility, and efficiency of core infrastructure and operations to consistently meet SLAs and provide the predictable performance needed to support core workloads—for today and tomorrow.

Moving to a composable infrastructure is an incremental process. If you'd like to gain a deeper understanding, read "[HP Composable Infrastructure: Bridging traditional IT with the New Style of Business.](#)"

Learn more at
hp.com/go/ComposableProgram

Sign up for updates
hp.com/go/getupdated



Share with colleagues



Rate this document





Work smarter

At Insight, we'll help you solve challenges and improve performance with intelligent technology solutions.

[Learn more](#)

