Atlantic Council
SCOWCROFT CENTER
FOR STRATEGY AND SECURITY

CYBER STATECRAFT
INITIATIVE

FOUR MYTHS ABOUT THE CLOUD

THE GEOPOLITICS OF CLOUD COMPUTING

Trey Herr

Scowcroft Center for Strategy and Security

The **Scowcroft Center for Strategy and Security** works to develop sustainable, nonpartisan strategies to address the most important security challenges facing the United States and the world. The Center honors General Brent Scowcroft's legacy of service and embodies his ethos of nonpartisan commitment to the cause of security, support for US leadership in cooperation with allies and partners, and dedication to the mentorship of the next generation of leaders.

Cyber Statecraft Initiative

The **Cyber Statecraft Initiative** works at the nexus of geopolitics and cybersecurity to craft strategies to help shape the conduct of statecraft and to better inform and secure users of technology. This work extends through the competition of state and non-state actors, the security of the internet and computing systems, the safety of operational technology and physical systems, and the communities of cyberspace. The Initiative convenes a diverse network of passionate and knowledgeable contributors, bridging the gap among technical, policy, and user communities

Atlantic Council
SCOWCROFT CENTER
FOR STRATEGY AND SECURITY

CYBER STATECRAFT
INITIATIVE

FOUR MYTHS ABOUT THE CLOUD

THE GEOPOLITICS OF CLOUD COMPUTING

Trey Herr

ISBN-13: 978-1-61977-115-4

Cover: Vector isometric data center. Server room with hot and cold aisle containment, generator, UPS and battery rooms, CRAC unit with compressor, Network operations center and other equipment. Credit: tarras79/iStock illustration.

August 2020

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Cloud computing providers are more than companies—they govern vast utility infrastructure, play host to digital battlefields, and are magnificent engines of complexity. Cloud computing is embedded in contemporary geopolitics; the choices providers make are influenced by, and influential on, the behavior of states. In competition and cooperation, cloud computing is the canvas on which states conduct significant political, security, and economic activity.

This paper offers a brief primer on the concepts underneath cloud computing and then introduces four myths about the interaction of cloud and geopolitics. First, that all data is created equal – a discussion of how cloud providers build and operate these data intensive services and the impact of debates about how and where to localize these systems and their contents. Second, that cloud computing is not a supply chain risk – cloud providers play host to some of the most remarkable security challenges and widely used technical infrastructure in the world, their decisions impact the supply chains of millions of users and entail management of risk at sometimes novel scale. Third, only authoritarian states distort the cloud – a pernicious myth and one that continues to hold back a cogent Western stategy to defend the open internet and threatens to upend the economics supporting the public cloud. Fourth, that cloud providers do not influence the shape of the internet – this final section highlights both risk and opportunity for the internet which runs much deeper than speech controls and content takedowns.

Understanding the geopolitics of cloud computing demands that we approach select companies and treat them like states, understanding their influence exists in the domain of technology, society, economics, and politics, if not the most visible forms of warfare. The influence of cloud is more than theoretical and has significant implications for policy making across trade, foreign policy, national security, as well as technology policy. The myths we tell ourselves about these interactions risk distorting our perception of the events in front of us.

THERE IS NO CLOUD, JUST OTHER PEOPLE'S COMPUTERS

Cloud providers, in effect, rent out computers and networks to users around the world, from Fortune 500 companies to individuals. Providers build new services on top of their computing resources, like accessible machine translation, sophisticated databases, and new software development tools. Large internet companies increasingly use these cloud services in lieu of building their own technology infrastructure. The growth of cloud computing from an academic research project to a commercial product generating billions of dollars in sales has commoditized computing capacity, storage, and networking bandwidth, and led to a new generation of data-intensive startups.

Cloud computing ties corporate decision-making driven by business risk even more closely to national security risk as a single provider's supply chain decisions and internal security policies can impact millions of customers. This dynamic recalls the "era of big iron," when room-sized mainframes built by a handful of powerful firms were how most users accessed a computer. The language of that era persists today: the vast networks of servers that cloud providers build and operate are similarly cloistered in specialized and well-protected rooms, concentrated under a handful of corporate giants.[1] The decisions these giants take about what technology to buy, build, and operate shapes the technical environment for an increasing number of government and sensitive corporate entities.

These changes in technology have had political ramifications as the growing clout of major cloud service providers causes friction between regulatory models developed for personal computers and servers located in one jurisdiction and a cloud infrastructure that is globally distributed. As ever larger numbers of customers, including intelligence and security agencies, move their data and operations into cloud services, concerns arise over where the infrastructure underneath these services is built and how it is administered. Regulation of the different types of data in the cloud create flashpoints and misunderstanding between companies and governments. Add to that a healthy skepticism from non-Western states about the dominance of US cloud providers, and the conditions are ripe for friction.

Figure 1: Global Public Cloud Revenue

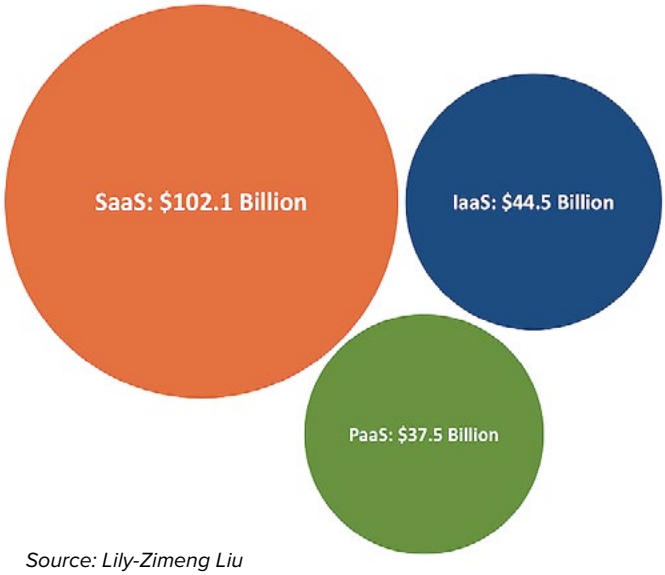

Source: Lily-Zimeng Liu

An unfortunate amount of material written about cloud computing discusses it with the awe reserved for magical woodland creatures and general artificial intelligence. In contrast, cloud computing is quite real—manifest as miles of metal racks housing sophisticated electronics connected by planetary-scale fiber and radio data networks all supported by specialist teams and massive cooling sources. The term cloud comes from networking diagrams where a system being described had a link to some far away set of computers, a line drawn up to the corner of the page toward a bubbly figure representing the "other." This bubbly cloud-like image became shorthand for computers and network services that were not in the scope of the diagram itself but remained accessible. Caring for these 'fleets' of machines demands constant attention and adjustment even the

1 Market analysts jokingly referred to the manufacture of massive mainframes as building "big iron." Working directly with cloud providers' servers, without provider-built software or provider-operated systems, is today referred to as getting close to "bare metal." The era of "big iron" carries into today as references to "metal" abound.

Figure 2: IaaS & SaaS, 2019 Public Cloud Market Share

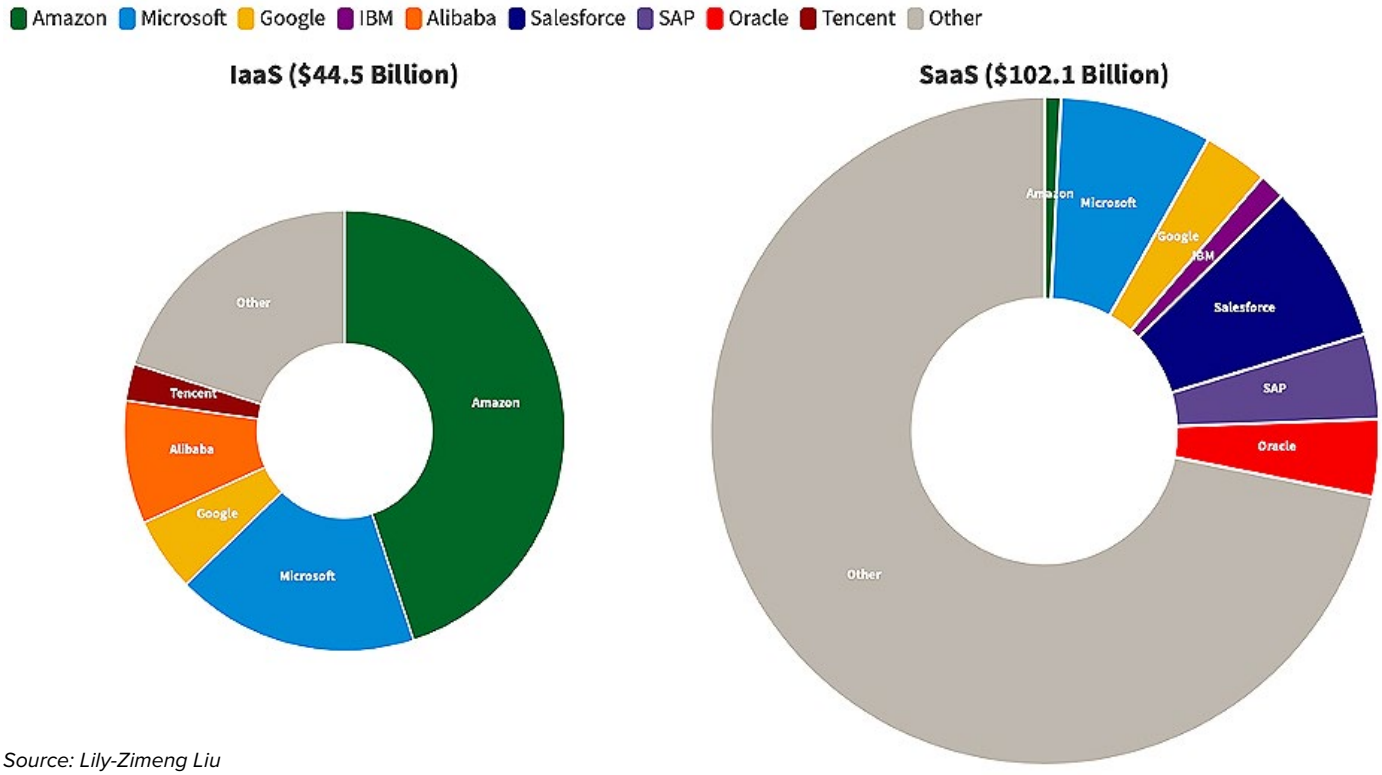

Source: Lily-Zimeng Liu

best run processes can suffer embarrassing failures, like a broken Google update that caused a temporary outage[2] through large swaths[3] of North America in June 2019, or a lightning strike[4] at a Microsoft data center that hobbled[5] Active Directory company-wide for hours.

At the root of the majority of cloud computing is the shared services model, where many users reside on a single physical machine.[6] Multitenancy is the term used to describe shared use, while the technology that makes it possible is called a hypervisor: software that supervises a computer and divides up its resources—processor time, memory, storage, networking bandwidth, etc.—like cake at a birthday party where every partygoer is blindfolded. Everyone gets to enjoy their slice of cake, unaware of those around them enjoying their own portions, too. The hypervisor keeps each user separate, giving them a turn to use the computer while creating the appearance that each is alone on a single machine. The hypervisor is critical to keeping users isolated from one another. Flaws in the hypervisor software can enable attackers to escape from their slice of the computer into that of other users or, worse, into the host machine's operating system controlled by the cloud provider.

In a cloud service, each of these computers runs additional software selected by the cloud provider and user and each is tied into a network. By building services which use these networked machines, cloud providers can take storage at a facility in Frankfurt, match it with processing in Texas, and deliver the result to a user in Tokyo. Web mail, search results, streaming video, photo storage and sharing, the computation that makes a digital assistant responsive to

2 Liam Tung, "Google Details 'Catastrophic' Cloud Outage Events: Promises to Do Better Next Time," *ZDNet*, June 7, 2019, https://www.zdnet.com/article/google-details-catastrophic-cloud-outage-events-promises-to-do-better-next-time/.

3 "Google Cloud Networking Incident #19009," Google Cloud Status Dashboard, accessed January 15, 2020, https://status.cloud.google.com/incident/cloud-networking/19009.

4 Kurt Mackie, "Microsoft's Cloud Recovering from Datacenter Lightning Strike," *Microsoft Certified Professional Magazine Online*, September 5, 2018, https://mcpmag.com/articles/2018/09/05/microsoft-cloud-datacenter-lightning-strike.aspx.

5 Richard Speed, "Microsoft Reveals Train of Mistakes That Killed Azure in the South Central US 'Incident,'" *Register*, September 17, 2018, https://www.theregister.co.uk/2018/09/17/azure_outage_report/.

6 Some cloud service providers sell direct access to servers without any additional services or software from the provider stacked on top. These "bare-metal" deployments may involve multitenancy within a single-user organization or avoid it entirely, offering exclusive use of the hardware.

Figure 3: Components of a Data Center

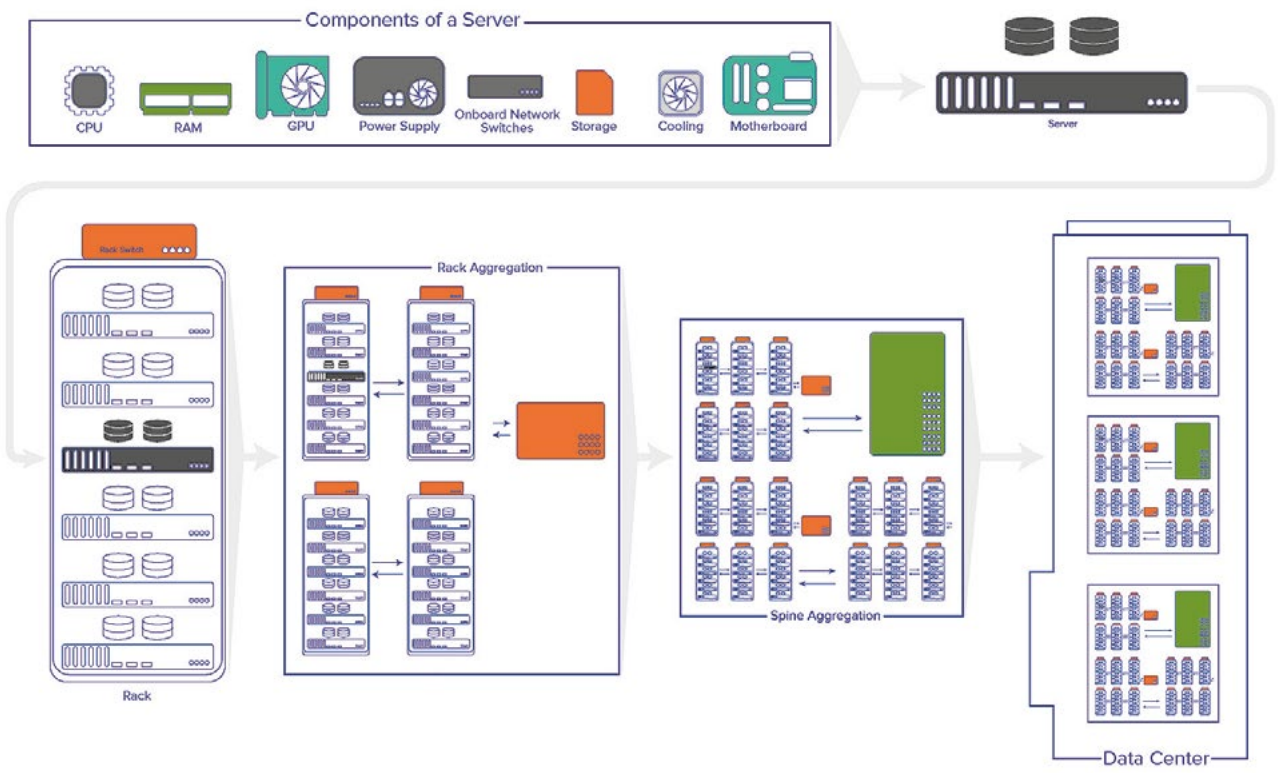

Source: Tianjiu Zuo and John Eric Goines

Atlantic Council

your voice—all of these are services layered on top of the cloud. In industry parlance, there are three basic models of cloud service:

1. **Infrastructure as a Service (IaaS):** These are the raw computing, storage, and networking elements that users can rent and consume like a service rather than a product but must largely set up and configure themselves. For example, renting a virtual machine to host an email server.

2. **Platform as a Service (PaaS):** This is the diversity of software and online services built on top of the cloud. Users access these services without managing the underlying infrastructure. For example, the machine learning service an aircraft engine manufacturing company integrates into its products to predict when they will fail.

3. **Software as a Service (SaaS):** These are the online services that require no deep administration from the user. These services are offered without substantial ability to rewrite, rebuild, or reintegrate them like PaaS. For example, sharing documents online or the image recognition service a hospital uses to identify tumors in a CT scan.

There are hundreds of cloud companies, most selling services in one model. A handful compete in all three and the largest of these are referred to as the hyperscale providers—Microsoft, Google, Amazon, and Alibaba. Cloud computing is an expanding constellation of technologies—some old, some repurposed, and some wholly new. Much of the innovation in cloud is in managing these fleets of machines and building the vast networks required to make them accessible for users, rather a single snazzy new product or feature. There is no one single model of cloud computing. The major providers all build their infrastructure in slightly different ways, influenced by market strategy and legacy technology investments, but the abovementioned three models help categorize what one might find in the cloud.

Security in the cloud is similarly a mix of the old and new. Old are the challenges of ensuring that users, and not

malicious actors, have access to their data, and systems are protected against myriad integrity and confidentiality threats that range from distributed denial-of-service (DDoS) attacks to power outages. New is the need to do all of this across tens of thousands of domains and millions of users every day; cloud presents an enormous challenge of the scale on which these longstanding security functions need to take place. This has driven increased automation and more user-friendly tools but also created recurring gaps[7] where the user and cloud provider are not in sync[8] about each other's security responsibilities.

Similarly, new is cloud providers playing host to a growing domain of conflict. There are instances where the origin and destination of an attack occurred in infrastructure owned by the same cloud provider; attacker and defender using the same cloud and observed (possibly interdicted) by the cloud provider. As ownership of information technology (IT) infrastructure concentrates, so does exposure to what two US academics labeled the "persistent engagement" of cyberspace—with fewer and fewer major providers, there is a higher likelihood of engagements that start and end within the same network.[9] Cloud providers have to balance the responsibilities of their global customer base with the demands of their home governments. These providers are put in the position of arbitrating between their terms of service and security commitments to customers and national intelligence and military activities taking place in their infrastructure, creating a tangled web of business and national security risk.

Figure 4: Illustrating the Multi-Tenant Model

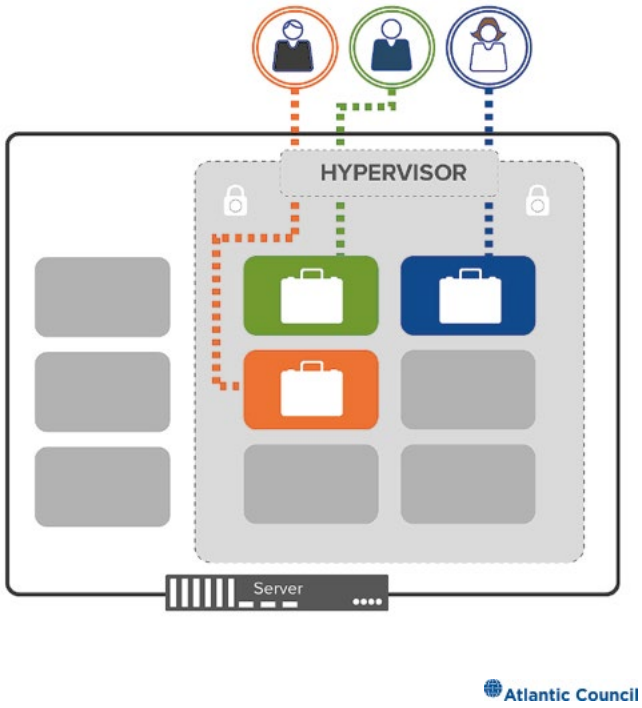

Source: Simon Handler, Trey Herr, and John Eric Goines

Deploying Hybrid Cloud

Driven by specific business, regulatory demands, or organizational discomfort with the cloud model, some users combine pubic cloud services with locally managed infrastructure. This half cloud, half local hybrid model varies by provider but generally finds cloud services deployed alongside a user's existing infrastructure to run the same software as in a public cloud data center. These two work alongside each other and give users access to some or all cloud services while still allowing them to run their own equipment. Hybrid cloud can preserve technical flexibility by allowing organizations to retain other equipment or allow a slower transition from self-run data centers to full use of a public cloud. It also involves some compromise on the economic model behind cloud, premised on widest possible use of shared infrastructure, and requires additional administrative capacity and competence from users. Hybrid cloud serves a real business interest, helping transition less cloud-friendly users, but has a real political benefit as well—making data localization easier. Hybrid cloud allows for sensitive services or data to be physically located in a specific jurisdiction, while remaining linked to public cloud services. This hybridity comes with additional cost, and large-scale deployments can be more technically complex while posing additional security risks as more ongoing technical and policy decisions are shared between user and provider.

7 Ericka Chickowski. "Leaky Buckets: 10 Worst Amazon S3 Breaches." *Business Insights Blog*, January 24, 2018, https://businessinsights.bitdefender.com/worst-amazon-breaches.

8 Kaushik Sen. "S3 Security Is Flawed By Design." UpGuard, last updated December 5, 2019, https://www.upguard.com/blog/s3-security-is-flawed-by-design.

9 Richard J. Harknett and Michael P. Fischerkeller, "Deterrence Is Not a Credible Strategy for Cyberspace," Foreign Policy Research Institute, June 23, 2017, https://www.fpri.org/article/2017/06/deterrence-not-credible-strategy-cyberspace/.

EXAMINING FOUR MYTHS

The challenge with telling any story about cloud computing is avoiding spectacle and hyperbole. The cloud will revolutionize, energize, and transform! Much of what is in the public domain about cloud computing comes from marketing material and is thus a creative interpretation of the truth.

Myths are useful in the telling of a story. They simplify complexity and add drama to otherwise mundane events. But myths can become a lie when repeated over and over. In cybersecurity, myths come from all parties: naïve regulators, fearful companies, excitable marketing, and collective misunderstanding. The myth at the heart of the cloud is that only technical factors determine how these services are delivered and how their infrastructure is built.

On the contrary, political realities exercise an ever greater influence over how the cloud is built, deployed, and used in global businesses and security enterprises. External forces, like pressure to store data within a particular jurisdiction or concerns over a supply chain, are equal in weight to internal forces as the concentration of ownership in cloud computing raises significant questions about the democratic accountability of technology. To shed light on the geopolitics of cloud computing, this issue brief examines the following four myths of the cloud:

- All data is created equal
- Cloud computing is not a supply chain risk
- Only authoritarian states distort the public cloud
- Cloud providers do not influence the shape of the internet

These are not the only myths about cloud computing, but they are some of the most persistent, persuasive, and unrealistic.

MYTH 1: ALL DATA IS CREATED EQUAL

It is challenging to talk about cloud computing without discussing data. For all its metal and concrete infrastructure, snazzy code, and marketing materials, cloud computing often comes down to managing huge volumes of data. Three categories are helpful:

- **User data**, what the customers of a service store in the cloud: emails, tax files, design documents, and more.
- **Derived data**, which allows cloud providers to learn about how users access and interact with these files: which documents do users from an office in Berlin tend to access first and should they be stored nearby to reduce latency?

Figure 5: Illustrating Data Types in the Cloud

Source: Lily-Zimeng Liu and John Eric Goines

- **System data**, or what cloud providers learn about their systems from the way users consume services: what causes a spike in processor utilization or a drop in available bandwidth to a data center or a security alert for malware?[10]

User data was the subject of intense attention from the European Union's General Data Protection Regulation (GDPR) resulting in changes by cloud providers. These changes have manifested in new tools and policies for users to limit access to their data, determine where in the world it can be stored, and transport it between cloud providers should the need arise. User account information—details like name, address, payment method, and what services the user consumes—is often given separate treatment from user data as it provides valuable information about user preferences and can be important to trace malicious behavior.

Derived data is incredibly rich and varied; what a Facebook user clicks on depending on the time of day, their most frequently liked posts, or how many times they start but delete a comment. Derived data about user behavior is the secret sauce behind most "behavioral-analytics" security tools and it powers the advertising machine at the heart of Google/Alphabet's revenue's engine: YouTube viewing habits, faces appearing together in photos, even the text of emails on the widely used Gmail. All of this data tells a story about users and their habits, beliefs, dreams, and desires—a story which is a commodity.[11]

For cloud providers, this derived data can help understand which services succeed or fail, how to more precisely price offerings, and where or when to introduce new features. Investigations and security responses produce even more data, combining the activities of multiple customer accounts to track an attacker as it moves across the provider's infrastructure. The result is extensive privacy training and data security measures within at least the three largest US cloud providers—Amazon, Google, and Microsoft. Security-focused staff at these providers are subject to additional controls on data retention, restrictions on sharing data, and regular privacy training. Some of these controls impose difficult limits on how long data can be retained. If a recently detected cyberattack on a cloud service is the product of months of careful reconnaissance and infiltration by the attacker, limits on retaining data no longer than 90 days can make investigations to determine the source of the breach highly challenging.

System data is everything a cloud provider can learn about its own systems from how they are used; how a server heats up in response to sustained requests for new video links or network traffic associated with a particularly massive shared spreadsheet. System data includes technical information like the network routes a particular server uses to move data

Derived Data for Security: The Password Spray

Cloud providers use derived data for numerous security functions. All of the intelligence a company like Amazon gathers about its services—which customer accounts access the same data, how many times a user has logged in or tried the wrong password—are signals in the nervous system of the cloud.[1] *These signals are combined to understand "normal" behavior and identify malicious activity. Repeated attempts to log into the same account with many different passwords is an old way to try to break in. Many attackers now employ "password sprays"— attempting to access a large number of accounts with just a few commonly used passwords.*[2] *A cloud provider could track many or all of those attempts even if they take place across different users, companies, and accounts. Since the cloud provider builds and operates the infrastructure, it knows much of what goes on within, allowing it to combine data from hundreds of users to identify a single attacker or malicious campaign.*

1 Shlomo Sagir et. al, "Get instantaneous behavioral analytics and anomaly detection" Microsoft, June 28th, 2020, https://docs.microsoft.com/en-us/cloud-app-security/anomaly-detection-policy.

2 Andy P, "Spray You, Spray Me: Defending Against Password Spraying Attacks," National Cyber Security Centre, May 15, 2018, https://www.ncsc.gov.uk/blog-post/spray-you-spray-me-defending-against-password-spraying-attacks.

10 Cloud providers, and many online services, typically segregate account data. One of the major challenges for security teams reviewing an incident is to determine who is responsible for the activity they are tracking between or within a single account. Even matching logs from a single machine to customer accounts across products can prove tricky; the task is that much harder when accounts have been hijacked or purchased using stolen credentials or credit cards. There are a multitude of other taxonomies to categorize data, including several global standards, for example, ISO 8000, ISO/IEC 19944, and ISO/IEC 38505-1 and -2. What is presented here is a slightly abstracted version of the data types found in ISO/IEC 19944.

11 Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: PublicAffairs, 2020).

Figure 6: Data Protection Laws in Europe and Turkey

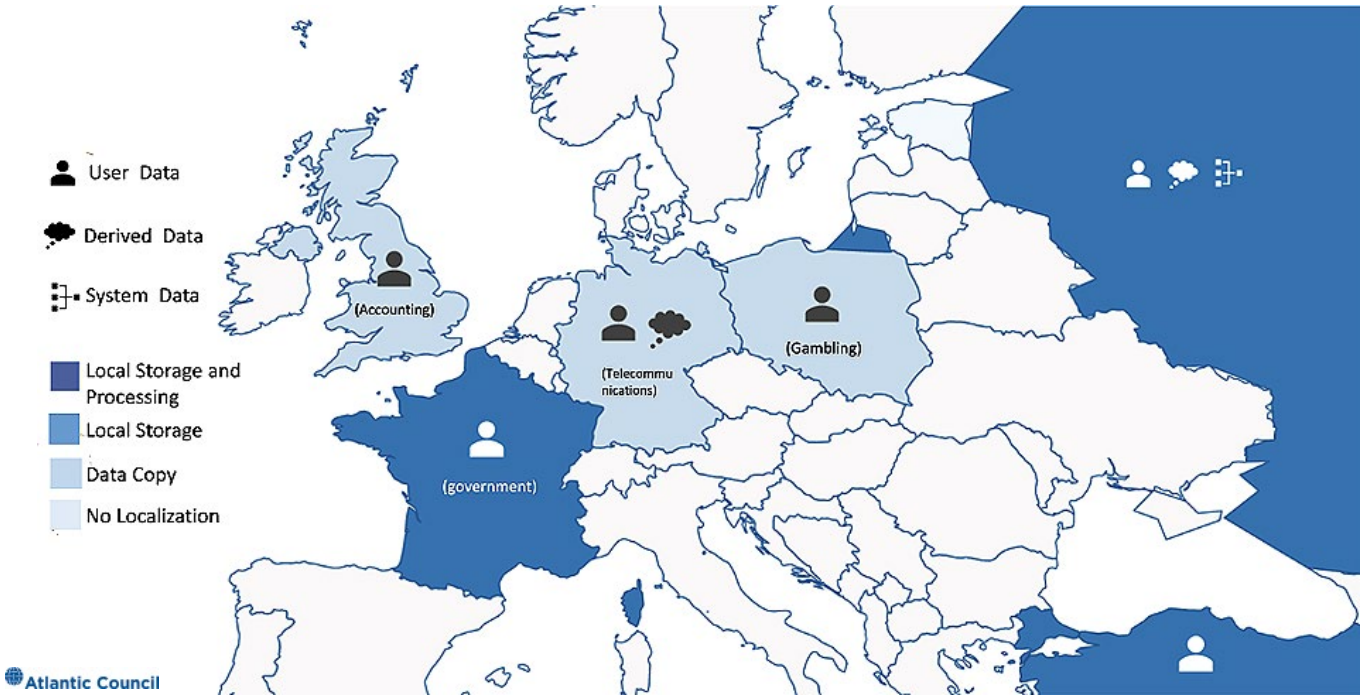

Source: Lily-Zimeng Liu (DLA Piper, "Data Protection Laws of the World," https://www.dlapiperdataprotection.com/.)

between facilities, logs of all the network addresses a single machine tries to access, results of testing for the presence of unverified or disallowed code, or the temperature in front of a rack in a data center. This kind of data is rarely tied to a specific user and is likely to be accessed by more specialized teams like Microsoft's Threat Intelligence Center[12] or Google's Threat Analysis Group.[13] Much of the derived and system data created are quickly destroyed, sometimes after only a few days or weeks. At the mind-boggling scale of global cloud providers, even the cost of storage becomes a limiting factor in how long data lives.

Differentiating between these types of data and their common uses becomes significant for debates about where this data should be located. One of the most frequent debates about cloud computing, especially when providers first enter a market, is where providers will store and process data. Efforts to force data to be hosted in a single jurisdiction are collectively referred to as data localization.[14] Each of these data types serves a distinct business purpose within cloud providers that is rarely captured with nuance by localization requirements. Many localization requirements focus on user data, yet many of these same requirements state their intent is to restrict the widespread misuse of data in online advertising which typically leverages derived data. The most difficult category of the three types of data for providers to specifically locate is system data and yet this is rarely carved out from localization requirements.

Two laws illustrate the challenge of ignoring different types of data and the workings of the public cloud. Vietnam passed a new and expansive data localization requirement in 2018 for all manner of internet service firms, including cloud providers, to store Vietnamese users' data in Vietnam for a period of time defined by the government. The law also requires any processing, analysis, or combination of this data must take place on Vietnamese soil.[15] This produced a flurry of protests from cloud providers without giving Vietnamese information about how to address their chief underlying concerns about user data. In contrast, in October 2019, Indonesia modified a longstanding requirement that mandated internet service companies locate their data centers in the country. The

12 Patrick Howell O'Neill, "Inside the Microsoft Team Tracking the World's Most Dangerous Hackers," *MIT Technology Review,* November 6, 2019. https://www.technologyreview.com/2019/11/06/238375/inside-the-microsoft-team-tracking-the-worlds-most-dangerous-hackers/.

13 Robert McMillan, "Inside Google's Team Fighting to Keep Your Data Safe From Hackers," *Wall Street Journal*, January 23, 2019, https://www.wsj.com/articles/inside-googles-team-battling-hackers-11548264655.

14 Jennifer Daskal and Justin Sherman, *Data Nationalism on the Rise: The Global Push for State Control of Data*, Data Catalyst, June 23, 2020, https://datacatalyst.org/reports/border-control-the-rise-of-data-nationalism/.

15 Ngo Duy Minh, "Vietnam Data Protection Overview," OneTrust DataGuidance, July 30, 2019, https://www.dataguidance.com/notes/vietnam-data-protection-overview.

Figure 7: Data Protection Laws in Asia and Australia

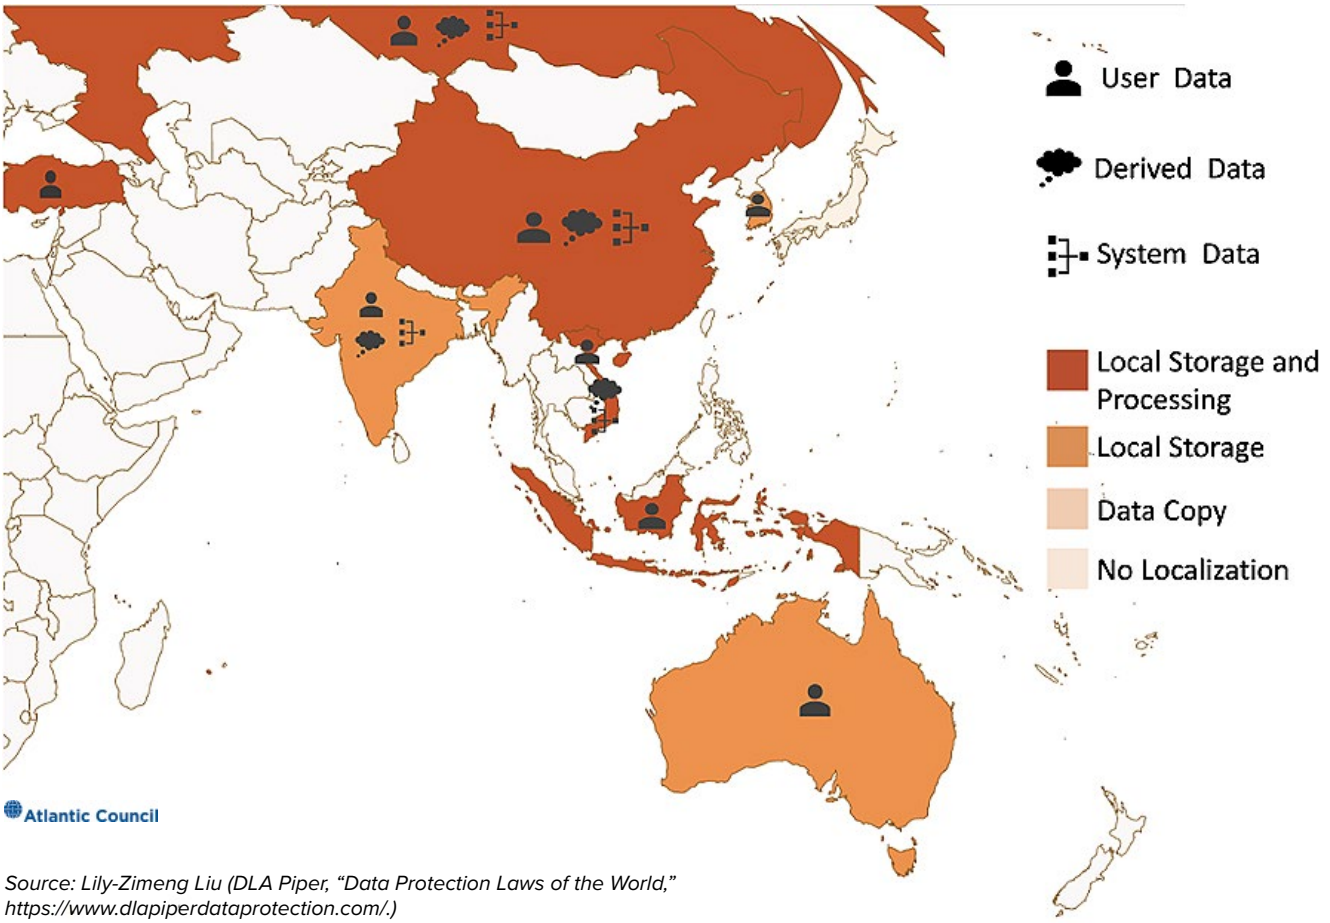

Source: Lily-Zimeng Liu (DLA Piper, "Data Protection Laws of the World," https://www.dlapiperdataprotection.com/.)

rule is now limited to entities that provide public services and exempts those that use storage technology that is not available in Indonesia.[16]

The Indonesian law focused on data centers—linking the location of infrastructure with the location of data.[17] Given how the public cloud model can locate processing and storage of data in widely disparate locations, data could be copied to this local infrastructure while processed and stored elsewhere. Even if user data was covered by the law, it is not clear that derived or system data would be, the latter in particular because it is largely generated by the provider about itself. The Vietnamese law, in contrast, appeared to cover all three types of data with its blanket restrictions on processing and data combination. This implied where infrastructure might be located and further limited what cloud providers could do with what Vietnamese users stored on their service and what the provider learned about these users.

Some governments look to localization to help improve the security of their users' data despite little evidence to support the premise, and even evidence to challenge it. Forcing cloud providers to build variants of otherwise standard infrastructure designs and operating models could introduce unknown flaws or gaps in security practices. The public cloud is premised on global networks of computing and storage resources, not isolated local clusters. While data centers and associated power/cooling infrastructure do have to be located in a specific physical location, the ability to pass derived data about malicious activities and track abnormal events across the world is where cloud can improve users' security. The same is true of user and system data—linking data centers in different parts of the world can provide more redundant backups and force providers to maintain consistent security and operational practices across these facilities. Figure 8 displays the location of every publicly acknowledged data center from the five listed cloud providers as of late 2019.

16 Hiswara Bunjamin and Tandjung, "Indonesia's Electronic Systems and Transactions Regulation Replaced and Data Regulation Amended," October 25, 2019, https://www.hbtlaw.com/latest-thinking/indonesias-electronic-systems-and-transactions-regulation-replaced-and-data.

17 Kay Vasey, "Indonesia Moves Towards Comprehensive Data Law—How Will It Impact Your Business?" March 30, 2017, http://www.connectedasia.com/indonesia-moves-towards-comprehensive-data-law-how-will-it-impact-your-business/.

Figure 8: Global Coud Data Center Map

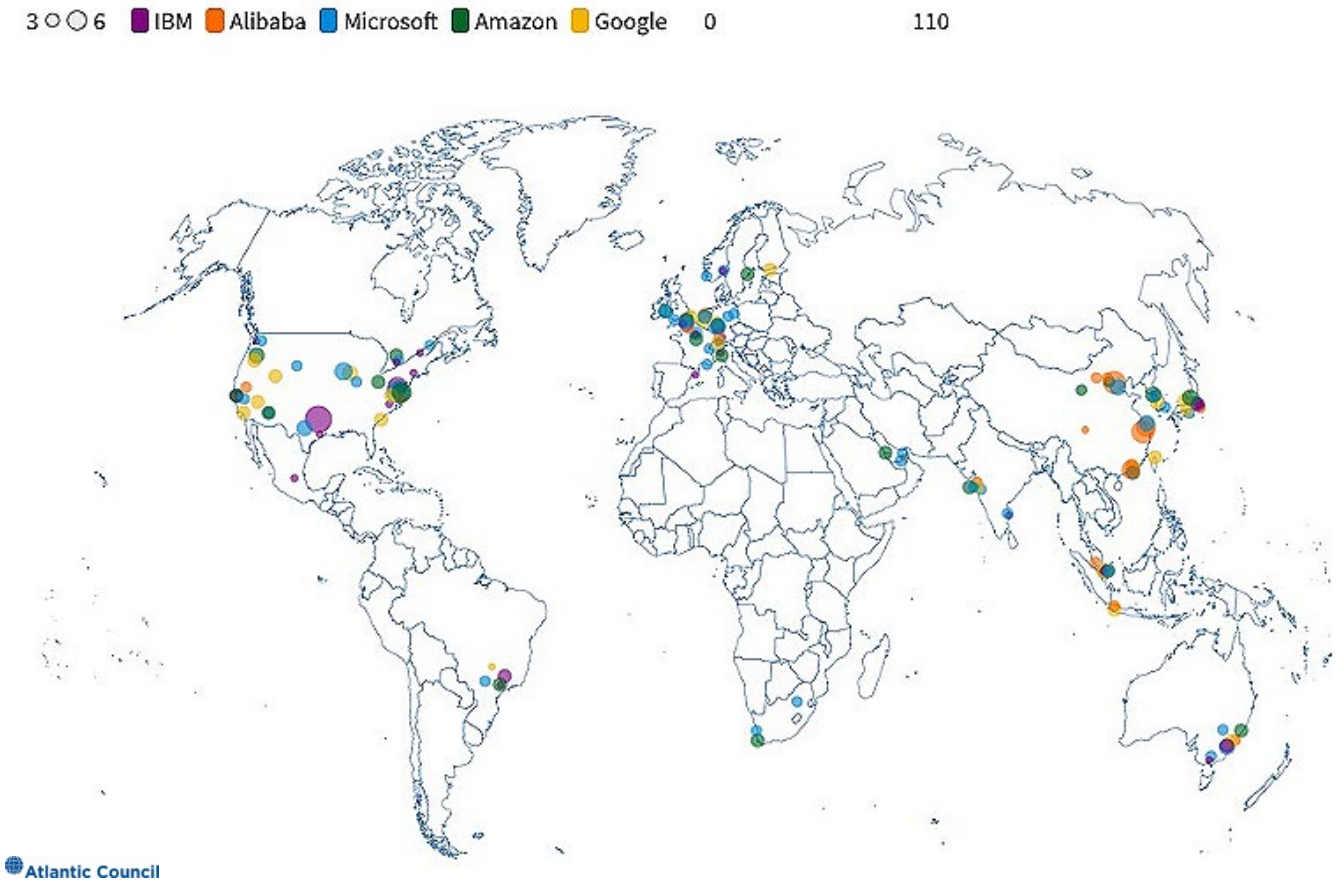

Source: Lily-Zimeng Liu[1, 2]

1 Data Centers, Alibaba Cloud's Global Infrastructure, accessed August 10, 2020, https://www.alibabacloud.com/global-locations; AWS Global Infrastructure Map, AWS, accessed August 10, 2020, https://aws.amazon.com/about-aws/global-infrastructure/; Google Cloud, accessed August 10, 2020, https://cloud.google.com/about/locations/#regions; Global Locations for your Global Business, cloud data center locations, IBM, accessed August 10, 2020, https://www.ibm.com/cloud/data-centers/; Azure Geographies, Microsoft Azure, accessed August 10, 2020, https://azure.microsoft.com/en-us/global-infrastructure/regions/.

2 This map and the graphics derived from it are based on publicly available data (early August 2020) from the cloud providers themselves as noted above. Vendor definitions of region and availability zone vary and not all datacenters are built and maintained directly by each corporation, some are leased in whole or in part.

Many of these are clustered around major cities and internet exchange or cable landing points.

Cloud providers have evolved[18] their policies[19] and tools[20] in light of demands for greater control and new regulatory requirements. This is progress and it reflects no small amount of investment on the part of these companies, especially the hyperscalers. But it also reflects a baseline expectation—users own their data and control should follow with ownership. It can be challenging to keep that linkage of ownership and control tight across globally distributed networks and constantly evolving infrastructure, but it is table stakes as cloud becomes the default model for much of computing.

18 Julie Brill, "Increasing Transparency and Customer Control over Data," *Microsoft On the Issues*, April 30, 2019, https://blogs.microsoft.com/on-the-issues/2019/04/30/increasing-transparency-and-customer-control-over-data/.

19 Sunil Potti, "Advancing Control and Visibility in the Cloud," Google Cloud, November 20, 2019, https://cloud.google.com/blog/products/identity-security/new-security-tools-for-google-cloud-and-g-suite.

20 AWS, "AWS Control Tower Is Now Generally Available," June 24, 2019, https://aws.amazon.com/about-aws/whats-new/2019/06/aws-control-tower-is-now-generally-available/.

Figure 9: Relative Global Data Center Count by Provider/Region

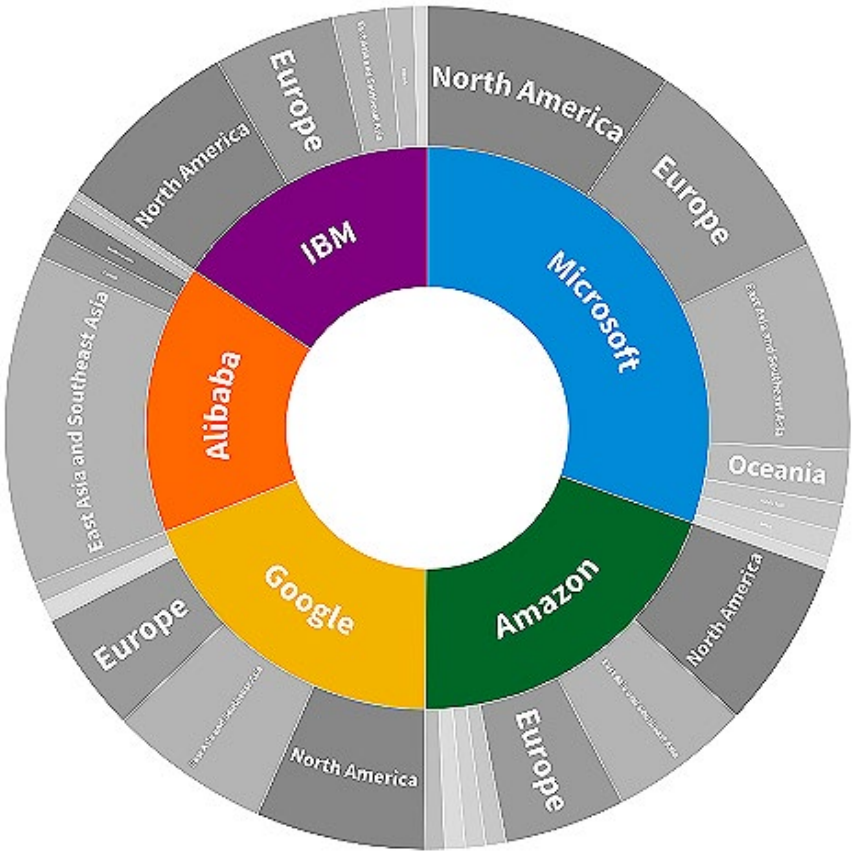

Source: Lily-Zimeng Liu

Localization is the visible tip of a much larger debate about how to govern data, including its use, storage, retention, combination, and lawful access by governments; debate whose values should and will govern data and how those values will be enforced internationally. As the cloud grows more common and sophisticated in widespread use, countries must rapidly evolve policies to avoid a gap between technical design and commercial practice on one side and security and normative priorities of society on the other. Initially, this gap is a business cost to cloud providers that may pale in comparison to data privacy or national security concerns. Over the long run, such a gap has the practical effect of distorting the economic and technical underpinnings of the public cloud. Isolating markets through localization requirements makes new cloud infrastructure less cost-effective. This, in turn, reduces the incentive for manic competition that helps drive companies like Google to commit billions of dollars to update products, offer interesting services, and build in better security as a way to compete with larger rivals like Amazon and Microsoft.

WHO SHOULD CARE?

Regulators around the world must address how they govern the different types of data used by cloud providers. The alternative is costly breakdowns in the public-private partnership over cloud security and deployment where companies opt out of smaller markets and offer services which are more costly and less flexible in those larger. Innovation requires investment and attention to the rules that govern intellectual property ownership, access to data and requisite research facilities, and the ease of bringing new ideas to market, among many other factors. Many of the US firms under fire in Europe have made their business monetizing data derived from their users because it was possible, profitable, and popular. These conditions, however, may be changing. The European Union (EU) would do well to create a robust framework for data governance, including limiting requirements for localization in the EU, that facilitates the next generation of internet entrepreneurs rather than assailing the exemplars of this one (looking at you, France[21]).

21 KPMG, "France: Digital Services Tax (3%) Is Enacted," July 25, 2019, https://home.kpmg/us/en/home/insights/2019/07/tnf-france-digital-services-tax-enacted.html; Winston Maxwell and Mathilde Gérot, "Ignoring GDPR, French Senate Votes for a Data Localization Amendment," *Chronicle of Data Protection*, June 1, 2016, https://www.hldataprotection.com/2016/06/articles/international-eu-privacy/ignoring-gdpr-french-senate-votes-for-a-data-localization-amendment/.

RECOMMENDATIONS

- **[White House]** Do not forget about machine learning. Finally appoint a chief technologist to the Federal Trade Commission. Charge this person with leading a multistakeholder group to define a policy framework for the status of machine learning models, their outputs, and associated data, including legal ownership and classification under major data governance regimes. Work with Congress to implement key portions of this framework into law and help lead the global policy debate.

- **[The European Union Agency for Cyber Security]** ENISA should develop an updated Network and Information Security (NIS) Directive and rules following on the Cybersecurity Act to ease the adoption of the cloud in regulated industries and security-specific national agencies. This should build on the work of the EU's Cloud Service Provider Certification (CSPCERT) Working Group and identify rules which could be rolled back or revised.

MYTH 2: CLOUD COMPUTING IS NOT A SUPPLY CHAIN RISK

Supply chain policy has been dominated by a focus on telecommunications, especially 5G, over the past several years. Efforts by the United States to block certain technology providers, and pushing others to do the same, emphasize the risk posed by allowing untrusted hardware and software into "core" networks, the broad public utility of 5G, and the significance of trusted telecommunications services for both public and private sectors well into the future.

Cloud computing poses a similarly broad and far more immediate source of supply chain risk. Indeed, many cloud service providers are moving into the 5G market, partnering with traditional telecommunications firms, while others take advantage of the increasing virtualization of telephony hardware into software; Amazon is partnering with Verizon to provide 5G services[22] and Microsoft acquired virtualized telecommunications provider Affirmed Networks earlier in 2020.[23] Indeed, several cloud providers are effectively

Table 1: Global Data Center Count

	Total	Alibaba	Amazon	Google	IBM	Microsoft
Total	387	60	76	73	60	118
East and South East Asia	133	49	21	24	8	31
North America	119	4	25	25	28	37
Europe	90	4	18	18	18	32
Oceania	20	2	3	3	4	8
South America	10	0	3	3	2	2
Middle East	8	1	3	0	0	4
Africa	7	0	3	0	0	4

Source: see above, map

22 Jon Fortt and Annie Palmer, "Amazon Just Partnered with Verizon to Improve 5G Speeds," CNBC, December 5, 2019, https://www.cnbc.com/2019/12/03/amazon-verizon-partner-on-new-5g-wave-length-product.html.

23 Frederic Lardinois, "Microsoft Acquires 5G Specialist Affirmed Networks," TechCrunch, March 26, 2020, https://techcrunch.com/2020/03/26/microsoft-acquires-5g-specialist-affirmed-networks/.

Figure 10: Data Center Server Supply Chain

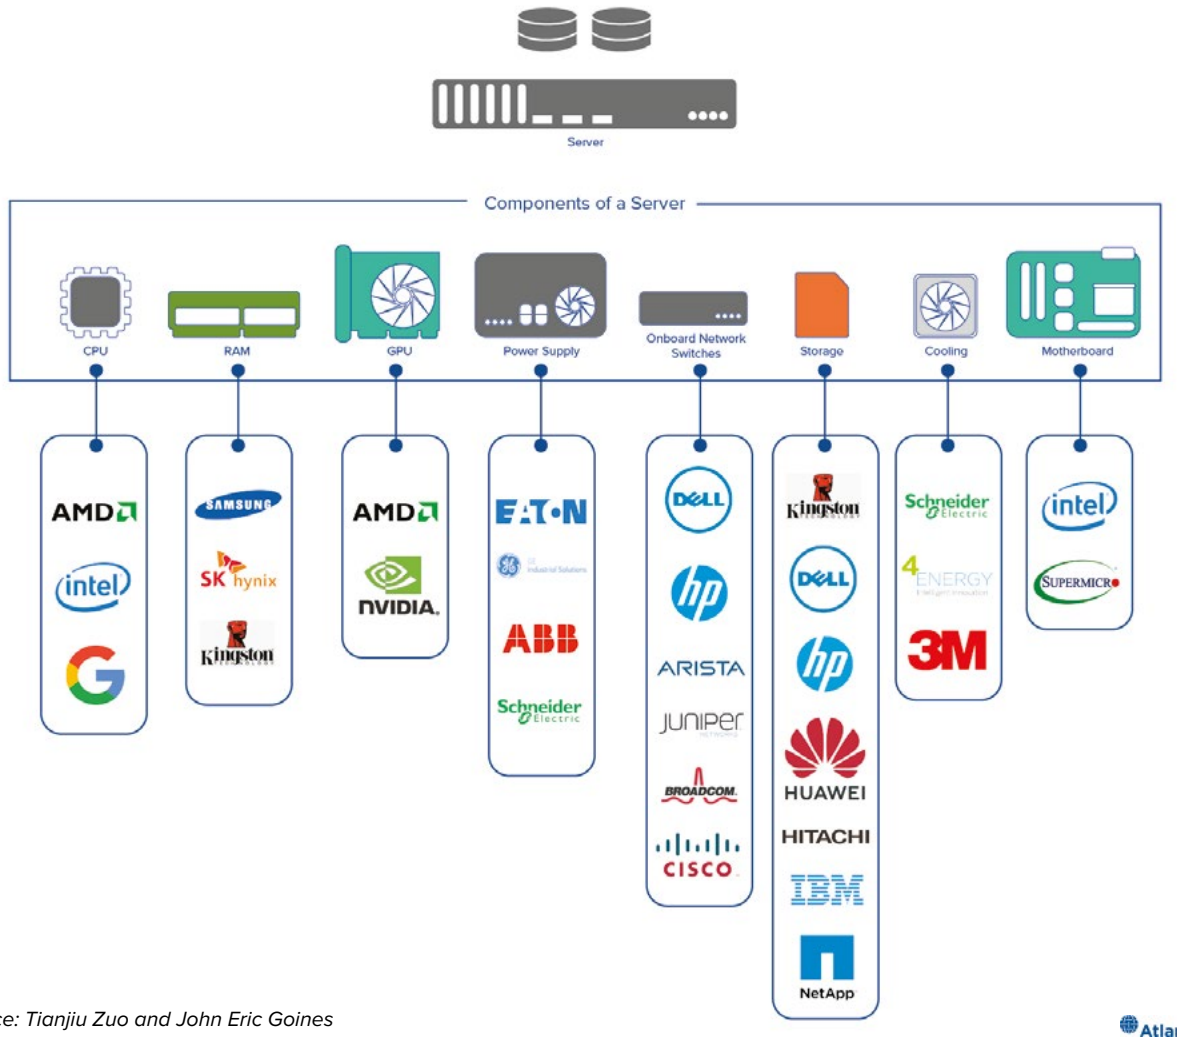

Source: Tianjiu Zuo and John Eric Goines

Atlantic Council

telecommunications companies themselves with large global networks of undersea and overland fiber optic cables.[24]

Cloud providers buy and maintain vast computing infrastructure and host customers around the world. They are targeted by the same diversity of sophisticated threats that target telecommunications firms. Intelligence and defense agencies across the world, along with the financial sector and nearly every Fortune 500 company, use cloud computing. Cloud computing is already a pervasive infrastructure and an important source of supply chain risk.

Major cloud providers aggregate the risk from commodity computing and networking technologies by purchasing processors, server boards, networking switches, routers, and more in galactic quantities. Each of these chips, cables, software packages, and servers comes with potential vulnerabilities. Cloud services rely on massive quantities of software, including code developed by third parties and open-source projects. As previous work has shown,[25] these software supply chains are vulnerable.

Cloud providers must manage many different kinds of risk—they operate large, sometimes global, telecommunications networks; manage huge corporate enterprise IT networks; develop and maintain massive quantities of software each year; and build out large data centers stuffed with computing, network, and storage gear along with the physical plants to

24 Rich Miller, "Cloud Players Are Redrawing the Subsea Cable Map," Data Center Frontier, December 4, 2018, https://datacenterfrontier.com/cloud-players-are-redrawing-the-subsea-cable-map/.

25 Trey Herr, June Lee, William Loomis, and Stewart Scott, *Breaking Trust: Shades of Crisis Across an Insecure Software Supply Chain, Atlantic Council*, July 26, 2020, https://www.atlanticcouncil.org/breaking-trust/.

Hacking the Hypervisor

Cloud computing supply chain risk is not limited to hardware or closely related software (firmware). Hypervisors, the software that allows a single computer to be used by multiple entities at once while keeping them separated, are critical to the functioning and security of cloud services. A vulnerability in a hypervisor could allow attackers to escape their virtual machine and gain access to other users or even the host sever like the Venom flaw discovered in KVM, Xen, and the hypervisor QEMU in 2015.[1] Google uses a variant of the open-source KVM as increasingly does Amazon, shifting away from the also open-source Xen. As part of its switch to KVM, Amazon has deployed what it calls Nitro—a set of hardware-based hypervisor and security functions. Each function is designed into a single chip, wringing out maximum efficiency and performance.[2] Microsoft runs an internally developed product called Hyper-V, which is built into Windows Server. Hypervisor escapes pose a great challenge to cloud services as this software provides the foundation for isolation between users and the provider. One vulnerability broker has offered up to $500,000 for vulnerabilities leading to an escape based on demand from customers.[3] There are notable examples of vulnerabilities disclosed in open[4] research[5] and competitions,[6] which suggest still more vulnerabilities being used or sold without disclosure.

1 Wolfgang Kandek, "Venom Hypervisor Vulnerability," *Qualys Security Blog,* June 3, 2020, https://blog.qualys.com/laws-of-vulnerabilities/2015/05/13/venom-hypervisor-vulnerability; "VENOM: QEMU Vulnerability (CVE-2015-3456)," Red Hat Customer Portal, August 25, 2016, https://access.redhat.com/articles/1444903.

2 Paul McLellan, "HOT CHIPS: The AWS Nitro Project," *Breakfast Bytes Blogs*, Cadence, October 2, 2019, https://community.cadence.com/cadence_blogs_8/b/breakfast-bytes/posts/the-aws-nitro-project.

3 Catalin Cimpanu, "Zerodium Offers Big Bucks for Cloud Zero-Days," *Zero Day, ZDNet*, March 5, 2019, https://www.zdnet.com/article/zerodium-offers-big-bucks-for-cloud-zero-days/.

4 Lindsey O'Donnell, "Black Hat 2019: Microsoft Protocol Flaw Leaves Azure Users Open to Attack," threatpost, August 7, 2019, https://threatpost.com/black-hat-2019-microsoft-protocol-flaw-leaves-azure-users-open-to-attack/147045/.

5 Rafal Wojtczuk, "Poacher Turned Gatekeeper: Lessons Learned from Eight Years of Breaking Hypervisors," presentation to Black Hat USA 2014, https://paper.bobylive.com/Meeting_Papers/BlackHat/Europe-2014/eu-14-Wojtczuk-Lessons-Learned-From-Eight-Years-Of-Breaking-Hypervisors.pdf.

6 Anmol Sachdeva, "Chinese Hackers Win $382,500 For Hacking KVM On Ubuntu, Edge, Adobe Reader," Fossbytes, November 18, 2019, https://fossbytes.com/chrome-safari-edge-hacked-chinese-hacking-competition/.

keep these data centers powered and cooled. The largest hyperscale providers in the West—Amazon, Google, and Microsoft—attempt to manage their risk through extensive security audits of purchased products, setting and assessing security standards for technology vendors, and detailed threat intelligence collection and reporting. Building trust in the secure design and support process of these vendors is crucial for cloud providers. In some cases, the cost or complexity of establishing this trust is too great and cloud providers bring a firm in house or even replace key vendors by building their own component devices.[26], [27] Nonetheless, risk can only be managed, not eliminated.

Cloud supply chain risk shares some important similarities with other domains. The same flaws in chips or vulnerabilities in software that would impact a cloud data center also pose risks to traditional enterprises in managing their own infrastructure. Managing the security of a large corporate network at a Google or an IBM requires many of the same decisions about investing limited resources as it does at a Goldman Sachs or a Walmart. Software supply chain matters to an open-source project just as it does to government agencies or the military.

The difference with the cloud is in its scale. Where a company running its own data centers might buy hundreds or thousands of servers ever year, cloud providers are buying hundreds of thousands of servers and assembling or even designing their own. This equipment is shipped all over the world, as shown in Figure 11, to data centers owned or leased by each provider. Depending on the service model—IaaS, PaaS, or SaaS—cloud providers might make significant design and procurement decisions about every single piece of the technology stack used to provide a cloud service or share more of that responsibility with customers.

This scale also informs a public interest in risk posed by cloud providers as part of the technology supply chain. Problems in the design of a chip or compromised firmware can become

26 Tom Simonite, "New at Amazon: Its Own Chips for Cloud Computing," *Wired*, November 27, 2018, https://www.wired.com/story/new-amazon-chips-cloud-computing/.

27 Ian King, "Cisco Enters Chip Market, Supplying Microsoft, Facebook," Bloomberg, December 11, 2019, https://www.bloomberg.com/news/articles/2019-12-11/cisco-enters-chip-business-begins-supplying-microsoft-facebook.

Figure 11: Global Distribution of Cloud Provider Data Centers

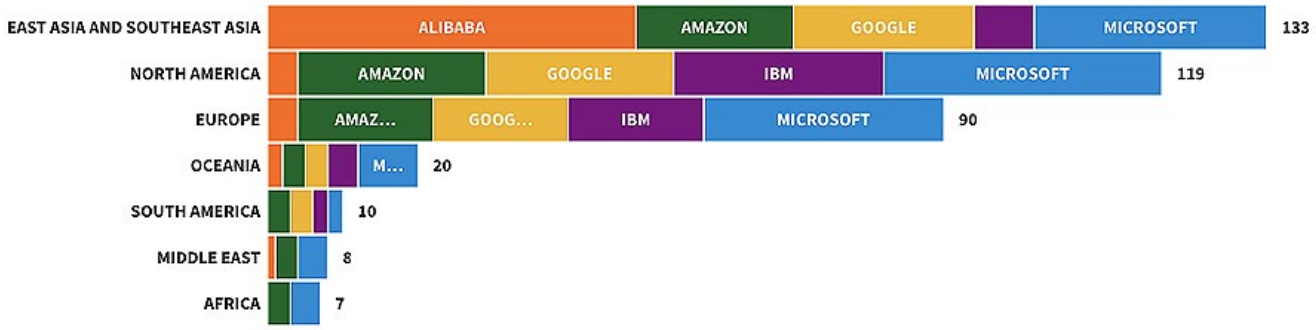

Source: Lily-Zimeng Liu

security risks for millions of users. Decisions made in a single company about how to evaluate the risk of certain vendors and the potential liabilities of their political system can ripple across an entire technology ecosystem.

The discussion of risk should not be taken lightly. The consequences from security events in the cloud grow ever larger as more security sensitive customers and their workloads move to the cloud, including operational defense information and intelligence community IT from the United States and the United Kingdom.[28] The aggregate number of companies and organizations that have shifted some or all of their IT infrastructure to the cloud continues to grow.[29] Cloud computing can mean greater efficiency, lower costs, and improved security, but it also comes with risk.

WHO SHOULD CARE?

The US Department of Homeland Security (DHS) and ENISA can do more to bring positive attention to the supply chain risks of cloud computing, working with their government partners and industry practitioners to make supply chain risk assessment and management programs more robust, consistent, and transparent to regulators and customers. Know-your-supplier rules in the industry are often still rooted in a 2001 public-private partnership program—Customs Trade Partnership Against Terrorism (CTPAT)[30]—focused on physical security and which only recently added a limited set of information assurance standards. Security auditing requirements for software, especially the specialized firmware used to control hardware, and popular container managers are inconsistent, developed by a vendor's internal security teams.

Certification and regulation of cloud services across the world, including in the United States, the UK, and Germany, focuses too much on enforcing outdated control sets and too little on the consistency of risk management programs between providers and working with industry to drive improvements in the security practices of key vendors like Intel and Cisco. Broader system security plans and private material submitted for bids on high-security clouds are not sufficient; providers should supply firm-wide risk management strategies and evaluations of performance that are accessible to all enterprise customers.

Customers can, and should, demand more information on their cloud providers' supply chain risk management programs.

28 Dan Gouré, "JEDI Will Be a Cloud Like No Other," RealClearDefense, June 7, 2019, https://www.realcleardefense.com/articles/2019/06/07/jedi_will_be_a_cloud_like_no_other_114485.html; US General Services Administration, GSA and DOD Award Defense Enterprise Office Solutions Cloud Contract, press release, August 29, 2019, https://www.gsa.gov/about-us/newsroom/news-releases/gsa-and-dod-award-defense-enterprise-office-solutions-cloud-contract;Cara McGoogan, "Ministry of Defence Switches to the Cloud As Microsoft Opens First UK Data Centres," *Telegraph*, September 7, 2016, https://www.telegraph.co.uk/technology/2016/09/07/ministry-of-defence-switches-to-the-cloud-as-microsoft-opens-fir/; Frank Konkel, "The Details About the CIA's Deal With Amazon," Atlantic, July 17, 2014, https://www.theatlantic.com/technology/archive/2014/07/the-details-about-the-cias-deal-with-amazon/374632/; AWS Public Sector Blog Team, "Announcing the New AWS Secret Region," *AWS*, November 20, 2017, https://aws.amazon.com/blogs/publicsector/announcing-the-new-aws-secret-region/; Naomi Nix and Ben Brody, "Microsoft Wins Lucrative Cloud Deal With Intelligence Community," Bloomberg, May 16, 2018, https://www.bloomberg.com/news/articles/2018-05-16/microsoft-wins-lucrative-cloud-deal-with-intelligence-community; "GCHQ Opts for Managed Cloud Services," *Intelligence Online*, February 13, 2019, https://www.intelligenceonline.com/government-intelligence/2019/02/13/gchq-opts-for-managed-cloud-services,108344708-art; Frank Konkel, "NSA 'Systematically Moving' All Its Data to The Cloud," Nextgov, June 21, 2018, https://www.nextgov.com/emerging-tech/2018/06/nsa-systematically-moving-all-its-data-cloud/149179/.

29 Gartner, Gartner forecasts worldwide public cloud revenue to grow 17.3 percent in 2019, press release, September 12, 2018, https://www.gartner.com/en/newsroom/press-releases/2018-09-12-gartner-forecasts-worldwide-public-cloud-revenue-to-grow-17-percent-in-2019; Shanhong Liu, "Public Cloud Computing Market Worldwide 2008-2020," Statista, August 9, 2019, https://www.statista.com/statistics/510350/worldwide-public-cloud-computing/.

30 US Customs and Border Protection, "CTPAT: Customs Trade Partnership Against Terrorism," February 23, 2020, https://www.cbp.gov/border-security/ports-entry/cargo-security/ctpat.

They should hold providers accountable through specific contractual requirements, including regular incident reporting summaries, disclosure of vendors used, and aggregate risk assessments for specific product and service chains. Sufficient clarity on these accountability and transparency standards would provide a solid basis for additions to the existing constellation of NIST and ISO supply chain risk management standards. Cloud providers themselves can do more to address these risks and publicly signal their intent through additional investments in cloud supply chain risk management programs, including new staff and a commitment to cross-industry standards. These investments would build on the positive, but uneven, progress across the industry over the past half-decade in confronting supply chain risks.

Cloud computing is verging on a utility. Its security against supply chain threats requires more than public-private working groups or rushed one-off risk assessments. Cloud computing security and the management of supply chain risk requires a stable and consistent public-private partnership with near-term aims to improve transparency and performance of risk management practices and long-term goals to drive proactive operational collaboration against high-consequence threats. The private sector holds critical expertise on how these technologies are built and operated. It is responsible for much of the capital investment and innovation that supports commercial cloud computing. The public sector bears direct responsibility for ensuring the safety and security of the public. It has legal means both to shape the markets for technology and pursue sources of risk extraterritorially. Cooperation is critical.

RECOMMENDATIONS

- **[The National Institute of Standards and Technology]** NIST should include cloud computing as a prominent case in the next revision of SP 800-161. It should further work with cloud service providers and the Federal Risk and Authorization Management Program (FedRAMP) office to strengthen supply chain family controls in the next revision of SP 800-53.

- **[DHS and ENISA]** DHS and ENISA should convene cloud providers on both sides of the Atlantic within the next year to specify supply chain risk management standards appropriate to the industry, including positive commitments of support for open-source software developers. They should formalize and publish these standards. DHS and ENISA should further collaborate to host representatives from participating firms annually for a multi-day session to review and update these standards.

- **[Cloud Customers]** Major cloud customers need to ask their providers for detailed supply chain risk management plans *and* auditable metrics of performance. This should include the biggest customers (Microsoft: Walmart,[31] Accenture,[32] and Boeing;[33] Google: Snap,[34] Ascension,[35] and Target[36]; Amazon: Netflix,[37] Adobe,[38] and Apple[39]). Customers should hold providers accountable for handing over this information through contractual requirements, including regular incident reporting summaries, disclosures of vendors being used despite negative security reporting or audit results, and aggregate risk assessments for specific product and service dependency chains.

- **[Cloud Service Providers]** The leadership of major cloud providers must recognize the diversity, difficulty, and persistence of supply chain risk to their infrastructure and appropriately resource their security and risk management efforts. This includes major providers like Amazon, Dell's VMWare, DXC, Google, IBM, Microsoft, Oracle, SAP, and Salesforce. These firms should commit to providing free security tools and services to any open-source projects integrated into their products; to developing cross-industry standards on supply chain risk management, including vendor security practices, auditing, and blacklisting; and sharing the proportion of total security dollars spent on supply chain risk management with customers on a quarterly basis.

31 Leo Sun, "IBM and Microsoft Are Upgrading Walmart's Digital Supply Chain," Motley Fool, September 30, 2018, https://www.fool.com/investing/2018/09/30/ibm-and-microsoft-are-upgrading-walmarts-digital-s.aspx.

32 Accenture, "Hybrid Cloud for Microsoft Azure," accessed August 5, 2020, https://www.accenture.com/us-en/service-hybrid-cloud-solution-microsoft.

33 Dawn Kawamoto, "Microsoft Azure Wins Boeing's Cloud Business," *InformationWeek*, July 19, 2016, https://www.informationweek.com/cloud/infrastructure-as-a-service/microsoft-azure-wins-boeings-cloud-business/d/d-id/1326316.

34 Tess Townsend, "This Is What Snap Is Paying Google $2 Billion for," *Vox*, March 1, 2017, https://www.vox.com/2017/3/1/14661126/snap-snapchat-ipo-spending-2-billion-google-cloud.

35 Tariq Shaukat, "Our Partnership with Ascension," *Google Cloud Blog*, November 19, 2019, https://cloud.google.com/blog/topics/inside-google-cloud/our-partnership-with-ascension.

36 Samantha Ann Schwartz, "Target Taps Google Cloud as Vendor Refuses to Be 'Distant Third' behind AWS, Microsoft," CIO Dive, July 25, 2018, https://www.ciodive.com/news/target-taps-google-cloud-as-vendor-refuses-to-be-distant-third-behind-aws/528503/.

37 Amazon, "Netflix on AWS," accessed May 18, 2020, https://aws.amazon.com/solutions/case-studies/netflix/.

38 Amazon, "Adobe Systems Case Study," accessed May 18, 2020, https://aws.amazon.com/solutions/case-studies/adobe/.

39 Nick Statt, "Apple's Cloud Business Is Hugely Dependent on Amazon," *Verge,* April 22, 2019, https://www.theverge.com/2019/4/22/18511148/apple-icloud-cloud-services-amazon-aws-30-million-per-month.

MYTH 3: ONLY AUTHORITARIAN STATES DISTORT THE PUBLIC CLOUD

Moves by democratic states to require data localization and isolated government clouds distort the public cloud and risk long-term harm to users and cloud technology. This risk is larger for small and medium-size countries without the internal cloud market sufficient to provide substantial leverage over large cloud providers. The public cloud is what cloud computing most often refers to: a globally accessible network of services whose various forms of data and equipment are spread around the world according to the demands of technical and financial efficiency. Public does not suggest all clouds are the same—each provider owns and operates their own infrastructure—but instead that a variety of users access that same infrastructure. This is cloud computing's core premise: greater efficiency from pooled resources. This same server or high-performance computing cluster can be in use nearly all of the time, shared in fractions between many people.

Countries and some companies are increasingly pushing back on this global accessibility, arguing that infrastructure used by their governments and sensitive industries should be isolated from the public cloud and located in specific regions or jurisdictions.[40] This concern is rooted in the recurring idea that the public cloud is more dangerous than infrastructure reserved for a single organization or set of users or that to dominate the cloud industry, countries must penalize foreign competitors and demand local infrastructure. Accuracy aside, this perception drives many government organizations and new cloud customers in regulated industries to demand their own separate infrastructure.

The myth of this distortion of the public cloud is that it is pursued only by governments with a preexisting interest in disrupting or controlling their citizens' use of technology—for example, Chinese requirements that all cloud infrastructure be hosted in China by local companies or joint ventures[41] and Russia's 2015 data localization laws.[42] The reality is more complicated. Germany and France both have data localization rules, in administrative and legislative form, that restrict the storage or transmission of certain data outside their borders. The French rule, contained in a ministerial circular produced in 2017, requires any system involved in surveillance of electronic communications to be based in France and blocks any data subject to court proceedings from leaving the country.[43] In 2000, France was one of the first countries to bring suit against a major internet service firm, Yahoo, to bar French users from visiting its auction sites that sell Nazi memorabilia as part of a broader push to block hate speech on the internet.[44] More recently, France has passed laws allowing a state agency to filter for objectionable or pirated content, including a 2011 technical proposal to install surveillance equipment in routers used

Cloud Deployment Models

While the public cloud is the most dominant form of cloud computing, there are other deployment models. Private clouds take the same design of infrastructure and management systems but deploy them for only one organization, operated either by the provider or that consuming organization. Community clouds are effectively public clouds whose tenants are users of multiple associated organizations, like Microsoft's Azure Government—accessible to federal agencies as well as state and local governments but not the private sector. The economics of cloud computing improve as the user base grows, so public cloud deployments tend to have the widest variety of services and the most cutting-edge features. One of the biggest sources of growth in private and community clouds are national security and intelligence agencies, including in the United States and in the UK, as these groups have expressed discomfort with their lack of control over the public cloud's security standards and inherent multitenancy.

40 Sheenagh Matthews, "A German Cloud for German Companies," Bloomberg, May 19, 2016, https://www.bloomberg.com/news/articles/2016-05-19/a-german-cloud-for-german-companies; Liam Tung, "Meet GAIA-X: This Is Europe's Bid to Get Cloud Independence from US and China Giants," *ZDNet*, June 8, 2020, https://www.zdnet.com/article/meet-gaia-x-this-is-europes-bid-to-get-cloud-independence-from-us-and-china-giants/.

41 *China: Doing Business and Investing in China Guide* (United States: International Business Publications, 2007).

42 Duane Morris, "Russia's New Personal Data Localization Law Goes into Effect in September 2015," June 15, 2015, https://www.duanemorris.com/alerts/russia_new_personal_data_localization_law_into_effect_september_2015_0615.html.

43 France's Ministry of Interior and Ministry of Culture and Communication, "An informational note of April 5, 2016 regarding cloud computing," accessed February 28, 2020, https://francearchives.fr/file/f7ace4517613a246583fd2dd673a0e6d0f86c039/static_9151.pdf; Nigel Cory, "Cross-Border Data Flows: Where Are the Barriers, and What Do They Cost?" Information Technology and Innovation Foundation, May 1, 2017, https://itif.org/publications/2017/05/01/cross-border-data-flows-where-are-barriers-and-what-do-they-cost.

44 Kieren McCarthy, "Yahoo! Legally Obliged to Ban the French?" *Register*, November 6, 2000, https://www.theregister.com/2000/11/06/yahoo_legally_obliged_to_ban/.

Figure 12: European Data Center Map

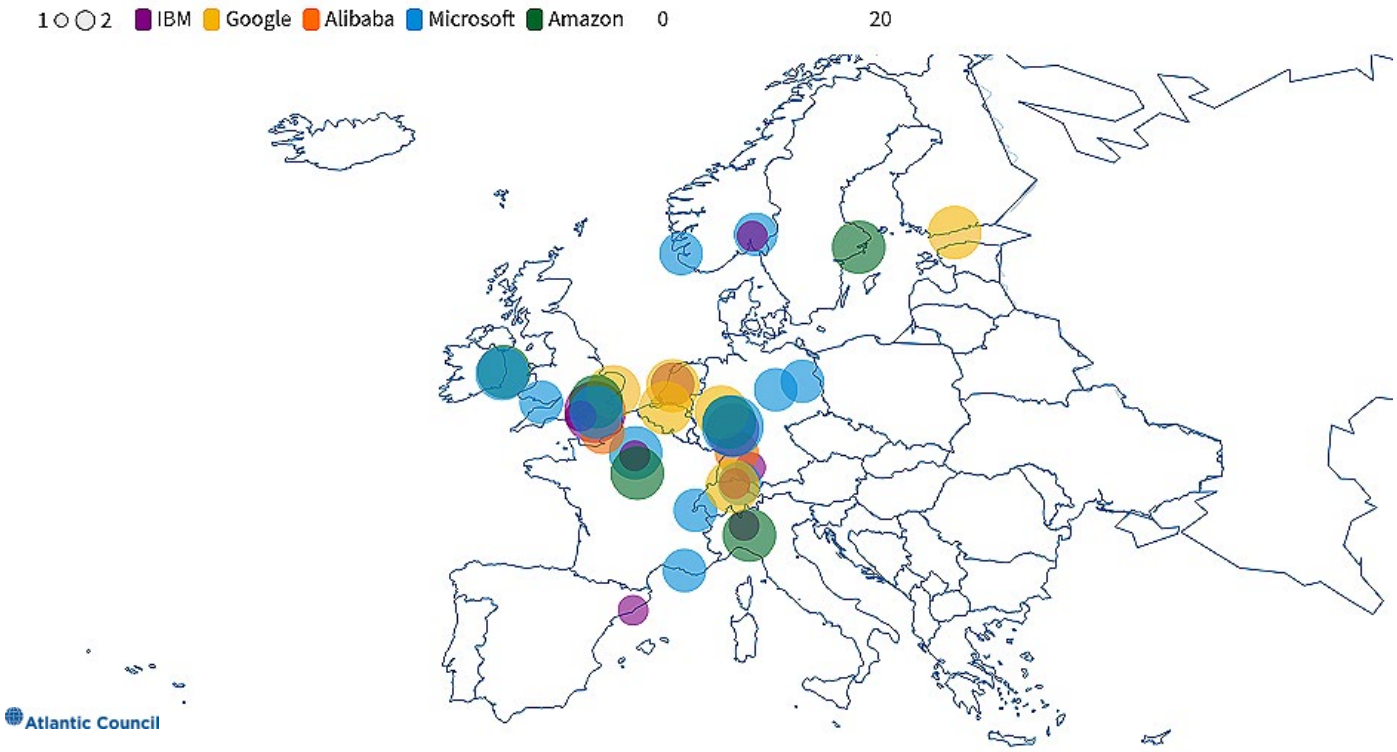

Source: Lily-Zimeng Liu

by internet service providers around the country.[45] France is not the only offender. In 2018, Germany implemented the NetzDG law which forces for-profit websites to remove "offensive" content within 24 hours or face the consequence of fines for a failure to act.[46]

The EU's GAIA-X initiative proposes to create a European cloud, suggesting that public clouds from US providers or built to US originated standards are less secure, less trustworthy, or reflect values at odds with those held in Europe.[47] Its not clear if the EU has a sufficient domestic market of cloud service providers to build a 'European cloud' with European companies so the end result may just be an American built cloud with slightly different designs but isolated from the public cloud these providers offer to the rest of the world.[48] This isolation rarely provides the supposed cybersecurity benefits and yet is a driver for private clouds in other domains.

45 Nicola Lucchi, "Regulation and Control of Communication: The French Online Copyright Infringement Law (HADOPI)," April 20, 2011, *Cardozo Journal of International and Comparative Law (JICL)*, Vol. 19 (2011), Max Planck Institute for Intellectual Property & Competition Law Research Paper No. 11-07, available at SSRN: https://ssrn.com/abstract=1816287; Christophe Auffray, "Moyens de Sécurisation: La Hadopi Lance Une Nouvelle Consultation," *ZDNet France*, April 20, 2011, https://www.zdnet.fr/actualites/moyens-de-securisation-la-hadopi-lance-une-nouvelle-consultation-39760125.htm.

46 Diana Lee, "Germany's NetzDG and the Threat to Online Free Speech," Yale Law School, October 10, 2017, https://law.yale.edu/mfia/case-disclosed/germanys-netzdg-and-threat-online-free-speech.

47 Federal Ministry for Economics Affairs and Energy, "GAIA-X: Policy Rules and Architecture of Standards" (Federal Ministry for Economic Affairs and Energy Public Relations Division, 2020), https://www.bmwi.de/Redaktion/EN/Publikationen/gaia-x-policy-rules-and-architecture-of-standards.html.

48 Janosch Delcker and Melissa Heikkilä, "Germany, France Launch Gaia-X Platform in Bid for 'Tech Sovereignty'," June 5, 2020, https://www.politico.eu/article/germany-france-gaia-x-cloud-platform-eu-tech-sovereignty/.

Pressure on cloud providers is significant when it comes to isolation of government data. Authorities in the United States,[49] the UK,[50] Germany,[51] France,[52] and Australia[53] require cloud providers to build local data centers to deliver services and host data for public organizations. The debate continues on whether or not physically isolating government data is necessary as a core security practice or (at great cost) simply makes other policies, like requiring specific cryptographic standards, easier to enforce. Like localization, there is little to suggest physical isolation offers inherent security benefits. It becomes easier for governments to levy exotic or exacting requirements when they demand the use of separate infrastructure, but this is altogether different from the benefits of such isolation alone.

Democratic governments have distorted the public cloud—the push for varying levels of localization and isolated government clouds directly harms the economic viability of the public cloud. Where government practices trickle down to the private sector this harm is magnified; the government's approach, especially for security and intelligence agencies, can become perceived as the gold standard for high-security models of cloud computing. This perception can ripple through regulated industries and far beyond government.

The problem posed by these ripples is that the economics of the public cloud underlie much of its organizational utility and long-term technical potential. Tying together millions of customers across the world creates a market for both marginal improvements and innovation in technology at sustainable cost to private enterprise. The maintenance of competition between firms like Amazon, Google, and Microsoft is part of what has accelerated the adoption and functionality of cloud computing services over the past decade. Localization and undue isolation requirements threaten the broader economics of the public cloud.

These requirements also pose a risk to users of cloud services. As the public cloud market becomes more fragmented, countries risk advantaging specific providers and unintentionally buying into some degree of vendor lock-in. Companies may invest in localized infrastructure for particular markets, resulting in short-term gains but long-term stagnation for that jurisdiction. Where data cannot cross borders and continents, much of the flexibility in the public cloud model and potential value in a competitive market is drained away.

WHO SHOULD CARE?

This phenomenon should be of particular concern to policy makers in small and medium-size countries who recognize that their influence over cloud providers is comparatively less than that of larger neighbors. Argentina, the Netherlands, and South Africa, for example, may not have the economic clout alone to drive major policy change. Collective action, however, by regions or supranational groupings can drive more considered and coherent policy dialogue with cloud providers. The Three Seas region could be an exemplary model and has the potential for a jointly regulated cloud computing market to set standards with cloud providers, rather than leaving each country to play for itself against an Amazon or Microsoft.

Cloud providers should also take note of this trend. Politics is dictating technical requirements and engineering decisions, sometimes because of genuine, but unmet, concern about security and control over the underlying infrastructure, and at other times for no better reason than a lack of well-informed regulators. Cloud providers should prioritize the availability of information about their infrastructure and operation of their services, including moderately technical explanations of system design and internal security practices. Too much engagement today is mired in lobbying efforts and marketing miasma.

RECOMMENDATIONS

- **[Three Seas Initiative and Similar Organizations]** Small and medium-size states should work together in appropriate regional collectives to build collective negotiating positions with cloud providers. Crafting a single negotiating entity will help concentrate market power for leverage and provide a level playing field between neighboring countries to tap into investment and service offerings from cloud providers. Where these collectives emerge to govern the handing of data and use of cloud services in a supra-national fashion they will represent further progress.

- **[Cloud Service Providers]** Invest in global policy teams and share information on the design, construction, and operation of cloud computing services with policy

49 Tom Keane, "Announcing New Azure Government Capabilities for Classified Mission-Critical Workloads," Microsoft Azure, October 17, 2017, https://azure.microsoft.com/en-us/blog/announcing-new-azure-government-capabilities-for-classified-mission-critical-workloads/.

50 Amazon, "G-Cloud UK," accessed January 15, 2020, https://aws.amazon.com/government-education/g-cloud-uk/.

51 Catherine Stupp, "Germany to Set up 'Bundescloud,'" EURACTIV, August 20, 2015, https://www.euractiv.com/section/digital/news/germany-to-set-up-bundescloud/.

52 Pat Brans, "French Public Sector's Never-Ending Struggle with the Cloud," *ComputerWeekly*, July 18, 2016, https://www.computerweekly.com/news/450300488/French-public-sectors-never-ending-struggle-with-the-cloud.

53 Digital Transformation Agency, Australian Government, "Secure Cloud Strategy," 2017, https://dta-www-drupal-20180130215411153400000001.s3.ap-southeast-2.amazonaws.com/s3fs-public/files/cloud/secure-cloud-strategy.pdf.

What About China?

China is home to the fourth of the four hyperscale cloud companies. Alibaba, whose dominant domestic market is many times larger than its closest competitor, has begun to expand abroad. The company has announced it will spend $28 billion over the next three years to compete with major Western providers and has already added data centers on both coasts of the United States as well as two each in Australia, Germany, India, Japan, the United Arab Emirates (UAE), and the UK.[1] *No major Western cloud provider provides services directly in China. China levies a set of requirements that these foreign providers participate in the mainland market through joint ventures with domestic companies that allow for majority Chinese ownership over the sale of foreign products and services.*[2] *These joint ventures, which use Western cloud providers' system designs and infrastructure but work within a different legal regime than other countries, are operated by Chinese personnel and impose small, but significant, technical changes (e.g., domestic encryption schemes) to how the cloud services are built.*[3] *Running through such joint ventures is a condition for access to the Chinese market. As a result, companies like Amazon*[4] *and Microsoft*[5] *have found local partners like Sinnet and 21 Vianet, respectively.*

1 Arjun Kharpal, "China's Alibaba to Invest $28.2 Billion in Cloud Infrastructure As It Battles Amazon, Microsoft," CNBC, April 20, 2020, https://www.cnbc.com/2020/04/20/alibaba-to-invest-28-billion-in-cloud-as-it-battles-amazon-microsoft.html.

2 "Doing Business in China," World Services Group, September, 2006, https://www.worldservicesgroup.com/guides/Doing%20Business%20in%20China.pdf.

3 Max Emelianov, "China's New Cybersecurity Law Is Bad News for Data Center Security," *DCD*, January 7, 2020, https://www.datacenterdynamics.com/en/opinions/chinas-new-cybersecurity-law-bad-news-data-center-security/.

4 Cate Cadell, "Amazon Sells Off China Cloud Assets As Tough New Rules Bite," Reuters, November 13, 2017, https://www.reuters.com/article/us-china-amazon-cloud/amazon-sells-off-china-cloud-assets-as-tough-new-rules-bite-idUSKBN1DE0CL.

5 Jason Verge, "21Vianet to Operate Azure Cloud in China into 2018," Data Center Knowledge, April 2, 2015, https://www.datacenterknowledge.com/archives/2015/04/02/21vianet-will-operate-azure-cloud-in-china-into-2018/.

makers. The pressure to close sales and access new markets threatens to pull cloud providers into costly commitments and distortions to their technical architecture. Facilitating healthy dialogue between policy makers and corporate policy staff who can access engineers and act as honest brokers, without supporting specific sales engagements, will yield more capable policy makers and a better business environment.

MYTH 4: CLOUD PROVIDERS DO NOT INFLUENCE THE SHAPE OF THE INTERNET

The evolution of cloud computing is deeply intertwined with that of the internet, yet there is a pervasive myth that cloud providers do not shape that tangle of tubes. The internet provided pathways on which to share computing, storage, and networking resources across the globe. The bandwidth and accessibility of "broadband" internet is what made cloud computing a viable market beyond enterprise sales—broadening its utility and flexibility. The accessibility and cost effectiveness of IaaS has helped companies like Netflix rapidly grow to become substantial portions of the internet's total traffic.[54] The functionality embedded in PaaS offerings has offloaded a host of complex programming tasks from developers to cloud providers, concentrating dependence for tasks like machine translation and data storage in an increasingly small number of firms and paths to their networks. As the largest cloud providers' infrastructure grows, ever more of this traffic runs between data centers and globe-spanning networks controlled by a single firm permitting things like proprietary changes to core internet transport protocols.[55] Broad adoption of cloud computing has influenced the shape of the web itself.

This is not a new phenomenon. Cloud service originated in early mainframe computers and the era of "time-sharing," when several users would access the same machine. The first implementation of such a time-sharing system came in 1959 from Stanford computer scientist and A.M. Turing Award recipient John McCarthy.[56] The first nodes in networks that would lead to the internet were not switched on until 1967 in a pilot program in London and 1969 in the ARPANET in the United States.[57]

54 Todd Spangler, "Netflix Bandwidth Consumption Eclipsed by Web Media Streaming Applications," *Variety*, September 10, 2019, https://variety.com/2019/digital/news/netflix-loses-title-top-downstream-bandwidth-application-1203330313/.

55 Mattias Geniar, "Google's QUIC Protocol: Moving the Web from TCP to UDP," July 30, 2016, https://ma.ttias.be/googles-quic-protocol-moving-web-tcp-udp; Ian Swett and Michael Behr, "Introducing QUIC Support for HTTPS Load Balancing," Google Cloud, July 13, 2018, https://cloud.google.com/blog/products/gcp/introducing-quic-support-https-load-balancing.

56 John McCarthy, "Reminiscences on the Theory of Time-Sharing," Stanford University, 1983, http://jmc.stanford.edu/computing-science/timesharing.html.

57 Ivan P. Kaminow and Tingye Li, *Optical Fiber Telecommunications IV-B Systems and Impairments* (San Diego: Academic Press, 2002).

The internet had to change its shape and grow as cloud computing took off. Under the prevailing internet model of the 1990s, data flowed back and forth between users and servers operated by companies like Yahoo. This path—user-server-user—was abstracted on networking diagrams as a vertical flow of traffic into and out of a company's infrastructure and was referred to as north-south traffic, as opposed to data moving between machines operated by a single service, represented by a horizontal line and referred to as east-west traffic.

Prior to Amazon's introduction of the first major cloud computing service, Amazon Web Services (AWS), in 2006,[58] the bulk of web traffic was of the north-south variety—moving between user and service. With AWS and the next decade's growth of cloud computing, more and more traffic began to move between servers to process more complex requests before moving back to the user. Users were now accessing multimedia files, collaboratively editing documents, and accessing databases with hundreds or thousands of different users, often at the same time, instead of loading webpages. More and more traffic started to move between servers, east to west, with machines responsible for different aspects of the operation: storing a file, processing an image, or reconciling changes to a shared document. These servers were no longer located at different facilities and so this traffic moved across the globe between them, over infrastructure leased or owned outright by major cloud service providers. Data moving between cloud providers might even pass directly from infrastructure owned by one to the other, without ever touching the public internet.

Fiber optic cables carry the bulk of internet traffic and offer massive bandwidth compared to satellite or microwave radio transmission. Even that massive bandwidth, on the order of tens to hundreds of terabytes per second, is not enough.[59] The volume of traffic between cloud data centers has grown faster than traffic over the internet in almost every year since 2006.[60] To satisfy this immense volume, Microsoft and Google have built or leased massive amounts of bandwidth, including laying new transatlantic fiber optic cables.[61] Amazon has largely relied on existing routes and internet infrastructure, while Microsoft and Google have gone on to become among the largest operators of fiber cable networks in the world.[62]

Google has even proposed its own new internet protocols used within its cloud computing networks. The first, Firepath, is a routing protocol used to manage the flow of data within Google's internet-scale array of data centers.[63] The second, QUIC, developed in 2011, is a replacement for the transport protocol TCP and is used within Google's network and to communicate over the internet with core products like Chrome.[64] Despite still being on the path to formal adoption by the internet's core technical standards body, the Internet Engineering Task Force (IETF), QUIC is already used in upwards of 10 percent of internet traffic.[65]

The ownership and operation of physical infrastructure is only part of cloud providers' influence on the internet; they also shape the marketplace for organizations and developers to deploy their technology. The experience of Signal, a widely used encrypted communications application popular among political dissidents, journalists, and the privacy-minded, is

58 Ron Miller, "How AWS Came to Be," *TechCrunch*, July 2, 2016, https://techcrunch.com/2016/07/02/andy-jassys-brief-history-of-the-genesis-of-aws/.

59 Valey Kamalov, Ljupcho Jovanovski, Vijay Vusirikala, Eduardo Mateo, Yoshihisa Inada, Takaaki Ogata, Kenichi Yoneyama, Pascal Pecci, David Seguela, Olivier Rocher, and Hidenori Takahashi, "FASTER Open Submarine Cable," 2017 European Conference on Optical Communication (ECOC), Gothenburg (2017) 1-3, doi: 10.1109/ECOC.2017.8346076; Lucy Spencer, "Google and Facebook Invest in Trans-Pacific Infrastructure," *ITU News*, October 18, 2016, https://news.itu.int/google-and-facebook-invest-in-trans-pacific-infrastructure/.

60 Arjun Singh, Joon Ong, Amit Agarwal, Glen Anderson, Ashby Armistead, Roy Bannon, Seb Boving, Gaurav Desai, Bob Felderman, Paulie Germano, Anand Kanagala, Hong Liu, Jeff Provost, Jason Simmons, Eiichi Tanda, Jim Wanderer, Urs Hölzle, Stephen Stuart, and Amin M Vahdat, "Jupiter Rising: A Decade of Clos Topologies and Centralized Control in Google's Datacenter Network," Communications of the ACM, August 1, 2016, https://dl.acm.org/doi/abs/10.1145/2975159.

61 Microsoft Reporter, "A Cable Stretching 4,000 Miles Between the US and Spain Is the Key to a High-Speed Future," Microsoft News Centre Europe, September 22, 2017, https://news.microsoft.com/europe/2017/09/22/a-cable-stretching-4000-miles-between-the-us-and-spain-is-the-key-to-a-high-speed-future/.

62 Cade Metz, "Facebook and Microsoft Are Laying a Giant Cable Across the Atlantic," *Wired*, May 26, 2016, https://www.wired.com/2016/05/facebook-microsoft-laying-giant-cable-across-atlantic/;
Tyler Cooper, "Google and Other Tech Giants Are Quietly Buying Up the Most Important Part of the Internet," VentureBeat, April 6, 2019, https://venturebeat.com/2019/04/06/google-and-other-tech-giants-are-quietly-buying-up-the-most-important-part-of-the-internet/.

63 Cade Metz, "Revealed: The Secret Gear Connecting Google's Online Empire," *Wired*, June 17, 2015, https://www.wired.com/2015/06/google-reveals-secret-gear-connects-online-empire/.

64 Adam Langley, Alistair Riddoch, Alyssa Wilk, Antonio Vicente, Charles Krasic, Dan Zhang, Fan Yang, Fedor Kouranov, Ian Swett, Janardhan R Iyengar, Jeff Bailey, Jeremy Dorfman, Jim Roskind, Joanna L Kulik, Patrik Westin, Raman Tenneti, Robbie Shade, Ryan Hamilton, Victor Vasiliev, Wan-Teh Chang, and Zhongyi Shi, "The QUIC Transport Protocol: Design and Internet-Scale Deployment," SIGCOMM '17: Proceedings of the Conference of the ACM Special Interest Group on Data Communication (August 2017): 183–196, https://dl.acm.org/doi/abs/10.1145/3098822.3098842; James Sanders, "For the Internet to Keep Growing, We Need a Next-Gen TCP," TechRepublic, January 26, 2016, https://www.techrepublic.com/article/for-the-internet-to-keep-growing-we-need-a-next-gen-tcp/.

65 Jan Rüth, Ingmar Poese, Christoph Dietzel, and Oliver Hohlfeld, " A First Look at QUIC in the Wild," *Passive and Active Measurement* (2018): 255-268, https://link.springer.com/chapter/10.1007/978-3-319-76481-8_19.

Figure 13: Regional Distribution of Data Centers by Provider

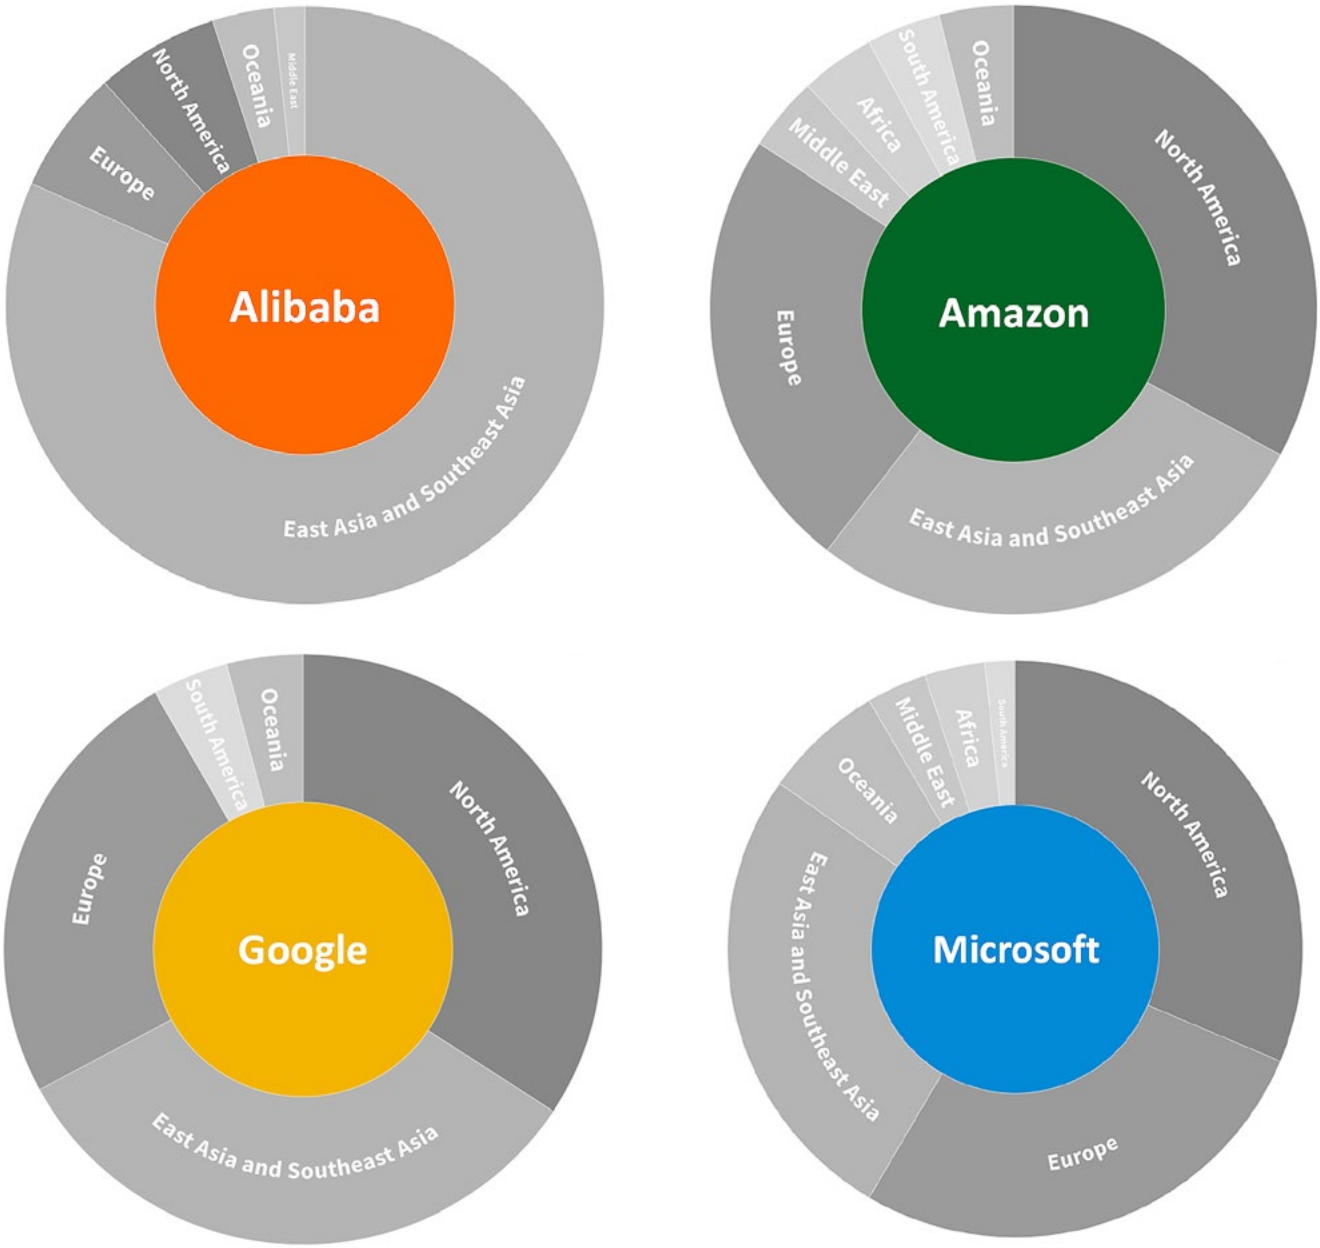

instructive. Because of its use to preserve the secrecy of communications, Signal has become popular in jurisdictions where censorship and targeting of activists is disturbingly frequent, like Egypt, Iran, Qatar, and the United Arab Emirates (UAE). These countries work to block access to censorship-circumvention technology, like encrypted communications apps, presenting a challenge to the developers.

In several of these countries, Signal leveraged a quirk in how software communicates over the internet to disguise itself. When firewalls and censorship programs in these countries went to filter and block internet traffic from banned applications like Signal, the quirk allowed Signal's traffic to appear to be headed to a larger internet service firm like Google or Amazon. The quirk, a technique called domain fronting, effectively allowed Signal to disguise its traffic. This forced the censoring country into a tough position. Blocking traffic from a popular but relatively small application like Signal posed a very different problem to blocking all traffic to Google or Amazon. This technical workaround protected

access to Signal, and other services like it, in places like Egypt, Iran, Qatar, and the UAE where internet traffic is regularly censored and some applications blocked.[66]

Signal's strategy worked well for years but ran into trouble in May 2018 when Google changed its terms of service and tweaked its networking technology to ban and block the practice. The politics behind Google's change of heart, which merit extended discussion elsewhere, are supposed to have come down to a debate over how to manage the company's risk of violating US sanctions against Iran while providing some internet-based services to users there. Signal attempted to shift operations to Amazon and take up the same practice but was quickly stymied, this time by a policy change and threats the application would be kicked off the service entirely if it failed to adhere. Amazon at the time faced immense pressure from the Russian government (as did Microsoft and to a lesser extent Google) to block service to the encrypted communications service Telegram. Russian authorities blocked access to many Amazon services and internet addresses in the country in an effort to compel cooperation and the company sought some way to relieve pressure on the non-trivial Russian market.

Domain fronting is a well-researched[67] technique employed by some malicious actors[68] as well as a variety of censorship circumvention tools[69] operating across the world, including several in China.[70] Amazon, Google, and others tried to label the change a technical correction—fixing something they never intended to support in the first place—but the decision had concerning policy implications and real consequences for people using Signal in these countries. Cloud providers hold this kind of immediate and far-reaching influence in setting the rules for services on their platform; rules which reflect a long tradition of corporate autonomy to set their terms of use but which fail to account for the awesome influence hyperscale cloud firms have on the internet.

Other comparably sized cloud computing companies, Alibaba and Microsoft, were effectively silent on the issue. It is not clear whether Alibaba would have had sufficient traffic flowing into the Arab Gulf states and Egypt, where this protection from censorship was sorely needed, to play a role in hiding Signal or other similar apps' traffic. Microsoft appears to have remained very still, attempting to escape notice, a strategy the company has deployed with remarkable success over the last several years of "tech-lash."[71] Despite a strongly worded congressional letter to Amazon and Google asking for details on the decision to block domain-fronting, there was little overt policy response and the debate faded away.[72] Signal was out of options in its search for companies that could both host the service and act as a front for its users.

Amazon, Google, and Microsoft wield influence which is stunning in its scope and finality over the architecture and even some operation of the internet. Debates over content moderation and regulation are warranted, but they focus on a surface level of the internet where competition and substitution are far easier. Who remembers Yammer or Yabbly or Yahoo? Facebook and YouTube have grown to be dominant, but their reign is not permanent. Control over the infrastructure on which these services run and influence over their deployment and operation are far more compelling. Changes and censorship at this level are harder to observe, changes are stickier, and while content can generally find a new platform or service to call home, there are few internets to choose between.

66 Barney Warf, "Geographies of Global Internet Censorship," *GeoJournal* 76 (November 2010): 1-23, https://link.springer.com/article/10.1007/s10708-010-9393-3 ; Jonathan Zittrain, Robert Faris, Helmi Noman, Justin Clark, Casey Tilton, and Ryan Morrison-Westphal, *The Shifting Landscape of Global Internet Censorship*, Berkman Klein Center Research Publication No. 2017-4, Harvard Public Law Working Paper No. 17-38 (June 2017), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2993485; Jessica Conditt, "Encrypted Chat App Signal Circumvents Government Censorship," Engadget, December 21, 2016, https://www.engadget.com/2016-12-21-signal-egypt-uae-censorship-block-domain-fronting.html; Sandra Pattison, "Internet Censorship 2020: Find Out Where Repression Reigns: An Overview of Internet Censorship in 149 Countries Around the World," Cloudwards, June 12, 2020, https://www.cloudwards.net/internet-censorship/; Moxie Marlinspike, "A Letter from Amazon," Signal, May 1, 2018, https://signal.org/blog/looking-back-on-the-front/; Helmi Noman, "Internet Censorship and the Intraregional Geopolitical Conflicts in the Middle East and North Africa," Internet Monitor, January 15, 2019, https://thenetmonitor.org/bulletins/internet-censorship-and-the-intraregional-geopolitical-conflicts-in-the-middle-east-and-north-africa.

67 David Fifield, Chang Lan, Rod Hynes, Percy Wegmann, and Vern Paxson, "Blocking-resistant Communication Through Domain Fronting," *Proceedings on Privacy Enhancing Technologies*, 2015 (2), https://content.sciendo.com/configurable/contentpage/journals$002fpopets$002f2015$002f2$002farticle-p46.xml.

68 Chris Brook, "APT29 Used Domain Fronting, Tor to Execute Backdoor," threatpost, March 27, 2017, https://threatpost.com/apt29-used-domain-fronting-tor-to-execute-backdoor/124582/.

69 Access Now, Message to Google and Amazon on domain fronting: You break it, you bought it, press release, May 2, 2018, https://www.accessnow.org/message-to-google-and-amazon-on-domain-fronting-you-break-it-you-bought-it/.

70 Fifield et al. "Blocking-resistant"; Access Now, Message.

71 Eve Smith, "The techlash against Amazon, Facebook and Google—and what they can do," *The Economist,* January 20, 2018, https://www.economist.com/briefing/2018/01/20/the-techlash-against-amazon-facebook-and-google-and-what-they-can-do.

72 Ron Wyden and Marco Rubio, Letter to Jeff Bezos and Larry Page, U.S. Senate, July 17, 2018, https://www.wyden.senate.gov/imo/media/doc/Wyden%20Rubio%20Letter%20to%20Amazon%20+%20Alphabet%20re%20Domain%20Fronting%20Ban.pdf.

WHO SHOULD CARE?

The concentration of ownership over cloud infrastructure and the ability for hyperscale firms to aggregate expertise while supporting capital-intensive innovation is a benefit to users and the technology ecosystem. The influence of these same firms over the shape of the internet and the opacity of their internal decision-making should do more than raise eyebrows amongst the policy community and civil society. This includes national security organs in the EU and the United States, including agencies working in support of the US National Cyber Strategy, which calls for an open, interoperable, reliable, and secure internet.[73] Policy makers should be building ties within these companies, staying in regular contact to understand whose portfolio and decisions impact which aspects of the political and social landscape. The same diligence applied in mapping the hierarchy of a partner diplomatic organization or foreign military apparatus is necessary here to realize the social network of influence over the internet.

Civil society organizations should work with major cloud providers to design and implement more effective transparency measures. These should provide a clear sense of the public good to a firm's leadership and greater confidence to users in the processes underlying major decisions like those which led to a loss of Signal in already censorship happy regions of the world. The strength of these cloud providers as sources of technical innovation and their capacity to manage massive amounts of infrastructure is an asset to the public and the open-source ecosystem. At the same time, these firms are struggling to take on the mantle of leadership required of companies that provide such a common and critical good.

Most importantly, the cloud providers themselves must recognize this magnificent influence. More than an opportunity to exercise well-worn tropes of corporate social responsibility, cloud providers must take care to provide transparency in their expectations, operations, and decision-making around their service. Cloud as a utility service is more than a product on the open market. It demands governance with a semblance of democratic accountability. Companies like Amazon, Google, and Microsoft can recognize the torturous confusion and real business harm endured by the likes of Facebook and Twitter in the first round of "negotiations" between policy makers and corporate leaders over how these companies should impact society. There is an opportunity to get ahead of this issue and assert leadership in a vacuum.

RECOMMENDATIONS

- **[Civil Society]** Civil society organizations should leverage the increasing focus of major cloud providers on large government contracts to push for reporting transparency on content or services removed for terms of service violations, public and regularly reported impact assessments of technical changes that affect software or services focused on encrypted communications and censorship circumvention, and encourage the appointment of an ombudsman at each major provider with adequate budget and authority to investigate changes and activity that could harm the public interest.

- **[Cloud Service Providers]** Cloud providers should bring engineers to the table. For any conversations with policy makers, government affairs teams must always be present with technical staff who can speak to specific practices, threats, and opportunities. Anything less risks the integrity of the public-private partnership and resulting pendulum between disruptive intervention and neglect by policy makers.

- **[Congressional Research Service]** The CRS should produce an appropriately resourced study on the national security and economic implications of ownership of the network infrastructure in particular the role of hyperscale cloud computing and major internet service firms.

73 Executive Office of the President of the United States, "National Cyber Strategy of the United States of America," September 2018, https://www.whitehouse.gov/wp-content/uploads/2018/09/National-Cyber-Strategy.pdf.

CONCLUSION

Cloud providers are geopolitical actors. Their influence reaches beyond technology markets. They influence the pace and direction of economic growth, shape international security competition, and mediate access to technologies which today inform changes in the global balance of power. The cloud is itself influenced by these same geopolitics—all of these wires and cables and boxes and bodies and their customers have to live somewhere and these jurisdictions have rules and goals all their own. The geopolitics of cloud computing demands a longer parchment than is present here, but this paper serves to highlight, and disprove, four important myths in the relationship: 1) all data is created equal; 2) cloud computing is not a supply chain risk; 3) only authoritarian states distort the public cloud; and 4) cloud providers do not influence the shape of the internet.

Technology shifts social and political dynamics. Cloud computing is no different, but it is more than a massive commercial phenomenon—it influences the trajectory of nations and the conduct of statecraft. Conflicts in cyberspace are already being fought in and through the cloud. As providers continue to concentrate unparalleled quantities of computing resources and user data, they will only grow in importance—as governors, as battlefields, and as magnificent engines of complexity. In the meantime, policy makers must arm themselves with more than myths as they seek to grapple with the geopolitics of the cloud.

ACKNOWLEDGMENTS

The author would like to thank Simon Handler, Lily Liu, and TJ Zuo for the graphics, design, and editing and Stewart Scott, Kath Kennelly, David Hoffman, and several anonymous reviewers for their comments. Thank you to Amanda Craig and Ryan Socal for a high-speed mentorship on cloud policy and special thanks to Angela McKay for the invitation to first spelunk into this world.

ABOUT THE AUTHOR

Dr. Trey Herr is the Director of the Cyber Statecraft Initiative under the Scowcroft Center for Strategy and Security at the Atlantic Council. His team works on the role of the technology industry in geopolitics, cyber conflict, the security of the internet, cyber safety, and growing a more capable cybersecurity policy workforce. Previously, he was a Senior Security Strategist with Microsoft handling cloud computing and supply chain security policy as well as a fellow with the Belfer Cybersecurity Project at Harvard Kennedy School and a non-resident fellow with the Hoover Institution at Stanford University. He holds a PhD in Political Science and BS in Musical Theatre and Political Science

Atlantic Council Board of Directors

The Atlantic Council is a nonpartisan organization that promotes constructive US leadership and engagement in international affairs based on the central role of the Atlantic community in meeting today's global challenges.

Atlantic Council

1030 15th Street, NW, 12th Floor,
Washington, DC 20005

(202) 463-7226, www.AtlanticCouncil.org