



CORPORATE PERFORMANCE IMPROVEMENT

DOES CLOUD MEAN THE PRIVATE DATA CENTER IS DEAD?



DOES CLOUD MEAN THE PRIVATE DATA CENTER IS DEAD?

MASS MIGRATION: SHOULD ALL COMPANIES MOVE TO THE CLOUD?

Achieving digital efficiency through the optimization of technology infrastructure, energy and support costs can materially improve company profitability. Since technology powers nearly all aspects of business today, this goal should be top of mind for every organization. However, while Gartner recently declared “[The Data Center is Dead](#)” and estimated that by 2025 80 percent of companies will use public cloud solutions, **there is no one-size-fits-all infrastructure for achieving digital efficiency.**

To be sure, a case exists for many companies with private data centers (DCs) to move to the public cloud between now and 2025. However, **there is a point of digital efficiency between cost and server capacity beyond which a private data center can yield higher returns on investment than a public cloud solution.** The key is identifying that point, which is unique to every organization based on variables like current hosted capacity, regulatory requirements, company size and computing needs.

This article explains the technological changes that have allowed public cloud solutions to compete with private DCs and explores cost advantages and disadvantages of both public cloud solutions and private DCs.

THE PUBLIC CLOUD BUILDUP

Public cloud hosting is possible because of lower-cost low-latency bandwidth, advancing security and efficient hardware sharing.

High-bandwidth internet connectivity with very low latency is now prevalent and affordable in the United States, the Pacific Rim and Europe. When combined with secure, virtualized or containerized applications, this advance has created the perfect conditions for cloud migration.

Today’s technology enables remote computing that is high-performing and reliable. Multiple users or businesses can easily share vast pools of hardware, infrastructure and a multitude of applications through Software as a Service (SaaS) and Platform as a Service (PaaS) models. This situation creates a very profitable business model for cloud providers.

Like utility companies, cloud providers use scale to their advantage.

Scale allows cloud providers to compete against private DCs on cost while still turning a significant profit. For example, Amazon makes nearly all its operating income from its Amazon Web Service (AWS) subsidiary. Just like large utility companies use scale to provide power at a fraction of the cost that any one municipality would need to charge, cloud providers can enjoy huge profits while remaining less expensive than a private DC.

Unlike large-scale cloud providers, private DCs frequently cannot efficiently employ all their computing and physical resources. Prior to virtual machines (VMs), which consolidate many system images on one piece of hardware, average utilization for DCs was often in the low single digits. Even now, with VMs, an average of 25-40 percent central processing unit (CPU) utilization is considered good. That includes utilization for backup, development and test machines. In contrast, cloud providers can sustainably run at 70 to 80 percent utilization.

Establishing and running a DC requires considerable investment.

Fitting out an empty DC or DC co-locations can easily cost more than \$1 million in capital with a four-year capital depreciation rate that is \$0.5 million per year, and two DCs are essential for disaster recovery. With a public cloud solution, you pay only for what you use and the startup costs are built in and shared among all cloud customers.

Public cloud solutions also reduce labor costs. To run a private DC, a company must invest in skilled labor to perform tasks that cloud providers include as part of their service. Qualified personnel are needed for hardware and software provisioning, wide area network (WAN) operations and regular maintenance and support. A well-run private DC might require a team of 10 to 20 people at a run-rate of \$1-4 million per year. Labor costs do not scale linearly with the number of computers. Rather, efficiency in staffing increases as the size of the DC operation increases.

Amortization or rent for the DC must also be considered. DC space is often rented in rows of 10 cabinets which can house 200 servers and cost over \$1 million per year.

Cloud solutions provide a clear cost advantage ... until they don't.

On the surface the cost difference between two private DCs (\$3.5 million to \$6.5 million per year) and a public cloud subscription for one VM (\$2,000 – \$8,000 per year) makes the decision a no-brainer. However, as the size of compute needs rises, the cost advantage erodes and eventually disappears.

THE CASE FOR PRIVATE DATA CENTERS

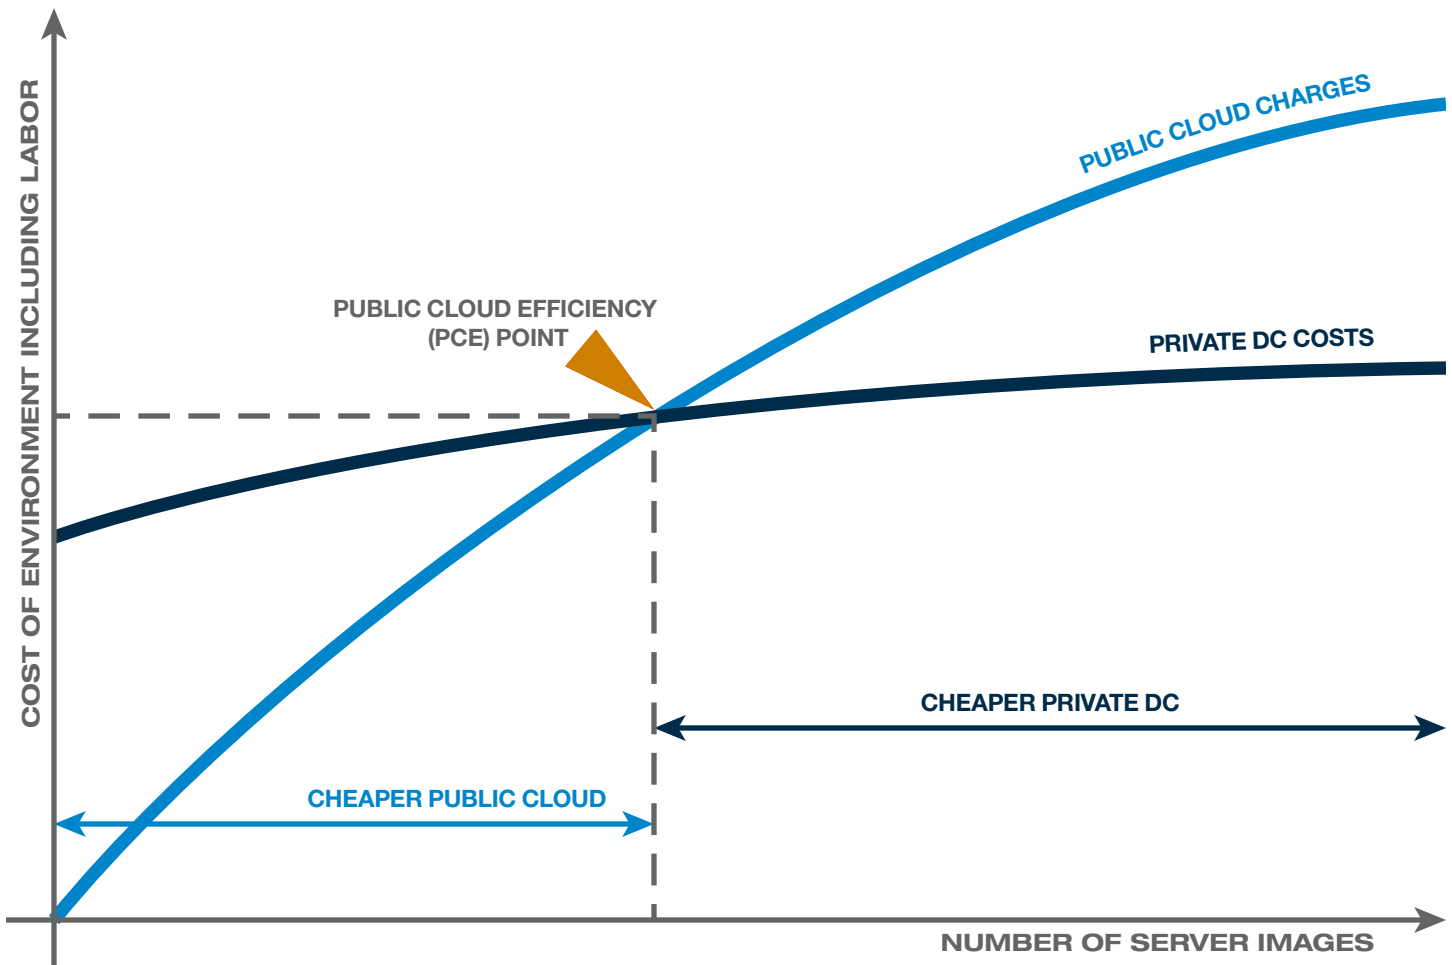
Cloud solutions are not the best option for every business.

There exists a point of digital efficiency between cost and server capacity that we can call the Public Cloud Efficiency (PCE) point (Exhibit 1). At this point, which is unique for every business, the number of server images needed pushes the total cost of the computing environment, inclusive of labor, to a level that is more expensive in the public cloud than it is in a private DC.

Companies that operate their own data centers at large scale will find the public cloud model increasingly less cost-efficient past the PCE point. This is because as the size of the compute problem scales, the relatively high fixed costs of a private data center amortize over the large installed compute plant, thereby lowering the unit cost for private DCs below that of the public cloud providers.

Exhibit 1

The Public Cloud Efficiency Point



Source: Alvarez & Marsal, LLC

Even if a company falls below the PCE, there are other considerations that may delay migration to the cloud.

Timing of contracts, leases and divestments, for example, can create roadblocks to immediate migration. DC leasing contracts typically have 5 or 10-year lifespans that may still be amortizing on the company books. Until divestment, renewal or replacement, these costs will persist even if a business transitions to the cloud, meaning it will have to bear the cost of both the DC and cloud solution simultaneously for a period. These combined expenses could effectively extend the PCE point.

Regulatory constraints may also impose delays to cloud migration. For example, regulations pertaining to personally identifiable information, health insurance portability and accountability (HIPAA) and payment data create high-risk barriers to companies in the financial and healthcare industries. Only recently have concerns about the stability and security of public cloud providers started to dissipate, primarily for the largest providers (e.g., Amazon, MSN and Google).

Application requirements must be considered as well. Many of the largest business-critical applications are not yet containerized, meaning they need dedicated cloud resources which can cost 10 to 20 times as much as variable cloud resources, potentially moving the companies needs beyond the PCE point until they are re-engineered.

Strategic planning is essential for timing cloud migration to maximize efficiency. Any transition plan should include considerations for the organization's current infrastructure, applications, regulatory requirements and support groups to make sure all are aligned and ready to move.

MAKING THE CALL TO MIGRATE

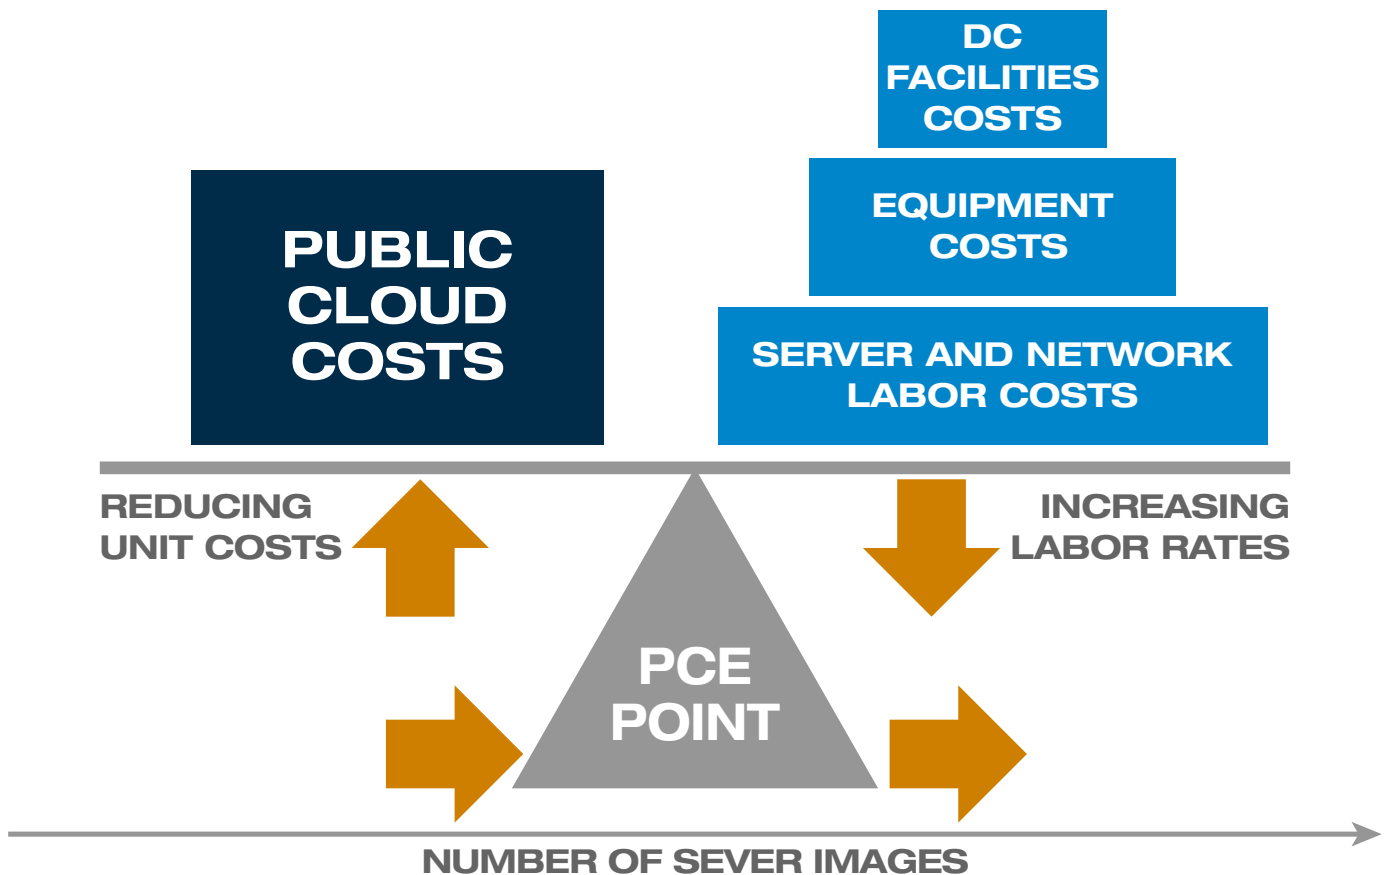
The PCE depends on the unique circumstances of a business, but it can also shift with advancements in technology.

For the sake of example, let us assume that an average business with no special considerations requires 4,000 VMs. By applying Moore's Law, named for the founder of Intel, one can predict the amount of compute power a single server provides will double every 18 months. This means that by 2025, the same footprint of servers today will provide eight times the capacity. In other words, the same DC footprint that accommodates 4,000 VMs today will allow for 32,000 to 64,000 in just seven years.

While server capacity is increasing at this rate, the prices that public cloud companies charge for compute power have been dropping roughly in line with Moore's Law, which means the PCE point will increase over time. Exhibit 2 illustrates how the balance between public cloud and DC costs rests on the fulcrum of the PCE point.

Exhibit 2

Impact of Public Cloud and Data Center Costs on PCE



Source: Alvarez & Marsal, LLC

Based on these predictions and trends, an average company running fewer than 32,000 VMs by 2025 will find that using the cloud is more cost effective than running a private DC. In contrast, a company with larger demand at that time will likely conclude that running their own DC is more cost-effective than a public cloud solution.

For businesses with VM demands that won't grow to that point and are already below today's 4,000 VM PCE that was used in our example, predicting that by 2025 they will have moved to the cloud is a safe bet.

Knowing this, all business leaders should ask and answer three questions:

1. Is your company above the PCE point today?
2. Will the combined compute and storage demand of your business be larger than the PCE point by 2025?
3. Do you have a strategic plan for digital efficiency that considers current infrastructure, applications, regulatory requirements, support groups and timing?

Technology changes quickly. Regular assessment is key to managing the change.

Even if a private DC might be the most efficient option for a company today, technology changes and its data and storage requirements may mean that a public cloud solution will be optimal in the future. Implementing a hybrid cloud solution may be one way to ease migration and distribute transition costs over time but every company's circumstances are unique.

Therefore, identifying a company's PCE over time and planning in advance are essential to ensuring digital efficiency now and into the future.

[Maximize your digital efficiency. Contact Alvarez & Marsal to get started.](#)

ABOUT THE AUTHOR



ADAM RICHARDS

SENIOR DIRECTOR, CORPORATE PERFORMANCE IMPROVEMENT

arichards@alvarezandmarsal.com

Adam Richards is a Senior Director with Alvarez & Marsal in San Francisco. He specializes in architecture and large systems design. With more than 25 years of industry experience, his primary areas of concentration are enterprise architecture, cloud migration and highly-available and scalable systems design.

This article was produced with research and support from the A&M Insight Center, which serves to provide A&M professionals and clients with relevant, industry-specific, actionable insights derived through proprietary studies and research.



ABOUT ALVAREZ & MARSAL

Companies, investors and government entities around the world turn to Alvarez & Marsal (A&M) when conventional approaches are not enough to make change and achieve results. Privately held since its founding in 1983, A&M is a leading global professional services firm that provides advisory, business performance improvement and turnaround management services.

With over 3000 people across four continents, we deliver tangible results for corporates, boards, private equity firms, law firms and government agencies facing complex challenges. Our senior leaders, and their teams, help organizations transform operations, catapult growth and accelerate results through decisive action. Comprised of experienced operators, world-class consultants, former regulators and industry authorities, A&M leverages its restructuring heritage to turn change into a strategic business asset, manage risk and unlock value at every stage of growth.

Follow us on:



When action matters, find us at: alvarezandmarsal.com

ALVAREZ & MARSAL