WHITE PAPER

Yes, the Cloud Is a Catalyst; It's Also a Competitive Benchmark

STREAM DATA CENTERS

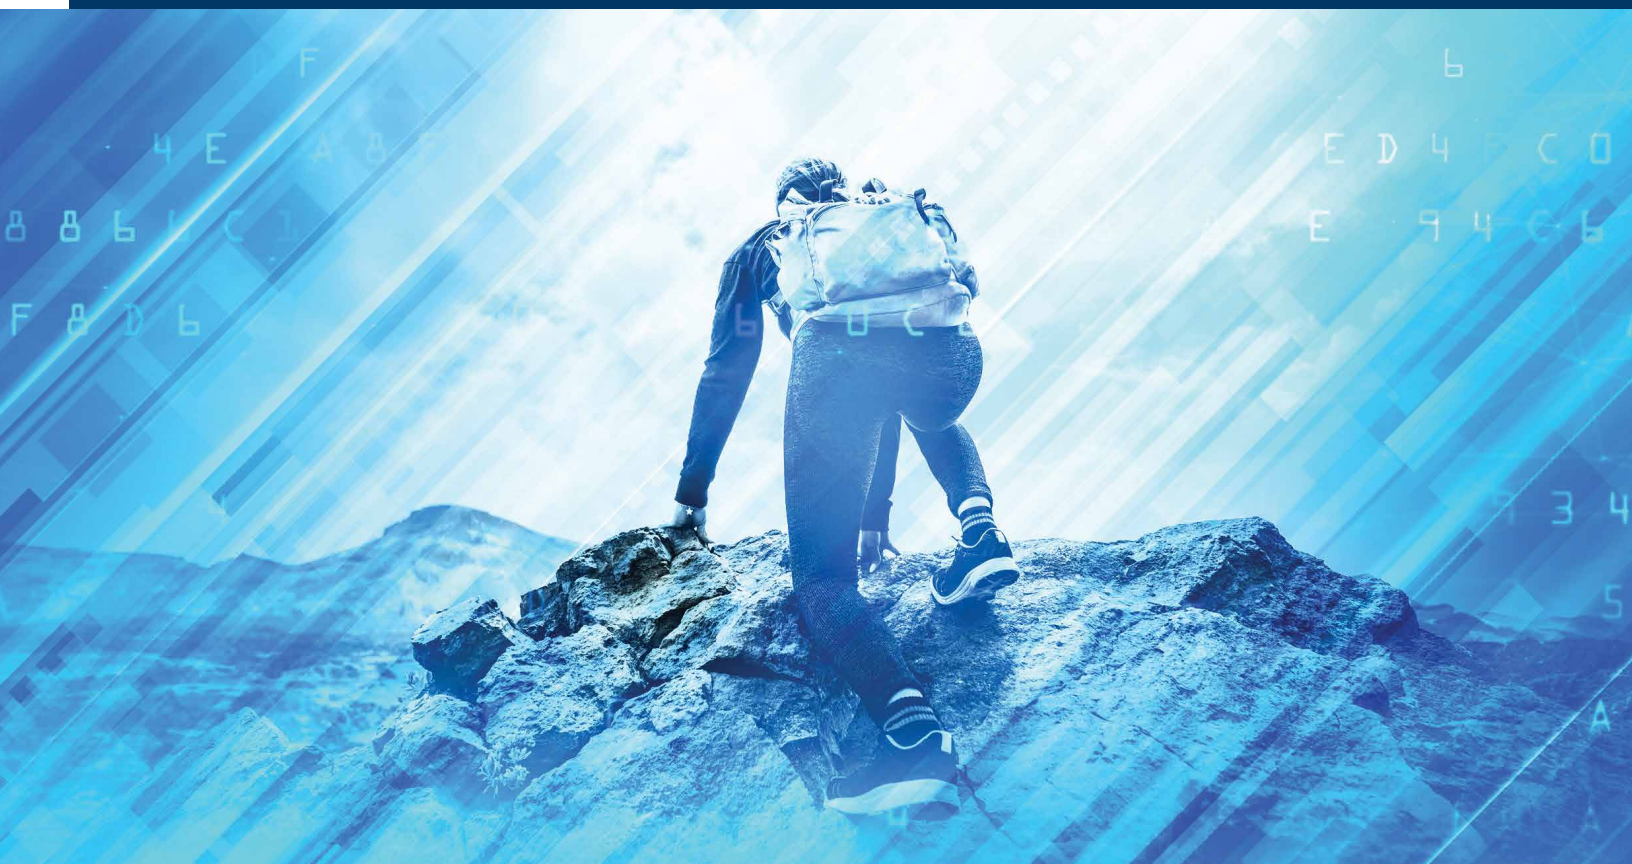

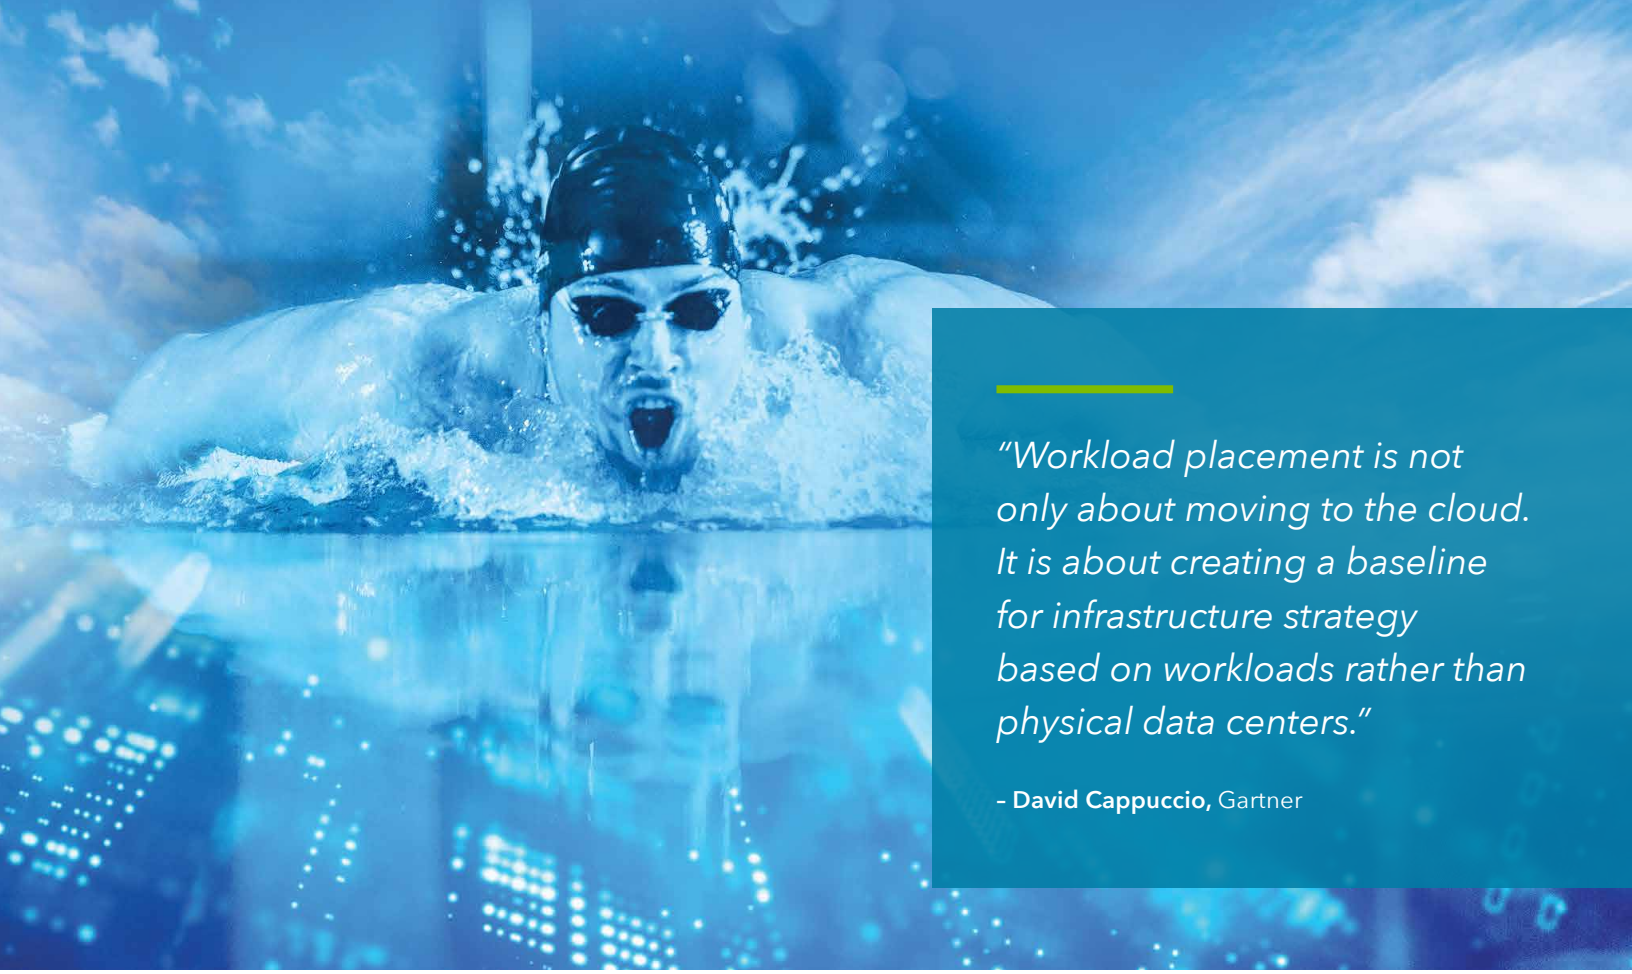

"Workload placement is not only about moving to the cloud. It is about creating a baseline for infrastructure strategy based on workloads rather than physical data centers."

– David Cappuccio, Gartner

1. Introduction

The answer to the question, "Where should this application live?" is always "It depends." It depends on the end users (in particular, where they are), on the application architecture, and on regulatory constraints governing the application. It also depends on the total cost of ownership, flexibility, performance, and 'righteousness' offered by the various prospective application hosts.

Approaching infrastructure strategy decisions from the perspective of the workload represents a significant shift for most enterprises. Brian Moynihan, chairman of the board and CEO of Bank of America explained it well: "We don't need to own the hardware, we just need to find out who can provide the right way."

Where Workloads Live

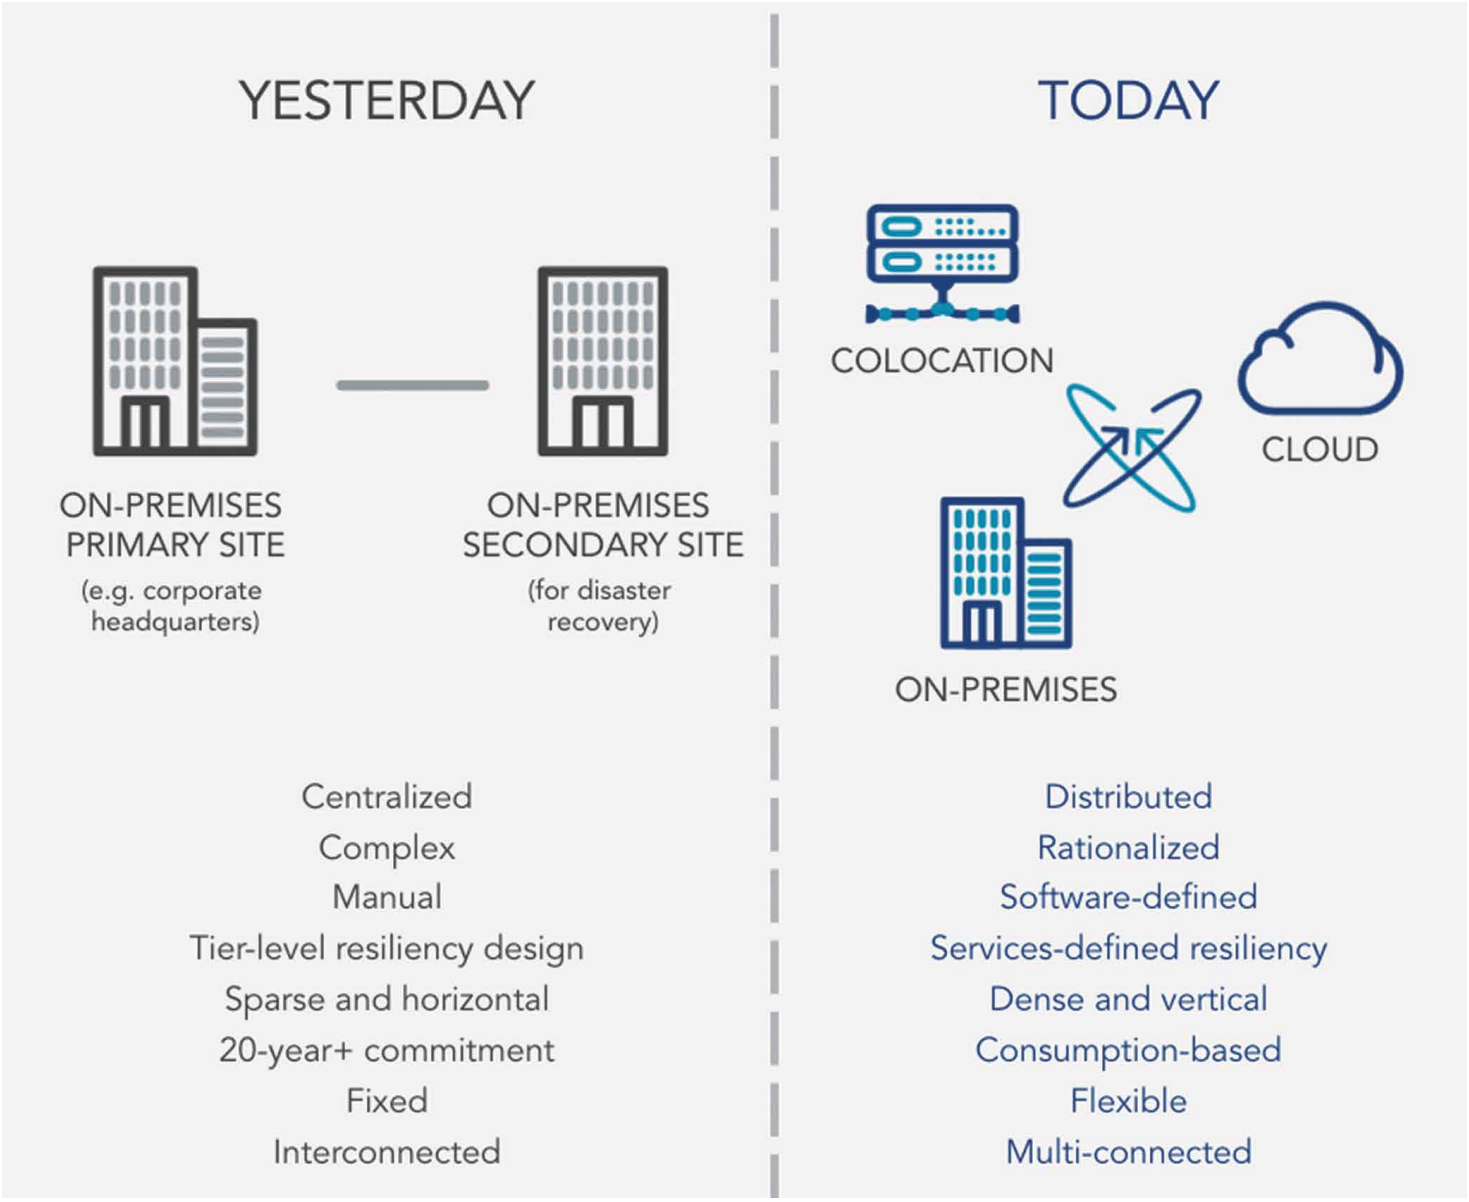

Source: Adapted from Gartner, **The Future of Enterprise Data Centers–What's Next,** April 2019

In opening more possibilities for where applications can live, the cloud has served as a catalyst for evolution in on-prem and colocation data centers. But the cloud is not only a catalyst. Because it's a great option for most applications, it has become a competitive benchmark. Enterprise IT leaders are no longer only competing against the TCO, flexibility, performance, and righteousness standards set by their direct competitors. They also must be good stewards of their company's resources by deploying and managing infrastructure in facilities that meet the same benchmarks as the cloud service providers do.

1.1 THE RISE OF THE CLOUD

Enterprise investment in incremental capacity within internal data centers has continued to shrink. At the same time, investment in cloud and colocation continues to rise. These shifts are driven by the unparalleled flexibility offered by cloud service providers as well as several challenges associated with in-house data center facilities. As on-prem facilities continue to age they lose the ability to keep pace with changes in density and energy efficiency. Budgeting becomes difficult due to the significant capital expenditures and managing what can be an unpredictable expansion forecast.

> *"In 2019, enterprise tech crossed a significant line. For the first time, enterprises spent more annually on cloud infrastructure services than on data center hardware and software."*
>
> **– Andy Patrizio,** Network World

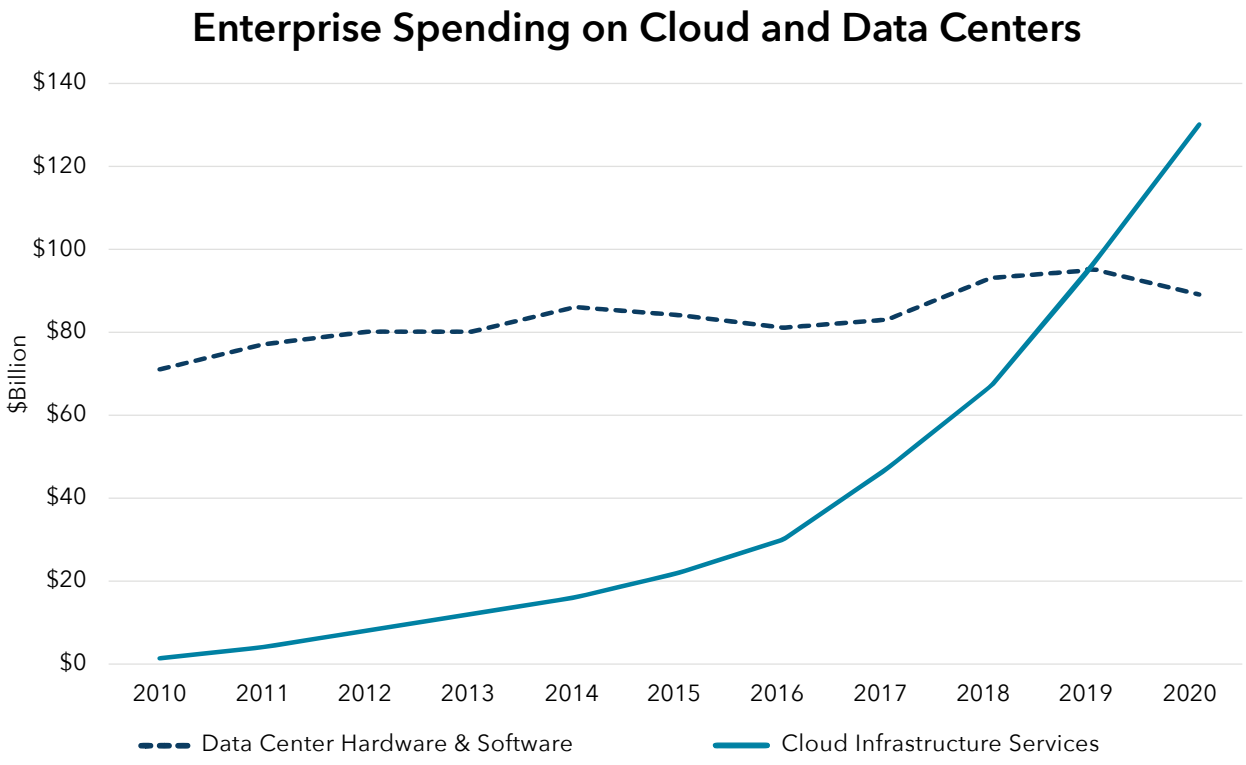

Source: Synergy Research Group via **FierceTelecom**

Clearly, the last decade has seen a dramatic rise of spending on cloud infrastructure services at the same time as spending on data center hardware and software remains relatively flat. (The latter includes spending on servers, storage, networking, security and associated software for both on-premises and colocation environments, as colocation tenants also purchase their own hardware and software even as they move from a capex to opex model at the facility level.)

Enterprise adoption of the cloud is a catalyst for evolution in colocation business models, designs, and locations. In order to serve the highly efficient cloud service providers, colocation providers evolved to deliver highly resilient, flexible and cost-efficient data center services. The ability to serve the hyperscale segment well is a significant benchmark for a colocation provider—which benefits both hyperscale and enterprise customers.

1.2 THE METRICS THAT MATTER

Much like the cloud providers, efficient enterprises understand the impact of geographic location on the metrics that matter. Cloud data centers are located where those metrics are optimized. Now, enterprises can also put their applications in the best locations irrespective of where headquarters is. Similarly, competitive cloud and colocation markets enable enterprises to choose from the best (i.e., lowest TCO, most flexible, highest performance, most efficient) provider in any given market.

So for deciding where the application should live, and with whom, the cloud is a benchmark on the metrics that matter:

- Total cost of ownership (which is impacted by land cost and/or rent, energy cost, network cost, labor cost, and tax incentives).
- Flexibility (especially with regard to scale).
- Performance (which is impacted by resiliency and connectivity).
- 'Righteousness' (namely, energy efficiency and use of renewables) – this is about doing the right thing for the planet and also maintaining a positive market perception.

2. TOTAL COST OF OWNERSHIP (TCO)

It used to be that proximity (e.g., to corporate headquarters) was an enterprise infrastructure leader's top criterion in site selection decisions. Hyperscale data center operators were the first to eschew that tradition in favor of application-level resiliency relying on multiple sites within a particular geographic region. Because of globally fantastic network performance and cost, TCO can be a much more important decision driver. The application no longer needs to be proximate to its end-users to deliver a great user experience.

A list of the top data center markets in the U.S. today makes clear the divergence between proximity to corporate headquarters and data center demand.

What Drives Data Center Demand?
It's Not Proximity to Corporate HQ

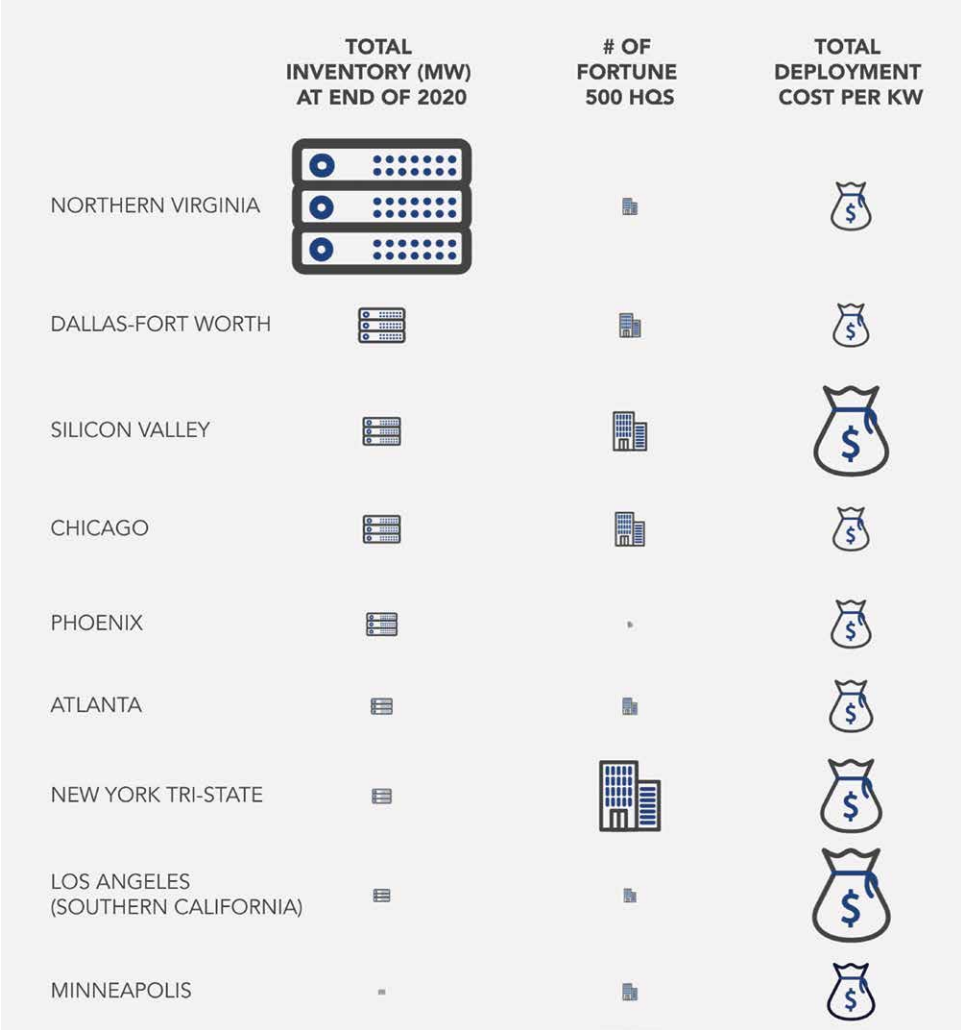

Note: Total deployment cost includes rent and power and accounts for tax incentives (where applicable)
Source: CBRE Research 2020; Cushman & Wakefield 202; Fortune Magazine 2019; author's calculations

Market	Total Inventory (2nd half 2020)	Inventory Growth 2015-2020	Total Deployment Cost per kW	# of Fortune 500 HQs
Northern Virginia	1376.7 MW	227.1%	$127.56	16
Dallas-Fort Worth	360.9 MW	219.4%	$115.14	23
Silicon Valley	292.1 MW	117.3%	$238.48	38
Chicago	288.5 MW	121.6%	$116.93	33
Phoenix	238.9 MW	67.9%	$127.98	5
Atlanta	163.8 MW	77.2%	$137.66	16
New York Tri-state	149.1 MW	19.7%	$194.25	68
Los Angeles (Southern California)	127.1 MW	nr	$246.78	13
Minneapolis	52.0 MW	nr	$151.44	16
Portland	nr	nr	$163.39	1

Note: Total deployment cost includes rent and power and accounts for tax incentives (where applicable)
Sources: **CBRE Research** 2020; **Cushman & Wakefield** 2021; Fortune Magazine 2019; author's calculations

Clearly it's not only total cost of ownership (TCO) that drives demand; Silicon Valley, for example, is a relatively high-cost market where data centers continue to flourish. The point is that there is no longer a correlation between the location of large corporate headquarters and the location of those enterprises' data centers.

The rise of Ashburn and decline of the New York tri-state area as data center hubs serve as perfect illustrations of the prominence of TCO over proximity as a location determinant—now that low-latency, low-cost network connectivity is so easy to come by.

> *"Proximity to New York City is no longer relevant in today's trading environment."*
>
> **– Hope Jarkowski,** Co-Head of Government Affairs for the NYSE

TWO DATA CENTER MARKETS IN CONTRAST: THE RISE OF NORTHERN VIRGINIA AND DECLINE OF NEW JERSEY

Northern Virginia

An "unmatched" level of connectivity is why the first telecom and internet companies were drawn to the Northern Virginia area now known as Data Center Alley. It was there that a group of network providers formed MAE-East, one of the first internet exchanges, in 1992. A year later, a grant from the National Science Foundation made MAE-East one of the United States' four original Network Access Points.

With that connectivity, alongside relatively low-cost land and power (and now, tax incentives), and close proximity to vibrant labor markets, by the late 1990s Northern Virginia had become the #1 data center market, a title it has held ever since.

New Jersey

Right across the river from the financial epicenter of the world, New Jersey has been the go-to location for Wall Street firms' data centers. Locating in Manhattan itself was prohibitively expensive from both a real estate cost and a power cost perspective, but until fairly recently network connectivity was not yet so robust that Wall Street firms could locate their data centers anywhere. They had to be close, so New Jersey it was.

That arrangement continues to persist, but now Wall Street firms have more options. Network performance has improved so much it's no longer necessary for the majority of applications to be in close geographic proximity to Wall Street. The New York Stock Exchange proved that to be true in late 2020 when for four days it operated with its production servers in Chicago rather than New Jersey.

The NYSE's test was a response to a proposed tax on high-volume electronic trading in New Jersey. Nasdaq's response to the tax proposal was a visit to Texas and meetings with the Lone Star governor to discuss possible incentives for a move to Dallas. The New Jersey governor has since backed away from his tax proposal, but the point has been made: even Wall Street firms that depend on extremely low latency no longer have to be in New Jersey to get it.

The most significant variable core components of TCO include land cost (applicable for self-build and build-to-suit) and rent (for colocation), energy cost, network cost, labor cost, and tax incentives. So it should come as no surprise that the data center markets where data center demand—from cloud providers and enterprises alike—is highest are the markets with the lowest land cost/rent, lowest energy, network, and labor costs, and best tax incentives.

2.1 LAND COST OR RENT

2.1.1 Land Cost

"Affordable land eases the cost of entry to a market and likely indicates considerable availability. Large sites are increasingly prized by data center operators for development across several phases."

– Cushman & Wakefield

The top seven U.S. markets by land price, as ranked for 2020 by **Cushman & Wakefield**, include:

1. Columbus
2. Phoenix
3. Dallas
4. Nashville
5. Atlanta
6. Denver
7. Chicago

2.1.2 Rent

Market	Median Rental Rate (kW/mo) 1-4 MW	5-10 MW
Silicon Valley	$130	$105
Southern California	$127.50	$107.50
Atlanta	$115	$105
New York Tri-state	$115	$105
Minneapolis	$115	$100
Portland	$110	$100
Phoenix	$105	$95
Chicago	$100	$90
Dallas-Fort Worth	$97.50	$87.50
Northern Virginia	$85	$80

Source: CBRE

2.2 ENERGY COST

Power is typically about 40% of a data center's total operating cost. So even a relatively small difference in power cost can make a big difference in TCO.

Market	Median Power Rate (per kWh)	Annual Power Cost for 10 MW Critical Load
Los Angeles	$0.120	$10,512,000
Silicon Valley	$0.115	$10,074,000
New York Tri-state	$0.091	$7,993,500
Minneapolis	$0.069	$6,044,400
Chicago	$0.060	$5,256,000
Northern Virginia	$0.060	$5,256,000
Phoenix	$0.060	$5,256,000
Portland	$0.060	$5,256,000
Dallas-Fort Worth	$0.046	$4,051,500
Atlanta	$0.043	$3,723,000

Note: This does not account for PUE
Source: CBRE

2.3 NETWORK COST

Leveraging interconnection points for peering is one way to reduce network transit costs. Same goes for leveraging a cloud on-ramp for dedicated (faster, more secure, and lower cost) cloud access. And the more network service providers there are in a given market, the more likely it is for cost to be low and performance to be high. (Kind of like Costco, markets with many network connectivity options tend to offer higher performance at a lower cost than markets with fewer options.)

2.4 LABOR COST

For most data center operators, labor costs are not a relatively significant share of operating expenditures. Nevertheless, labor cost variations in different markets do make some difference on the bottom line; for example, the total annual labor cost for a data center with 100 employees is almost $5 million higher in Silicon Valley than in Dallas.

Market	Employer Costs per Hour Worked
Silicon Valley	$57.99
New York Tri-State	$45.29
Portland (West Census Region)	$40.19
Northern Virginia (Washington-Baltimore-Arlington MSA)	$39.45
Chicago	$39.31
Minneapolis	$38.47
Los Angeles	$38.42
Phoenix	$36.75
Atlanta	$35.66
Dallas-Fort Worth	$34.48

Source: Bureau of Labor Statistics

2.5 TAX INCENTIVES

Well over half of the states in the U.S. (32) now offer data center incentives—often sales—or property-tax abatements for long-term investment.

"When the box is checked for real estate, taxes and incentives end up swaying the business case or level the playing field for communities on a data center's short list."

- CBRE

U.S. States with Data Center Tax Incentives

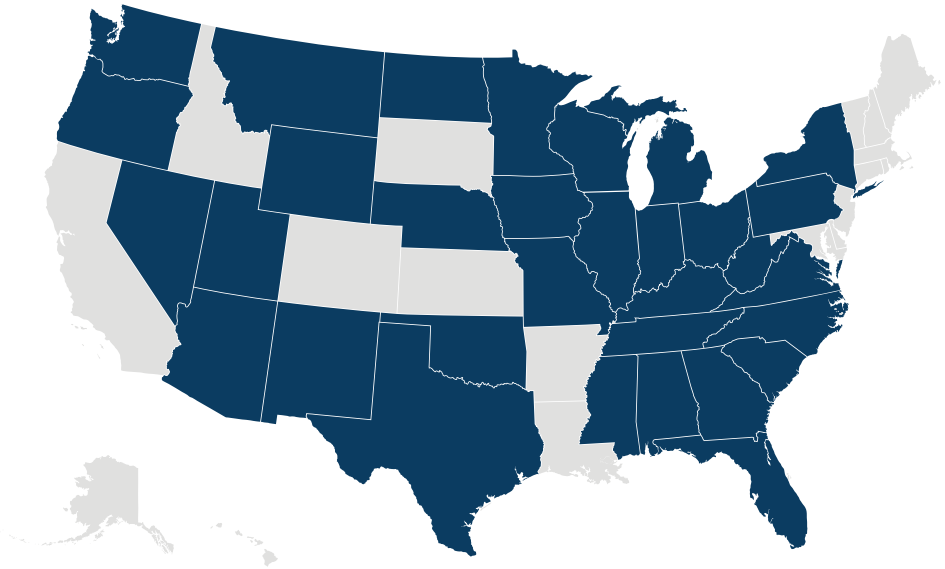

Sources: Cushman & Wakefield and Northern Virginia Technology Council

Market	Tax Incentive	Tax Impact on Data Center Deployment Costs
Chicago	10-year exemption on state and local sales tax (10.25% total)	27%
Phoenix	10-year or 20-year exemption from state, county, and city sales tax (8.8% total)	22%
Dallas-Fort Worth	10-year or 15-year exemption of state sales and use tax on tangible personal property (6.25%)	18%
Minneapolis	20-year state sales tax exemption on equipment and electricity (6.875%)	16%
Northern Virginia	State sales and use tax exemption (5.3%)	15%
Atlanta	State sales tax exemption on data center equipment (4%)	11%
New York Tri-state	New Jersey does not provide statewide tax incentives for data centers; New York provides a state sales tax exemption (4%)	8%
Portland	Oregon does not have a state sales tax; in addition, property taxes are abated for up to 5 years on qualifying capital assets within designated regional enterprise zones (approx. 1.5%)	4%
Los Angeles	None	na
Silicon Valley	None	na

Note: Tax impact percentages illustrate the impact of tax incentives on hardware and software purchases expressed as data center deployment savings
Sources: Noted within table as hyperlinks

"But cloud computing is bringing change to how companies approach uptime, introducing architectures that create resiliency using software and network connectivity. This strategy, pioneered by cloud providers, is creating new ways of designing applications."

– Rich Miller, Data Center Frontier

3. Flexibility

Flexibility is perhaps even more important than TCO to an enterprise IT infrastructure leader's ability to compete with the cloud. The cloud enables enterprises to "pay by the drink" (pay for just the capacity they use) and increase (and decrease) capacity on-demand. That kind of flexibility is hard to match on-premises; making standalone infrastructure investments of any magnitude very hard to get right in an uncertain world.

Indeed, the ability to right-size deployments versus trying to forecast 15 or 20 years into the future has many progressive IT infrastructure leaders working to 'get out of the data center business' by smartly leveraging a combination of cloud, internal data center and colocation—to drive optimal efficiency without tying up resources unnecessarily.

That was the case for one Fortune 500 enterprise notorious for its focus on cost efficiency. Speaking about the company's decision to move out of on-premises facilities into the cloud, the CIO said it wasn't lower TCO that drove the decision, but flexibility. It was the ability to buy right-size capacity today without a big capital commitment, and adjust that capacity as the business's needs changed.

In many cases, flexibility is about being able to leverage the cloud for short-term bursts. That was the case for the NFL in 2020. When COVID-19 rendered the organization unable to hold its annual draft live, it leveraged the cloud for the burst capacity required to livestream its virtual draft.

What's driving enterprise migration to the cloud is "not necessarily cost savings… The real benefits of cloud computing are in agility, scalability, and future potential."

– Eric Knorr, Editor in Chief, IDG

4. Performance

4.1 RESILIENCY

It used to be that enterprises relied on data center infrastructure over-engineering to ensure availability. Now that many applications support multi-site replication, resiliency is built into the application rather than the physical infrastructure. So data centers can be less complex, and also incredibly resilient and efficient.

"Delivering uptime has always been the prime directive for data centers. The industry was created to ensure that mission-critical applications never go offline. That goal has typically been achieved through layers of redundant electrical infrastructure, including uninterruptible power supply (UPS) systems and emergency backup generators.

It's like the difference between a battleship and a cruiser. It's not a difference of size, as cloud data centers are often huge (hence the term, hyperscale). Rather, it's about agility and resiliency. Like facilities within a particular geography that support multi-site replication, a cruiser doesn't have to be so heavily fortified and complex because it is so agile.

So to compete with the kind of resiliency offered by the cloud (while avoiding the cost and complexity of over-engineering) modern on-premises and colocation data centers are less complex while remaining highly resilient and efficient. For example, Stream's most current generation of data centers has the lowest per megawatt cost in our history, yet it still meets IEEE standards for delivering 99.9999% (six 9s) of uptime.

In addition to less complexity in the data center, the rise of application-based resiliency has driven a different kind of focus on location. In most approaches, supporting multi-site replication requires multiple data centers within certain geographic parameters. So the location decision becomes not only which market to go to, but what location(s) within any given market as well.

4.2 NETWORK PERFORMANCE

It used to be that proximity to end users was essential for delivering a great user experience. But cloud providers have demonstrated how, for most applications, proximity (e.g., to corporate headquarters) is not necessary for excellent end user performance—as long as network connectivity is strong. (That's why cloud strongholds are not always where corporate headquarters or even populations are concentrated; location decisions can be made based on TCO, resiliency, and other factors rather than proximity.)

Instead of proximity, then, the competitive benchmark is network connectivity—including interconnection options and dedicated cloud access. As the prominence of Northern Virginia as a data center market makes clear, connectivity is king.

> *"While a data center requires access to just one network to function, the greater the number of networks means the more the facility can prosper. Faster, denser fiber generally allows for the lowest latency and better overall network reliability, particularly when functioning with major cloud services."*
>
> **– Cushman & Wakefield**

Connectivity in the U.S.

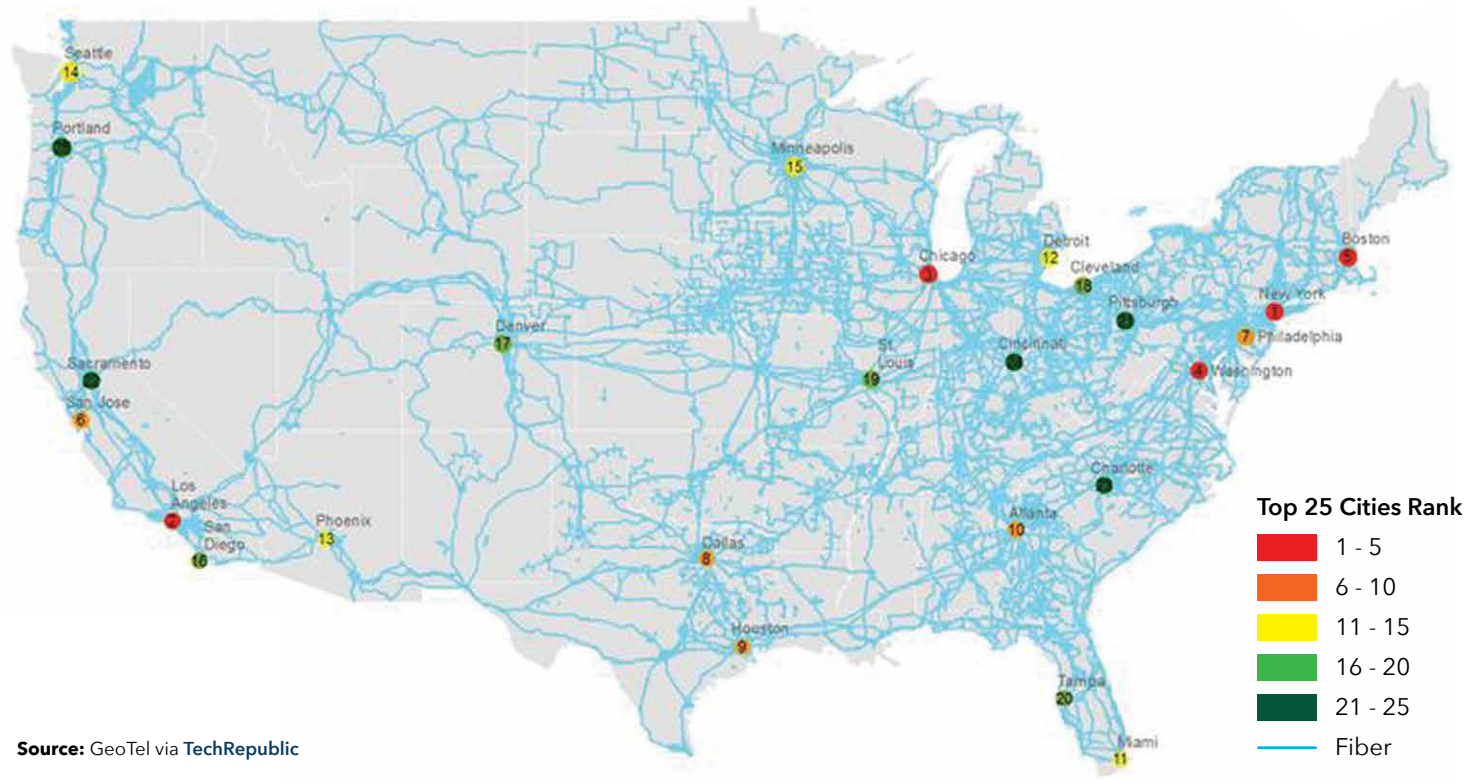

Source: GeoTel via **TechRepublic**

5. Righteousness

In the 'righteousness' category we include energy efficiency and renewable energy. There are many reasons why cloud sets the benchmark on these metrics. Greater energy efficiency, all else equal, means lower TCO. And while renewable energy is still more expensive in many markets, the holy grail of cost parity—or even, for renewable energy to be less expensive—is attainable. The rationale for cloud providers, and enterprises, to pursue 'righteousness' often also includes corporate responsibility commitments such as carbon reduction goals. Market perception matters, too. (For data center operators, getting called out by Greenpeace is the stuff of nightmares.)

5.1 ENERGY EFFICIENCY

Energy efficiency impacts TCO. All else equal, a data center with a lower PUE is a lower cost data center. Energy procurement strategies (including for renewables) is another lever to affect TCO, and cloud providers have very sophisticated such strategies (and immense buying power).

So in order to remain competitive with cloud, colo providers have had to become more energy efficient as well. And indeed they have.

According to the most recent data from Uptime Institute, average PUE has declined considerably since 2006, though efficiency gains have flattened out over the last decade. The Uptime Institute figures are based on surveys of global data centers, of varying ages, ranging in size from 1 MW to over 60 MW.

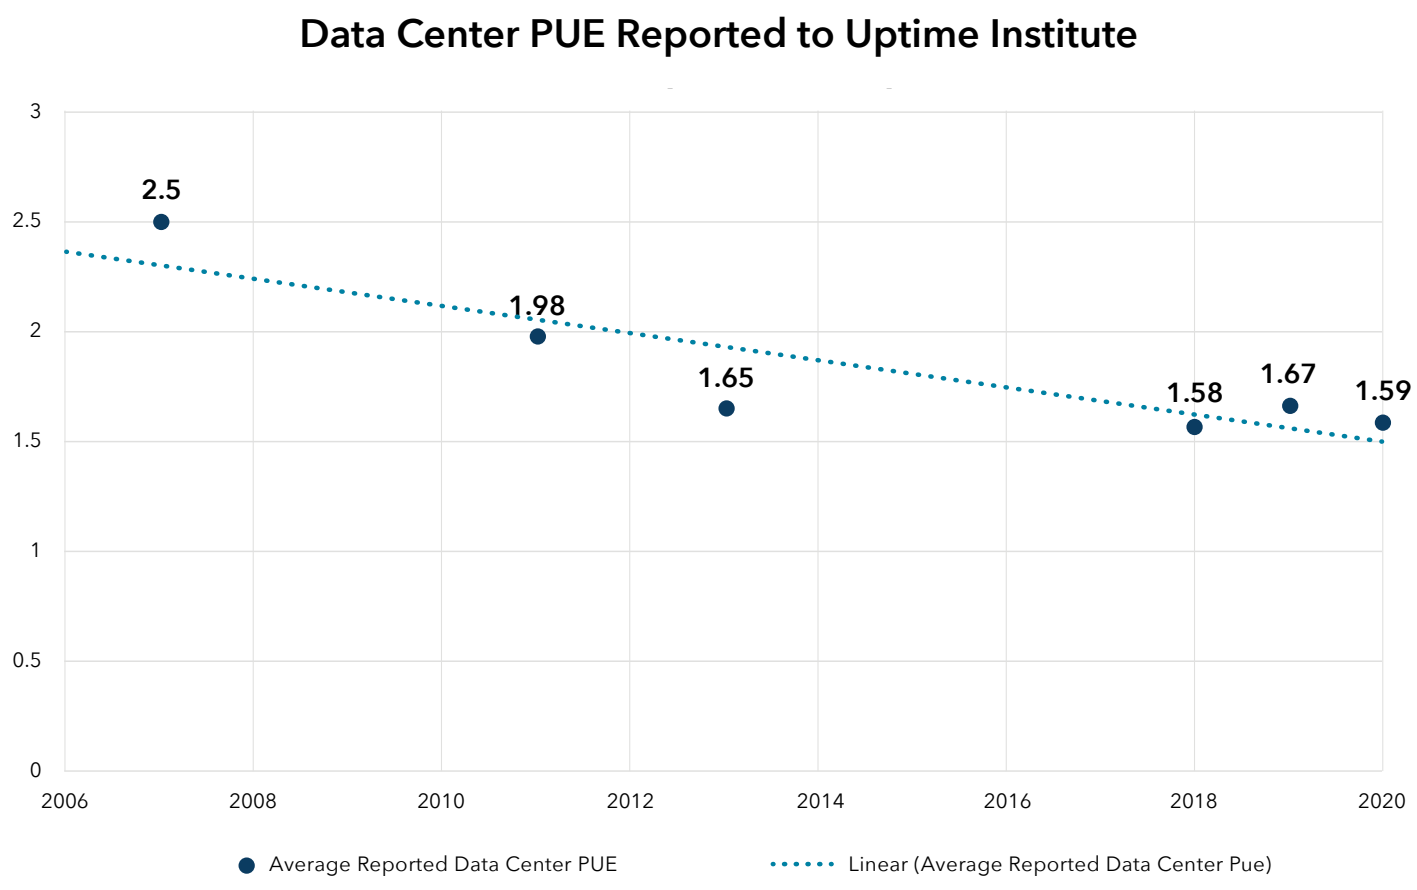

Source: Uptime Institute

There is no question that cloud providers are leading the massive improvements in data center energy efficiency, as cloud data center PUEs are typically much lower than average. (Often below 1.2, some driving toward 1.1.)

That's why, even as demands on the data center have increased by orders of magnitude, data center energy consumption has increased much less. Data centers are becoming more efficient.

Data Centers Are Doing More With Less
(RIGHTEOUS!)

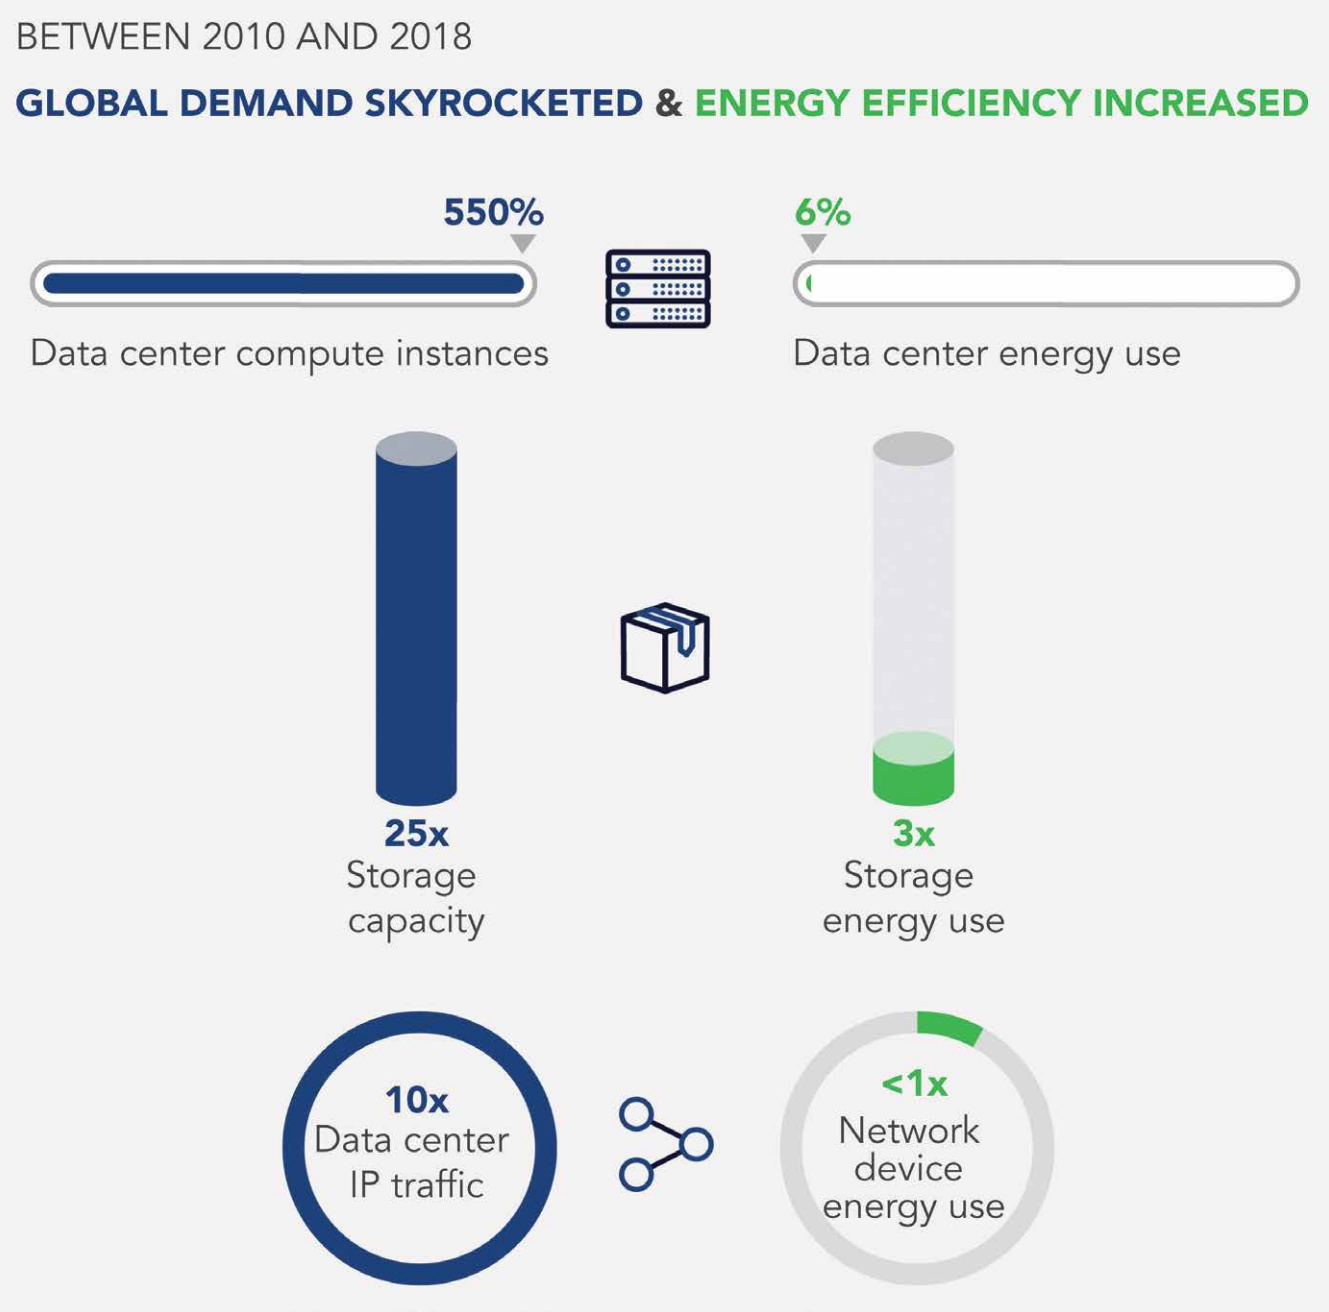

Source: Science

"As large corporations and local governments pursue their own sustainability goals, data centers will be required to follow suit to meet future regulatory concerns and obtain further business. Those facilities that maintain a low power usage effectiveness [PUE], use water sparingly and utilize renewable energy will benefit both in cost savings and partners."

– Cushman & Wakefield

5.2 RENEWABLE ENERGY

Most of the cloud providers have aggressive sustainability goals. In addition, they recognize the importance of sustainability to their customers. So cloud providers are among the leading buyers of renewable energy.

Renewable Energy Deals

Top 5 Buyers

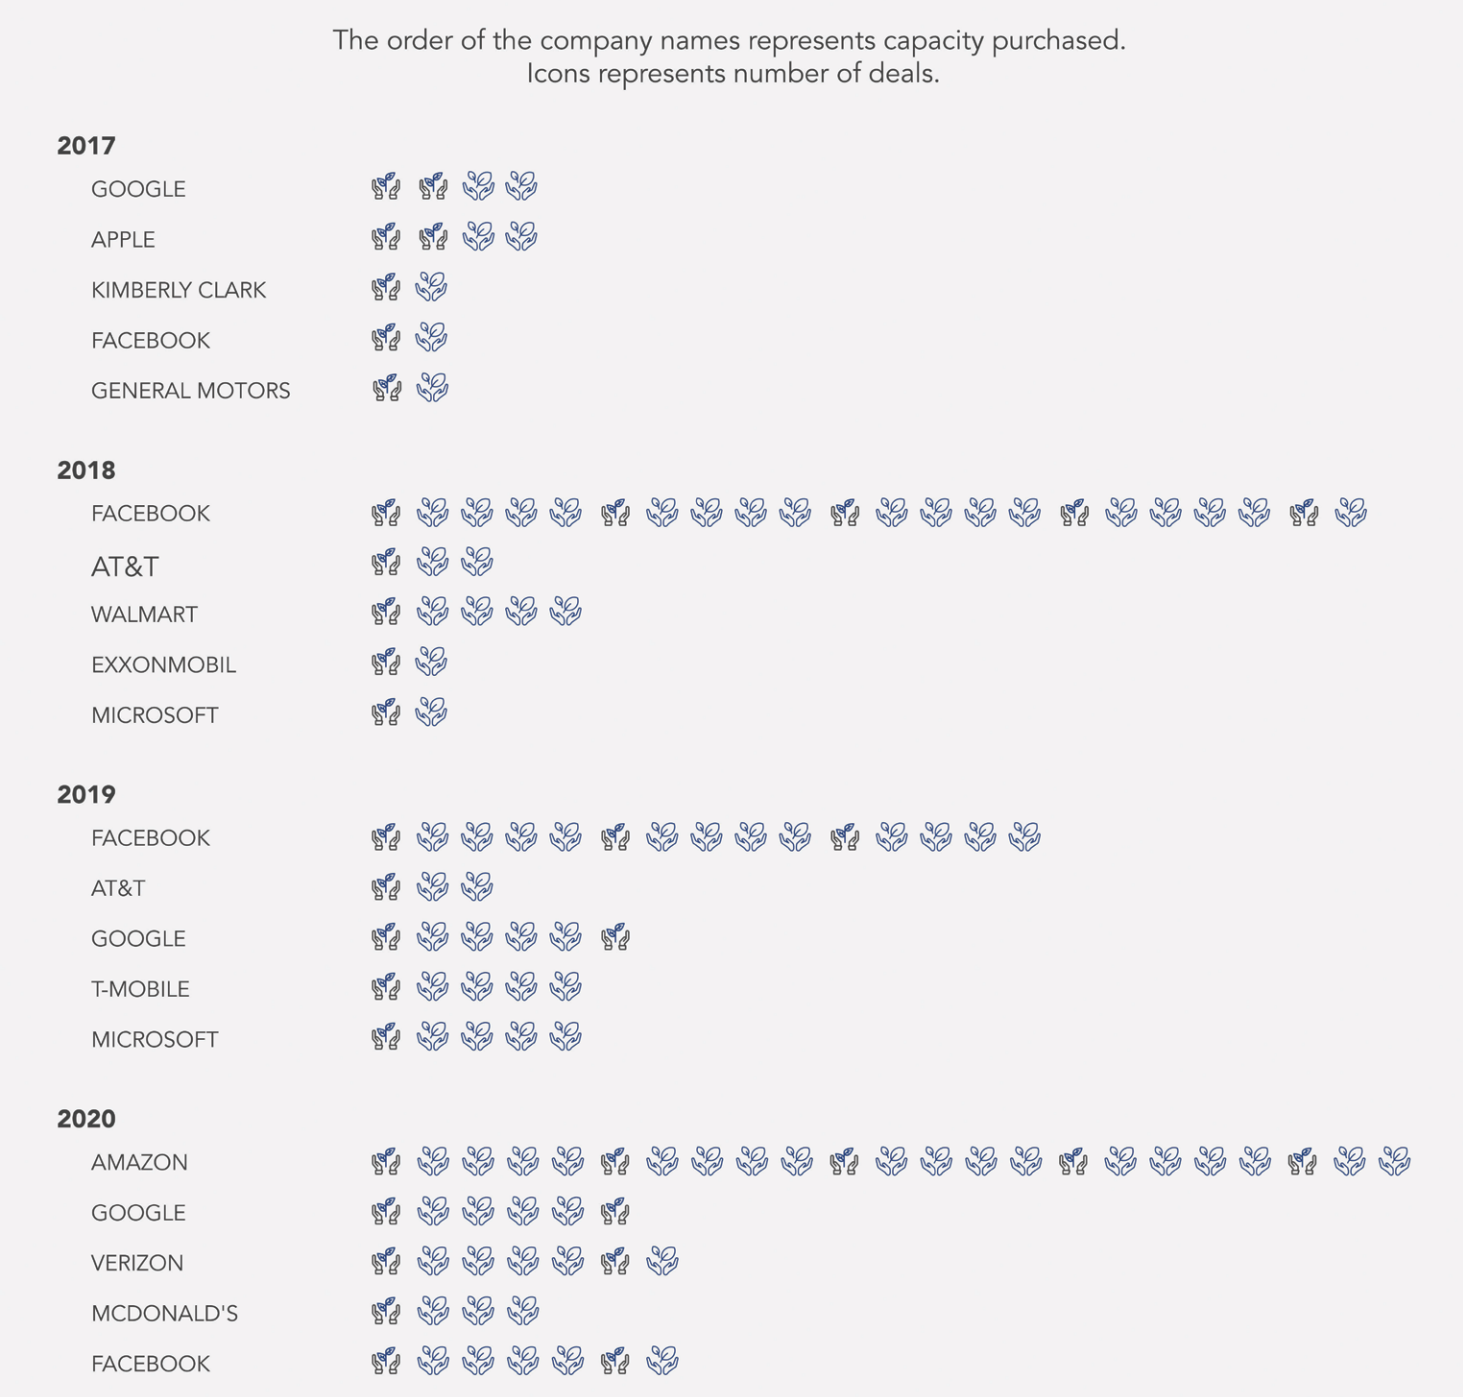

Source: REBA Deal Tracker

Deals are publicly announced corporate Power Purchase Agreements, green power purchases, green tariffs, and outright project ownership. Excludes off-site generation, deals with operating plants, and deals meant to meet RPS requirements.

Given that an enterprise can achieve sustainability goals through the cloud, any on-prem or colo alternative will be benchmarked against the righteousness of the cloud. That's why access to renewable energy is a significant consideration factor for many enterprise IT leaders when making the "where should this application live?" decision.

Utility Scale Renewable Energy Generation

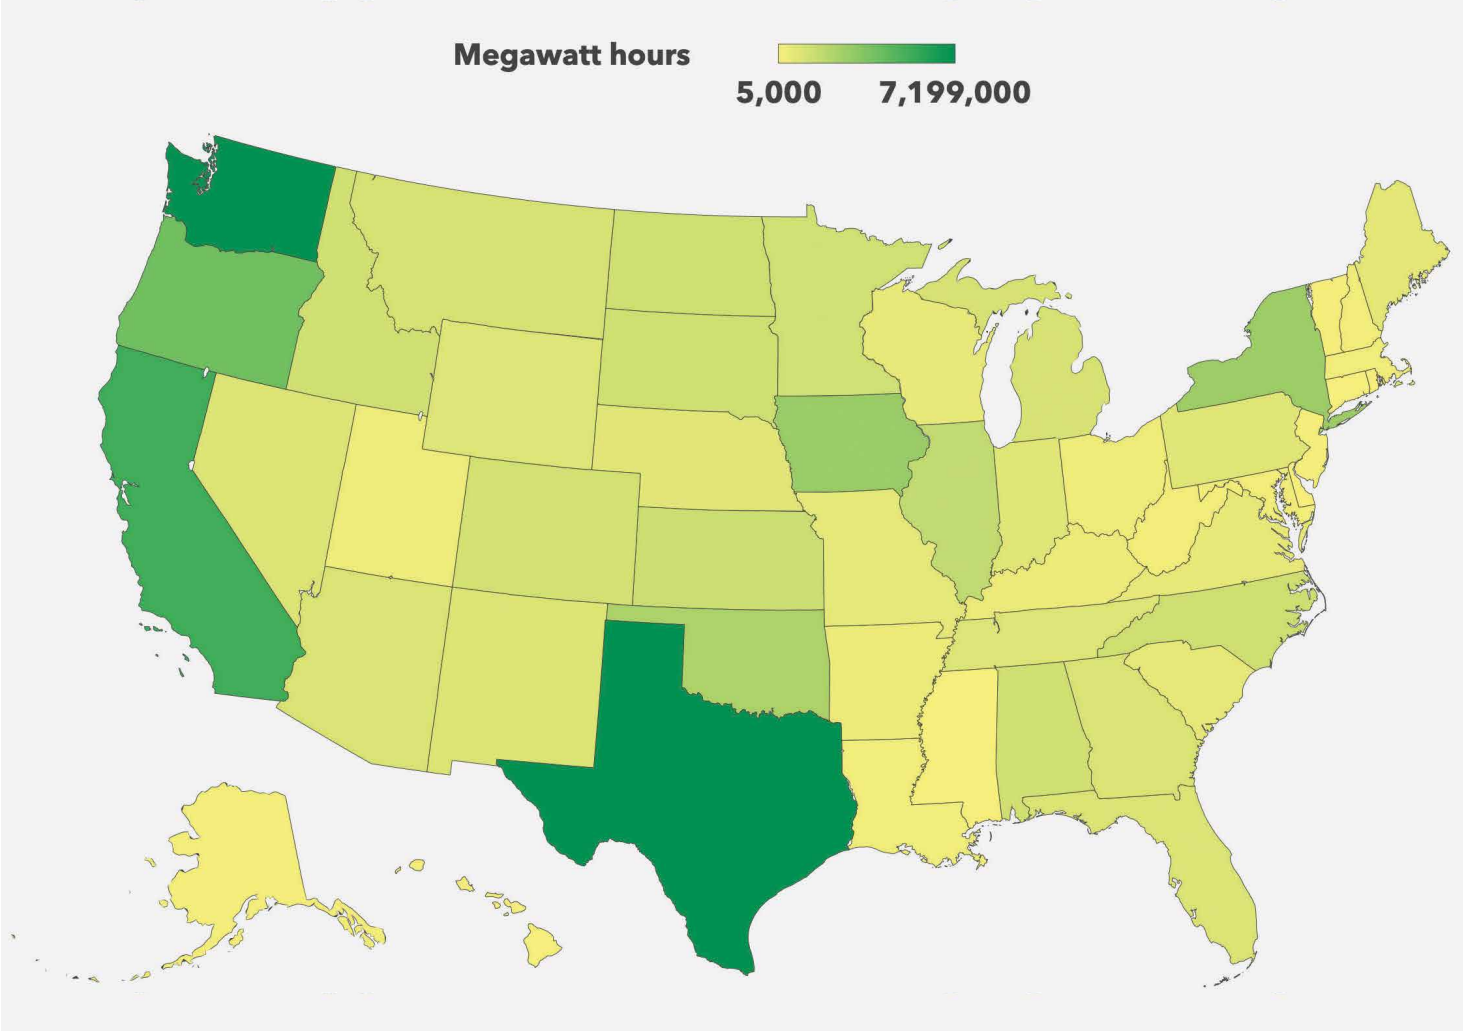

Note: Includes hydroelectric
Sources: Energy Information Administration

6. Bottom Line: You're competing with the cloud

With the rise of the cloud, enterprise IT leaders now have a good third option in the "Where should this application live?" decision making process. Uptime and security are table stakes. The pressure is on enterprise IT infrastructure leaders to house applications where TCO and performance are optimized, while remaining flexible to quickly scale up and down, and to do it righteously.

About Stream Data Centers

Stream Data Centers has provided premium data center services since 1999, with 90% of inventory leased to Fortune 500 customers. To date, the company has acquired and developed 24 data center facilities nationally while leadership has remained consistent for all 22 years.

Stream speculatively develops turnkey data centers for Hyperscale and Enterprise users. Additionally, Stream develops Build to Suit Data Centers and operates an Energy Services team with a focus on low cost, renewable energy solutions. All of Stream's facilities feature carrier-neutral, low latency connectivity to network and public cloud providers. Above all, Stream is dedicated to improving the data center experience through exceptional people and service.

Stream Data Centers is a subsidiary of **Stream Realty Partners**, a full service commercial real estate investment, development and services company. Stream Realty Partner's portfolio of 300+ million square feet of commercial properties and 900+ real estate professionals, gives Stream Data Centers outstanding market insight and reach.

STREAM
DATA CENTERS

FOR MORE INFORMATION

Learn more about Stream's build-to-suit data center development and colocation services at **www.streamdatacenters.com/.**