Atlantic Council
CYBER STATECRAFT INITIATIVE
DFRLab

Critical Infrastructure and the Cloud:
POLICY FOR EMERGING RISK

Tianjiu Zuo, Justin Sherman,
Maia Hamin, and Stewart Scott

CYBER STATECRAFT INITIATIVE

DFRLab

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.

The mission of the **Digital Forensic Research Lab (DFRLab)** is to identify, expose, and explain disinformation where and when it occurs using open-source research; to promote objective truth as a foundation of government for and by people; to protect democratic institutions and norms from those who would seek to undermine them in the digital engagement space; to create a new model of expertise adapted for impact and real-world results; and to forge digital resilience at a time when humans are more interconnected than at any point in history, by building the world's leading hub of digital forensic analysts tracking events in governance, technology, and security.

This report is written and published in accordance with the Atlantic Council Policy on Intellectual Independence. The author is solely responsible for its analysis and recommendations. The Atlantic Council and its donors do not determine, nor do they necessarily endorse or advocate for, any of this report's conclusions.

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

For more information, please visit
www.AtlanticCouncil.org.

ISBN-13: 978-1-61977-283-0

July 2023

Design by: Donald Partyka and Anais Gonzalez

Cover image by: Red Edge Graphics

CYBER STATECRAFT INITIATIVE

DFRLab

CRITICAL INFRASTRUCTURE AND THE CLOUD:

Policy for Emerging Risk

Tianjiu Zuo, Justin Sherman, Maia Hamin, and Stewart Scott

ACKNOWLEDGEMENTS

The authors would like to thank Trey Herr, Josephine Wolff, Rich Mogull, Marc Rogers, Scott Piper, Jim Higgins, and Frank Reyes for feedback on various versions of this document. The authors would also like to thank the numerous individuals who spoke with them on background for this report, as well as the participants, who shall remain anonymous, in three workshops on cloud security. Finally, the authors would like to thank Trey Herr, Simon Handler, and Lily Liu for their intellectual support on cloud issues.

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Cloud computing is so ubiquitous to modern digital and internet infrastructure that it often, perversely, eludes our notice. Cloud's benefits—cost savings, scalability, and outsourced management of infrastructure security and availability—have precipitated its rapid adoption. But, perhaps because the focus has been so strongly on these benefits, policy has lagged behind in reckoning with how essential cloud computing is to the functioning of the most critical systems and in the development of oversight structures commensurate with that new centrality.

The cloud, just like its on-premises predecessors, faces risks. In the Sunburst hack, the compromise of core cloud services—in this case, Microsoft Azure's Identity and Access Management services—was one vector that exposed multiple US government agencies to snooping by malicious actors.[1] The cloud, too, is vulnerable to the perennial fallibility of software systems: in a 2019 Google cloud outage, a misconfiguration cascaded into an hours-long brownout for services like YouTube and Snapchat, as Google's network became congested and the very network management tools needed to resolve the issue were themselves throttled.[2] The combination of the cloud's increasing role as fundamental infrastructure for many other services and its status as a complex, technical, system-of-systems imply a simple follow-on question: are the policy tools at hand suited to govern cloud's increasing complexity and criticality?

This report zeros in on an area where the stakes for cloud risk management are high: critical infrastructure (CI) sectors. The US government designates sectors as CI because their incapacity or destruction would have a "debilitating effect on security, national economic security, national public health or safety, or any combination of those matters."[3] The potential for a cloud compromise or outage to incapacitate, even temporarily, such a sector is one which policymakers must take seriously.

The research draws on public information to examine cloud adoption in five specific CI sectors: healthcare, transportation and logistics, energy, defense, and financial services. In doing so, it pays particular attention to three factors that can make the cloud an operational benefit or necessity for the ongoing functionality of a sector.

1. **Data storage and availability:** How much data, and of what kind, has a given critical sector put into the cloud? Could the sector maintain operations without access to this data? Are there on-premises data backups and/or regulatory requirements mandating them?[4]

2. **Scale and scalability:** Has a given sector come to rely on scale that only cloud computing can enable, or upon the cloud's ability to scale to larger workloads rapidly? Do the sector's core services now rely on such capacity?

3. **Continuous availability requirements:** Has a given sector permanently moved systems that require constant availability into the cloud without local backups? And, if there are back-ups, what delay do they require to resume function in the case of a cloud compromise or outage?

This report aims to raise awareness of the risks that a potential cloud compromise or outage poses to CI and, in so doing, to make the case that these risks necessitate the maturation of current policy tools, and creation of others, to address these risks. It does not seek to vilify cloud adoption by CI sectors or preach a return to on-premises data processing. Instead it suggests that CI sector regulators must consider cloud security and resilience a key question within their remit.

The report goes on to describe two features that make the risk profile of cloud computing markedly different from that of previous computing paradigms and that must inform the design of cloud risk management policy at a national level: **compounded dependence** and **delegated control and visibility**. Compounded dependence describes how widespread cloud adoption causes a

[1] Simon Handler, "Broken Trust: Lessons from Sunburst," *Atlantic Council* (blog), March 29, 2021, https://www.atlanticcouncil.org/in-depth-research-reports/report/broken-trust-lessons-from-sunburst/.

[2] Brian Barrett, "How a Google Cloud Catch-22 Broke the Internet," *WIRED*, June 17, 2019, https://www.wired.com/story/google-cloud-outage-catch-22/.

[3] 42 U.S.C. § 5195c,"Critical Infrastructures Protection," https://www.govinfo.gov/app/details/USCODE-2021-title42/USCODE-2021-title42-chap68-subchapIV-B-sec5195c.

[4] To the reader, term "on-premises data backup" refers to off the cloud and stored locally.

huge range of organizations to depend upon a few shared linchpin technology systems, including unglamorous subsystems within the cloud, where the failure of one node could precipitate a cascading collapse. Delegated control and visibility describe how organizations that adopt cloud services cede control of and lose visibility into the operations and failure modes of these technology systems, posing challenges for both businesses and policymakers seeking to measure and manage cloud risks.

The factors of compounded dependence and delegated control and visibility pose challenges to managing potential risks to the cloud with existing policy tools, which remain more focused on end products and services than their shared architecture and infrastructure. These risk factors will only become more pronounced as organizations accelerate their move to the cloud, and policy structures designed to manage them will be essential to smoothly navigating the ongoing transition towards cloud computing as the dominant computing paradigm.

The report concludes with policy recommendations to help policymakers gain more visibility into and eventually a better hold on cloud risks for CI sectors, building on the 2023 cloud security report from the US Department of the Treasury and the 2023 National Cybersecurity Strategy. These recommendations center on equipping Sector Risk Management Agencies (SRMAs)—the entities currently tasked with managing cybersecurity risk in CI sectors—with appropriate tools to understand cloud usage and risk within their sector, as well as mapping out the beginnings of a structure for cross-sector cloud risk management to facilitate greater transparency and oversight. These ideas are a start, rather than an end state, for cloud risk policy—visibility is a prerequisite for risk management, but other tools will be required in concert to fully confront the problem.

The conversation about cloud security is no longer just about the security of services, but about the durability of infrastructure underpinning fundamental economic and political activities. For policymakers, that recognition must now become as tangible as it is urgent.

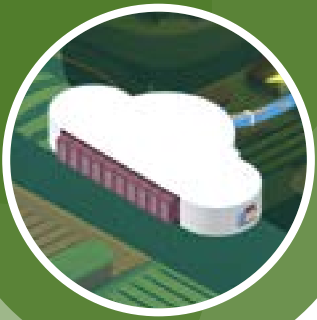

1

INTRODUCTION

A risk to the security or availability of cloud computing is a risk to US economic and national security. Over 95 percent of Fortune 500 companies use cloud systems,[5] and many sectors considered critical infrastructure (CI)—healthcare, transportation and logistics, energy, defense, and financial services, for example—are increasingly using cloud computing to support their core functionality. The government too is adopting cloud computing, with more and more critical governmental functions built in the cloud, from systems development at the US Department of Defense (DOD) to national public health crisis response systems.[6], [7]

The widespread and increasing use of the cloud, especially for high-value computing workloads, has also raised the stakes for cloud security. The cloud's centralization of data and computing capabilities has made it a target of, and battlefield for, creative, persistent threats engaging in economic espionage, offensive cyber operations, and even destructive attacks on civilian infrastructure, as well as a stage for arcane regulatory disputes and outmatched procurement processes. As a form of centralized infrastructure for computing, cloud deployments are exposed to both the security risks of their customers and the malintent of those customers' adversaries. Cloud service providers (CSPs) thus make architectural, operational, and security decisions with potentially vast, cascading effects across sectors. And still, they must build and operate this cloud infrastructure while straddling a highly contested global marketplace that crisscrosses political boundaries often fraught under the strain of immense technical complexity.

The aim of this paper is not, notably, to suggest that cloud adoption should be avoided or that cloud computing deployments innately bear more risk than their on-premises counterparts. Cloud computing offers real efficiency and cost benefits to organizations by obviating the need to maintain data centers and enabling flexible compute scaling in response to demand. Arguments can be made that on average cloud deployments are more secure than on-premises

[5] To the reader, as of 2018, 95 percent of Fortune 500 companies were already using Microsoft Azure in some capacity.
See: Arpan Shah, "Microsoft Azure: The Only Consistent, Comprehensive Hybrid Cloud," September 25, 2018,
https://azure.microsoft.com/en-us/blog/microsoft-azure-the-only-consistent-comprehensive-hybrid-cloud/.

[6] "For DOD, Software Modernization and Cloud Adoption Go Hand-in-Hand," *Federal News Network*, webinar announcement, September 26, 2022, https://federalnewsnetwork.com/cme-event/federal-insights/pushing-forward-on-dod-software-modernization/.

[7] GNC Staff, "HHS Protect: The Foundation of COVID Response," *Government Computer News* (GCN), November 16, 2020,
https://gcn.com/data-analytics/2020/11/hhs-protect-the-foundation-of-covid-response/315746/.

systems (though opinions are far from definitive).[8],[9] This paper instead seeks to illustrate that the risks posed by the widespread adoption of cloud computing are meaningfully *different* from the risks arising from the myriad of independent organization-specific computing systems. Moreover, it is a call to attention on the ways in which existing policy is not yet well-equipped to oversee and manage this novel risk landscape. In fact, it is precisely because cloud computing is so valuable that it is well worth attempting to grapple with these new risks rather than simply fleeing back to on-premises data systems.

From a policy perspective, one of the most challenging aspects of attempting to understand the cloud's role in CI is a lack of consistent visibility into the exact nature and depth of cloud adoption by individual organizations. There is great variety in how an organization might use cloud services—public, private, community, or hybrid clouds—as well as in the breadth of services and organization-specific usages of each service model, from software as a service (SaaS) to platform as a service (PaaS) and infrastructure as a service (IaaS). Different organizations have adopted the cloud at different rates and differ too in the degree to which they host operation-critical data and workloads in the cloud, versus merely using the cloud to host auxiliary data not necessary for their core operations.

Figure 1. Lego Blocks of Service Models

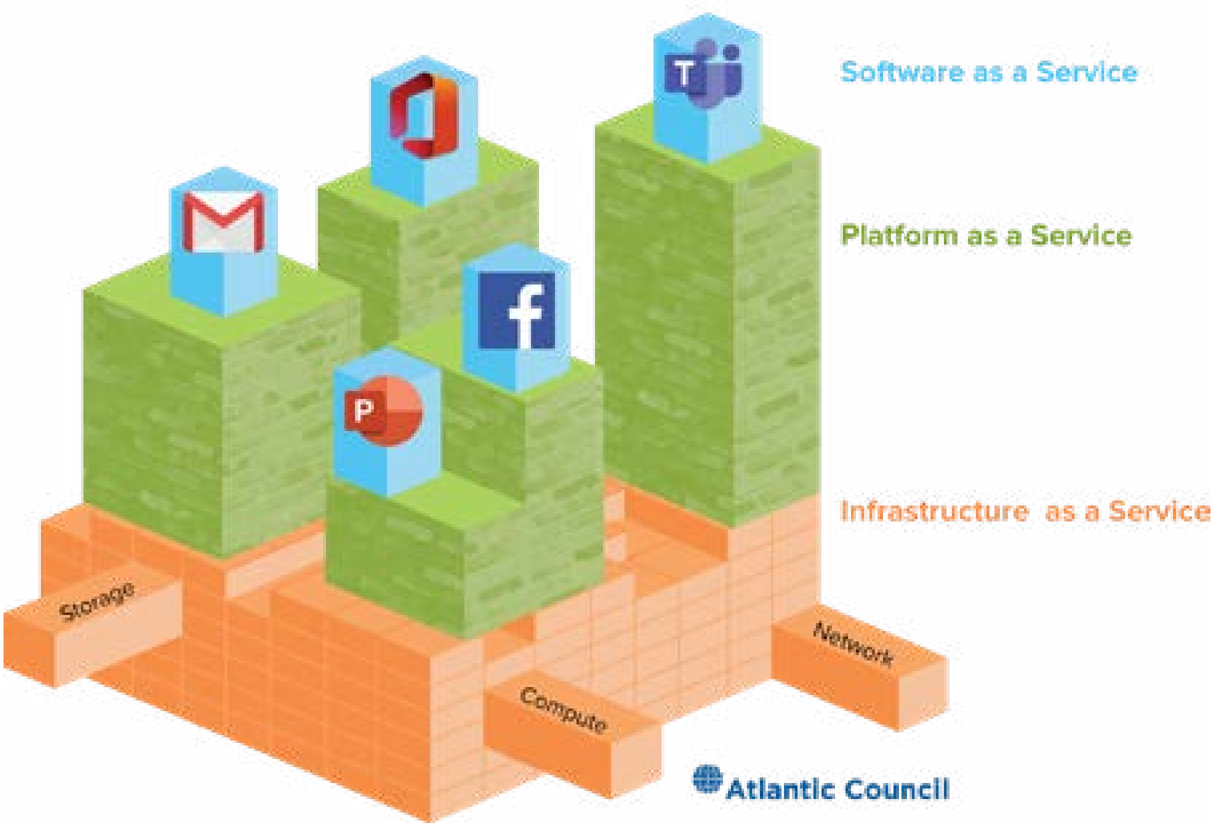

SOURCE: Lily Liu and Sarah Orio."

[8] Kevin Townsend, "More Than Half of Security Pros Say Risks Higher in Cloud Than On Premise," *SecurityWeek*, September 29, 2022, https://www.securityweek.com/more-half-security-pros-say-risks-higher-cloud-premise/.

[9] Dan Geer and Wade Baker, "Is the Cloud Less Secure than On-Prem?," Usenix (;login:), Fall 2019, https://www.usenix.org/system/files/login/articles/login_fall19_12_geer.pdf.

This report makes use of what information is public to show that the cloud is—already, and increasingly—embedded within CI sectors. It makes a series of policy recommendations intended primarily to help the government gain visibility into the complex, interdependent ecosystem of cloud risk and CI. These recommendations stop short of suggesting a holistic model for cloud regulation—the complexity of these products and how customers depend on different parts of them is still growing and too poorly understood across industry sectors for a one-size-fits-all approach.[10] But policymakers will not be able to arrive at a workable model without more visibility into and consideration of the landscape of cloud use and cloud risk.

The recent US National Cybersecurity Strategy invokes cloud services and calls for policy to "shift the burden of responsibility" for better cybersecurity.[11] This is an important and timely debate that must involve the largest CSPs, who shoulder so much risk and must be central in any renewed effort to govern the security of cloud services and infrastructure. One of the themes of this report is that policymakers will need to "shift the burden of transparency" onto CSPs, who are the only entities well-positioned to provide insight into their dependencies and the risks they face. Now, before a catastrophic incident, is the time for the government, with industry, to accelerate toward a healthier and more risk-informed regulatory model for cloud computing.

Cloud Computing and Cloud Risks

In cloud computing, CSPs offer customers the ability to connect, over the Internet, to data center servers and other computing resources for on-demand data storage, specialized services, and big data processing.[12] Cloud computing gives organizations, from non-profits and government agencies to the Fortune 500, the ability to run applications and work with data without building and operating their own physical data centers or buying computer hardware, as well as the flexibility to increase or decrease the amount of these services they pay for based on real-time needs.[13] Cloud computing allows organizations to offload many of the challenges that come with ensuring the security and ongoing operations of data infrastructure to a (nominally) well-resourced, technically mature provider.

The average internet user depends on cloud computing to edit documents on Google Drive, talk over a Zoom video call, or access their favorite retail or social media websites. Companies and organizations accelerated their adoption of enterprise cloud computing during the COVID-19 pandemic, as employees could use the Internet to interact with cloud-hosted organizational resources, regardless of their physical location. Cloud computing is an increasingly dominant component of the entire information technology (IT) ecosystem, even if its presence is functionally invisible to most end users.

Three CSPs dominate the market: Amazon Web Services (AWS), Microsoft Azure, and Google Cloud. Together, they make up over 65 percent of the global cloud market. These "hyperscalers" benefit from economies of scale: each operates hundreds of massive data centers and allocates and redistributes customer compute demands across them to use computing resources optimally. These economies of scale provide richer suites of functionality—and more robust security features—than many small customers could ever build in-house. But, the reliance of vast numbers of online services on cloud computing, combined with the highly concentrated nature of CSPs, mean that an outage at a single CSP can have strange and cascading effects across a wide range of applications. For example, a single AWS outage in 2021 stopped operations at Amazon delivery warehouses, blocked access to online education testing services, and immobilized smart home devices such as robot vacuums and app-controlled automatic cat feeders.[14]

The increasing reliance of many organizations on a few CSPs creates increasingly concentrated forms of systemic risk. When a CSP's distributed computing functions fail, the interruptions to service availability can cascade across the firm's services and clients. Most CSPs build their infrastructure from common, modular

[10] Scott Piper and Amitai Cohen, "The State of the Cloud 2023," WIZ (blog)," February 6, 2023, https://www.wiz.io/blog/the-top-cloud-security-threats-to-be-aware-of-in-2023.

[11] President Biden, "National Cybersecurity Strategy," The White House, March 1, 2023, https://www.whitehouse.gov/wp-content/uploads/2023/03/National-Cybersecurity-Strategy-2023.pdf.

[12] Simon Handler, "Dude, Where's My Cloud? A Guide for Wonks and Users," *Atlantic Council* (blog), September 28, 2020, https://www.atlanticcouncil.org/in-depth-research-reports/report/dude-wheres-my-cloud-a-guide-for-wonks-and-users/.

[13] Peter Mell and Timothy Grance, "The NIST Definition of Cloud Computing," NIST Special Publication (SP) 800-145, National Institute of Standards and Technology (NIST), US Department of Commerce, September 2011, https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-145.pdf.

[14] Annie Palmer, "Dead Roombas, Stranded Packages and Delayed Exams: How the AWS Outage Wreaked Havoc across the U.S.," CNBC, December 9, 2021, https://www.cnbc.com/2021/12/09/how-the-aws-outage-wreaked-havoc-across-the-us.html.

Figure 2: Cloud Computing Market Share in 2020

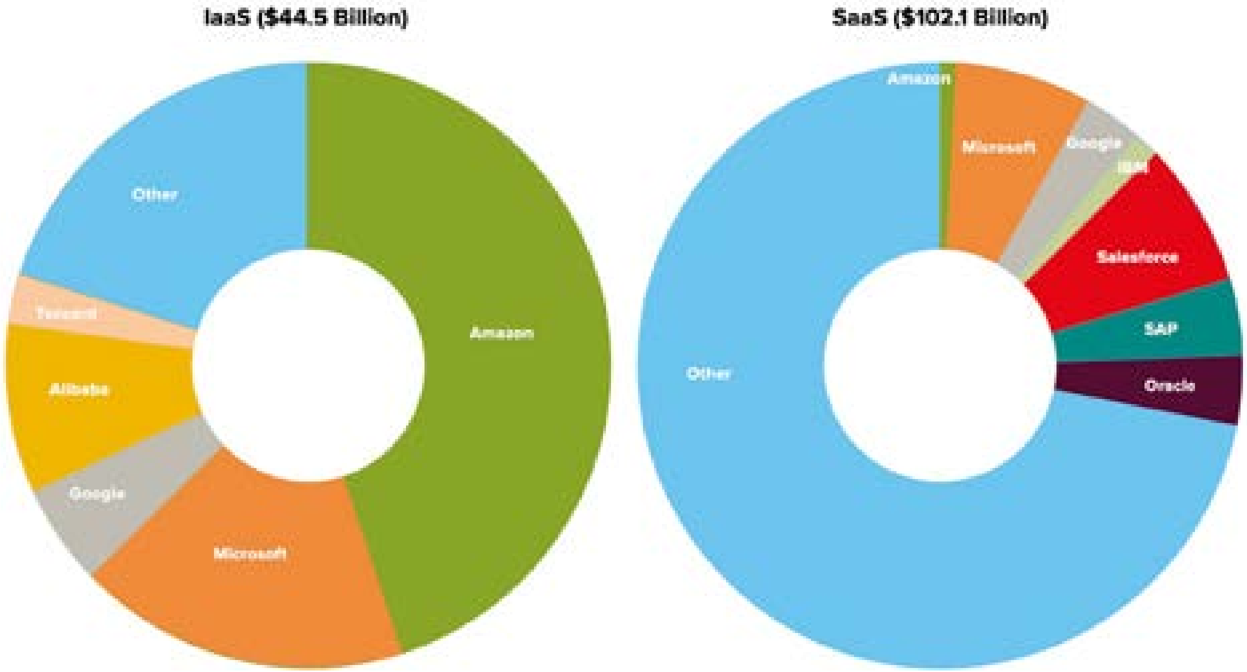

SOURCE: Lily Liu, 2020

architectures.[15],[16] The basic hardware and software packages that make up the cloud often share significant similarities—or are wholly identical—meaning one flaw can be present in many diverse locations. For instance, on December 14, 2020, Google Cloud suffered a widespread outage that made Gmail, Google Drive, YouTube, and many other Google services inaccessible globally for about 45 minutes.[17] The outage occurred because of an error in allocating storage resources for a User ID Service, which authenticates users before they can interact with Google services.[18] Because many different Google products rely on the same authentication service, this error brought down several major cloud applications worldwide. Even Google Nest smart home devices, like speakers and thermostats, also forcibly went offline, triggering a fail-safe mode where users could not access the device settings.[19] The number of systems and actors affected in this outage, stemming from a singular source, demonstrates the interconnectedness of cloud computing infrastructure.

Cloud incidents do not result only from internal software failures. Physical incidents such as a flood or storm can take out a cloud data center. For example, a severe storm near Washington, DC, cut power leading to an AWS data center in 2012, resulting in multi-hour downtimes of sites

[15] James Hamilton, "Architecture for Modular Data Centers," Cornell University, arXiv platform, December 21, 2006, https://doi.org/10.48550/arXiv.cs/0612110.

[16] "Seven Principles of Cloud-Native Architecture," Alibaba Cloud Community, January 5, 2022, https://www.alibabacloud.com/blog/seven-principles-of-cloud-native-architecture_598431.

[17] Alex Hern, "Google Suffers Global Outage with Gmail, YouTube and Majority of Services Affected," *The Guardian*, December 14, 2020, https://www.theguardian.com/technology/2020/dec/14/google-suffers-worldwide-outage-with-gmail-youtube-and-other-services-down.

[18] "Google Cloud Infrastructure Components Incident #20013," Google Cloud, accessed April 27, 2023, https://status.cloud.google.com/incident/zall/20013.

[19] Hern, "Google Suffers Global Outage."

such as Instagram and Heroku.[20] The failure of Heroku, itself a hosting service provider, caused further cascading failures for the websites built on its platform.[21] While many CSPs have protocols for rerouting traffic to other data centers in the event of an outage, called "failover," this event demonstrated that failing over safely and seamlessly is a significant challenge: the outage impacted Netflix even though it explicitly pays to spread traffic across multiple data centers to avoid just such a failure—because, per Netflix's Director of Architecture, AWS traffic routing "was broken across all zones" during the incident.[22] Additionally, because cloud services rely on the internet to connect customers to computing resources, attacks against underlying internet infrastructure can have ripple effects impacting the availability of cloud services. For example, a 2016 distributed denial of service (DDoS) attack against Dyn, a domain name service provider, caused outages at cloud services such as AWS.[23]

Cloud platforms can also suffer from unique cybersecurity risks. Because cloud services are generally multitenant environments—that is, a single instance of the software or infrastructure serves multiple unrelated organizations at the same time—a malicious actor who can escape the bounds of tenant isolation can access the data and resources of other customers. Security researchers who identified multiple bugs that them to access other tenants' data called out a "problematic pattern" in which CSPs are often non-standardized and non-transparent about their tenant-isolation practices, making risk management more challenging for customers.[24] And, because a single CSP typically serves many customers, cloud platforms are appealing targets for hackers seeking to compromise many different organizations, including through software supply chain attacks. These attacks involve implanting and/or exploiting vulnerabilities in a less secure software that the target software depends on. During the Solar Winds/Sunburst campaign, discovered in 2020, Russian groups managed to access Microsoft Azure's identity and access management (IAM) service, Azure AD. Critical applications, such as Office 365, Workday, AWS Single Sign-On, and Salesforce, are commonly integrated with Azure AD. As one of several techniques, the actors abused this access to move through different Office 365 user accounts to access highly confidential documents, emails, and calendars.[25]

Cloud computing risks have different characteristics than those of on-premises computing. Increasing adoption of cloud services by CI operators, therefore, necessitates more active involvement from the policy community to adapt to this new risk landscape. Accordingly, this report brings attention to five CI sectors that are in the process of forming deep dependencies on the cloud and provides guidance for how policymakers can create visibility into the cloud ecosystem to begin adapting policy to address the compounded dependence and delegated control and visibility prevalent within cloud infrastructure.

[20] Rich Miller, "Amazon Data Center Loses Power during Storm," Data Center Knowledge (Informa), June 30, 2012, https://www.datacenterknowledge.com/archives/2012/06/30/amazon-data-center-loses-power-during-storm.

[21] Nati Shalom, "Lessons from The Heroku Amazon Outage," Cloudify, June 18, 2012, https://cloudify.co/blog/lessons-from-heroku-amazon-outage/.

[22] Miller, "Amazon Data Center Loses Power during Storm."

[23] Sebastian Moss, "Major DDoS Attack on Dyn Disrupts AWS, Twitter, Spotify and More," DCD Media Center, October 21, 2016, https://www.datacenterdynamics.com/en/news/major-ddos-attack-on-dyn-disrupts-aws-twitter-spotify-and-more/.

[24] Amitai Cohen, "Introducing PEACH, a Tenant Isolation Framework for Cloud Applications," WIZ (blog), December 14, 2022, https://www.wiz.io/blog/introducing-peach-a-tenant-isolation-framework-for-cloud-applications.

[25] Handler, "Broken Trust."

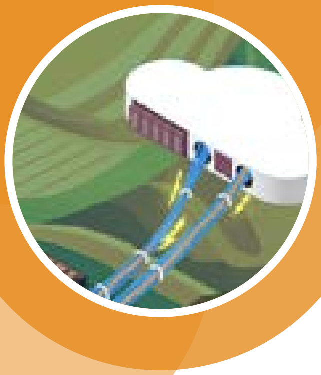

2

CRITICAL SECTORS USING THE CLOUD

This section examines cloud computing's proliferation across five critical sectors: healthcare, transportation and logistics, energy, defense, and financial services. It highlights how cloud computing already supports the maintenance of everything from patient data to home energy supplies. The sensitivity of data and services stored in the cloud varies among these sectors, yet, within each, the cloud is already or soon will be critical to US economic, national security, and general societal interests.

HEALTHCARE SECTOR

The healthcare sector has quickly recognized cloud computing's benefits.[26] One industry survey, for instance, reported that 35 percent of healthcare organization respondents already store more than half their data and infrastructure in the cloud.[27] In 2020, companies spent $28.1 billion on healthcare cloud computing, with the number projected to increase to $64.7 billion by 2025.[28]

The healthcare sector generates enormous amounts of sensitive data, much of which it stores in the cloud. Electronic health records (EHRs), which contain data such as a patient's medical history, diagnoses, and medications, are increasingly common in healthcare for their efficiency and interoperability, as are medical sensors and monitors that generate large amounts of data.

[26] US Cybersecurity and Infrastructure Security Agency (CISA), "Healthcare and Public Health Sector," CISA, accessed April 27, 2023, https://www.cisa.gov/healthcare-and-public-health-sector; Vinati Kamani, "5 Ways Cloud Computing Is Impacting Healthcare," Health IT Outcomes, October 2, 2019, https://www.healthitoutcomes.com/doc/ways-cloud-computing-is-impacting-healthcare-0001.

[27] Jessica Kim Cohen, "Report: Healthcare Industry Leads in Cloud Adoption," Health IT: Becker's Hospital Review, April 9, 2018, https://www.beckershospitalreview.com/healthcare-information-technology/report-healthcare-industry-leads-in-cloud-adoption.html.

[28] Research and Markets, "Global Healthcare Cloud Computing Market (2020 to 2025) – Emergence of the Telecloud Presents Opportunities," GlobeNewsWire, October 2, 2020, https://www.globenewswire.com/news-release/2020/10/02/2102876/0/en/Global-Healthcare-Cloud-Computing-Market-2020-to-2025-Emergence-of-the-Telecloud-Presents-Opportunities.html.

CSPs have created healthcare-specific tools for the cloud storage of EHR data, such as Microsoft Azure's "Fast Healthcare Interoperability Resources" (FHIR), a standard for transmitting EHRs and protected health information (PHI),[29] which has since integrated into AWS and Google Cloud.[30] A 2015 survey of small healthcare providers found that 82 percent of respondents in urban areas used a cloud-based EHR system[31] (often because they were cheaper than on-premises systems). Moreover, these findings show major existing EHR software providers have begun making deals to move client EHR systems to the cloud or even acquire them wholesale by CSPs, such as Oracle's recent acquisition of Cerner for $28.3 billion for its Millennium EHR platform to build a US national cloud database of EHRs.[32]

Other healthcare-adjacent systems—like insurance systems, Health Insurance Portability and Accountability Act (HIPAA) compliant communications, laboratories and testing labs, crisis coordination networks, and supply-chain management practitioners—have also largely transitioned to the cloud. Healthcare.gov, the US government's health insurance enrollment site, completely runs on AWS.[33] HIPAA-compliant email solutions used in healthcare settings are typically extensions of cloud-based emailing systems such as Outlook and Gmail. Radiology facilities have moved to cloud computing to share images and reduce storage costs.[34] Cloud computing helps healthcare providers translate great volumes of clinical information into "clinical decision support," which was previously impossible due to the limitations of on-premises computing infrastructure.[35] GE Healthcare Technologies uses Microsoft Azure's Edison Datalogue Connect to provide secure image and data exchange to physicians, reducing the need to duplicate tests at different facilities.[36] Technologies such as Amazon's Comprehend Medical standardize proprietary records so healthcare providers need not decrypt external patient records.[37]

The US Centers for Disease Control and Prevention (CDC) tapped AWS in 2014 to bolster its BioSense 2.0 program, an initiative to provide timely insight into the public health of US communities.[38] BioSense links local, state, and federal public health institutions to respond to public-health crises faster, which requires significant computing power and storage.[39] In the United Kingdom, the National Health Service signed a deal with IBM for a secure public cloud to improve service delivery.[40] Novartis, one of the largest healthcare companies in Europe, uses cloud services to improve data analytics and manage decisions about a complex global supply

[29] Gregory J. Moore, "Reimagining Healthcare: Partnering for a Better Future," Microsoft, December 2, 2019, https://cloudblogs.microsoft.com/industry-blog/health/2019/12/02/reimagining-healthcare-partnering-for-a-better-future/.

[30] Henner Dierks and Angus McAllister, "Using Open Source FHIR APIs with FHIR Works on AWS," Amazon Web Services (blog), August 28, 2020, https://aws.amazon.com/blogs/opensource/using-open-source-fhir-apis-with-fhir-works-on-aws/; "Cloud Health API: FHIR," Google Cloud, accessed April 27, 2023, https://cloud.google.com/healthcare-api/docs/concepts/fhir.

[31] John DeGaspari, "Cloud-Based EHRs Popularity Grows among Small Practices," Fierce Healthcare," June 3, 2015, https://www.fiercehealthcare.com/ehr/cloud-based-ehrs-popularity-grows-among-small-practices.

[32] Heather Landi, "Oracle, Cerner Plan to Build National Medical Records Database as Larry Ellison Pitches Bold Vision for Healthcare," Fierce Healthcare, June 10, 2022, https://www.fiercehealthcare.com/health-tech/oracle-cerner-plan-build-national-medical-records-database-ellison-pitches-bold-vision; Heather Landi, "Google, Epic Ink Deal to Migrate EHRs to the Cloud," Fierce Healthcare, November 16, 2022, https://www.fiercehealthcare.com/health-tech/google-epic-ink-deal-migrate-hospital-ehrs-cloud-ramp-use-ai-analytics; "A Cloud-Infrastructure Platform: Reimagining Remote Hosting Services," Veradigm (formerly Allscripts), accessed April 27, 2023, https://www.allscripts.com/service/allscripts-cloud/; "The CareVue EHR," Medsphere, accessed April 27, 2023, https://www.medsphere.com/resources/carevue-ehr-overview/.

[33] "Managing the Healthcare.gov Cloud Migration," Booz Allen Hamilton (case study), February 19, 2021, https://www.boozallen.com/s/insight/thought-leadership/managing-the-healthcare-gov-cloud-migration.html.

[34] Amy Vreeland et al., "Considerations for Exchanging and Sharing Medical Images for Improved Collaboration and Patient Care," HIMSS-SIIM Collaborative White Paper, *Journal of Digital Imaging* 5 (October 2016) 547–58, https://pubmed.ncbi.nlm.nih.gov/27351992/.

[35] Jennifer Bresnick, "Can Cloud Big Data Analytics Fix Healthcare's Insight Problem?" Health IT Analytics, December 1, 2015, https://healthitanalytics.com/news/can-cloud-big-data-analytics-fix-healthcares-insight-problem.

[36] Moore, "Reimagining Healthcare."

[37] "Healthcare Interoperability: Creating a Clearer View of Patients," Amazon Web Services, 2019, https://d1.awsstatic.com/Industries/HCLS/Resources/Healthcare%20Data%20Interoperability%20AWS%20Whitepaper.pdf.

[38] "US Centers for Disease Control and Prevention (CDC) Case Study," Amazon Web Services, 2014, https://aws.amazon.com/solutions/case-studies/us-centers-for-disease-control-and-prevention/.

[39] Kelley G. Chester, "BioSense 2.0," *Online Journal of Public Health Informatics* 5, no. 1 (April 4, 2013): e100, https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3692855/.

[40] Tammy Lovell, "IBM Deal to Provide the NHS with Quicker to Access Cloud Services," *Healthcare IT News*, August 3, 2020, https://www.healthcareitnews.com/news/emea/ibm-deal-provide-nhs-quicker-access-cloud-services.

chain of medicine manufacturing and distribution, among other functions.[41]

While it is challenging to determine how catastrophic the impacts of an outage of any single service or CSP would be, the cloud is increasingly critical to the efficient function of many healthcare organizations. While some practices may be able to revert to pen and paper in the event of a cloud outage, others may not, and most will suffer from the change.[42] In one incident, a ransomware attack on Allscripts' cloud-based EHR system forced healthcare providers to fall back to paper prescriptions, possibly delaying life-saving care and raising the risk of fraud and abuse.[43]

The cloud offers real benefits, especially for small providers: cost savings, ease of standing-up functionality without an in-house IT team, and (potentially) increased security over on-premises deployments. There is a reason why the various federal cloud strategies and policies, as well as the new National Cybersecurity Strategy, emphasize encouraging cloud adoption: Adoption must match more fulsome, fine-grained, and effective scrutiny of CSPs and their infrastructure. Healthcare's cloud transition will continue, so examining potential outage impacts and the degree of systemic vulnerability to a few points of failure are urgent priorities.

TRANSPORTATION AND LOGISTICS SECTOR

The transportation and logistics sector plays a vital role in US and global supply chains. For instance, the freight shipping industry moves some $19 trillion of goods over land in the United States each year.[44] The European Union similarly houses the world's largest ocean shipping fleet and controls around 40 percent of the world's tonnage, moving everything from oil and gas to cars and electrical appliances.[45] This report focuses on cloud adoptions by logistics companies and airlines and finds that, at present, this sector tends to use the cloud to enhance and optimize existing business functions, with a select number of firms making monumental shifts.

Several transportation and logistics firms are transferring their data to the cloud for easier management and future needs, though these are not generally "whole-of-business" shifts. The United Parcel Service (UPS) started its cloud transition in 2019 with Google Cloud, and it recently inked a deal to expand its cloud data storage.[46] UPS uses cloud services to "see and control how packages move through [its] network,"[47] which it cites as particularly critical functionality for high-volume periods such as the holidays and the COVID-19 pandemic, during which UPS needed to deliver more than a billion vaccine doses.[48] FedEx uses the cloud to plan its pickup and delivery

[41] Amazon Web Services, "AWS Announces Strategic Collaboration with Novartis to Accelerate Digital Transformation of Its Business Operations," Business Wire, December 4, 2019, https://www.businesswire.com/news/home/20191204005238/en/AWS-Announces-Strategic-Collaboration-Novartis-Accelerate-Digital/.

[42] "Emergency Preparedness: Be Ready for Unanticipated Electronic Health Record (EHR) Downtime," Institute For Safe Medication Practices," August 24, 2022, https://www.ismp.org/resources/emergency-preparedness-be-ready-unanticipated-electronic-health-record-ehr-downtime.

[43] Evan Sweeney, "Physician Practices Forced to Use Paper Records Lash out at Allscripts over Ransomware Response," Fierce Healthcare," January 22, 2018, https://www.fiercehealthcare.com/ehr/allscripts-ransomware-physician-practices-ehr-cybersecurity-e-prescribing.

[44] Bureau of Transportation Statistics, "2018 Freight Flow Estimates," US Department of Transportation, December 19, 2019, https://www.bts.gov/newsroom/2018-freight-flow-estimates.

[45] "Climate Change and Shipping ECSA Position Paper," European Community Shipowners' Associations (ECSA), January 2008, https://www.ecsa.eu/sites/default/files/publications/083.pdf.

[46] Isabelle Bousquette, "UPS Expands Deal with Google Cloud to Prepare for Surge in Data," *Wall Street Journal*, March 29, 2022, https://www.wsj.com/articles/ups-expands-deal-with-google-cloud-to-prepare-for-surge-in-data-11648551600.

[47] "UPS Extends Use of Google Cloud Data Analytics Technology," United Parcel Service, March 25, 2022, https://about.ups.com/us/en/our-stories/innovation-driven/ups-and-google-cloud.

[48] United Parcel Service, "UPS Extends Use of Google Cloud."

routes worldwide.[49] While these functions are important, without more information it is challenging to predict whether a significant cloud outage or compromise would be catastrophic or merely burdensome and inefficient for the ongoing operation of these entities.

American Airlines works with Microsoft Azure to host all its data and many software tools.[50] The airline uses cloud services to manage aircraft operations, like airport taxiing decisions, flight planning, and gating decisions at Dallas Fort Worth, one of its main hubs, as well as to run mobile apps and airport kiosks.[51] Other partnerships hint at future plans to utilize the cloud for more safety-critical tasks: NASA announced a partnership with General Electric Company on a project to integrate cloud computing into air traffic management systems,[52] and the Federal Aviation Administration (FAA) entered into a partnership in 2020 to work on cloud modernization for its safety systems.[53] A short-term cloud outage might cause catastrophic outcomes in any of these functions, but could cause delays or stoppages that trigger subsequent effects in downstream systems reliant on the smooth functioning of air transportation.

The transportation and logistics sector must contend with seasonal swings and weather emergencies, where situations call for additional computing power to solve challenging optimization problems on the fly. For example, Rolls-Royce's (RR) engine maintenance program downloads terabytes of data from airline fleets globally. RR relies on cloud computing to store and analyze this quantity of data, and the level of data ingestion and exfiltration is volatile, subject to fluctuations in global travel demand.[54] American Airlines depends on the cloud's elasticity to quickly rebook passengers during massive flight disruptions using services hosted by International Business Machines (IBM).[55]

New technology initiatives also demonstrate the sector's reliance on elasticity. UPS has attached radio-frequency identification chips (RFID) to packages for efficiency and optimization, which will increase data storage and processing demands significantly. The United States Postal Service's (USPS) rollout of machine learning tools depends on the capture of terabytes of package data from its processing centers, necessitating elastic data storage.[56] A report by international courier DHL states that more than 50 percent of logistics providers currently use cloud-based services, and an additional 20 percent will adopt it in the near future.

In sum, the transportation and logistics sector appears to currently use the cloud more for planning systems than real-time operational decisions, where failure could have devastating effects. However, even short-lived delays in shipping and transportation can have costly economic effects. Moreover, industry projections and cloud-feature development suggest that the cloud will become more critical to the sector's safe functioning in the future.

[49] "FedEx Uses Java on Azure to Modernize Route Planning for Pickup and Delivery Operations," Microsoft for Java Developers (YouTube video), 2022, https://www.youtube.com/watch?v=fJ_OUNdFXHs.

[50] Tobias Mann, "American Airlines Decides to Cruise into Azure's Cloud," *The Register*, May 19, 2022, https://www.theregister.com/2022/05/19/american_airlines_azure/.

[51] "American Airlines and Microsoft Partnership Takes Flight to Create a Smoother Travel Experience for Customers and Better Technology Tools for Team Members," American Airlines Newsroom, May 18, 2022, https://news.aa.com/news/news-details/2022/American-Airlines-and-Microsoft-Partnership-Takes-Flight-to-Create-a-Smoother-Travel-Experience-for-Customers-and-Better-Technology-Tools-for-Team-Members-MKG-OTH-05/default.aspx.

[52] "Air Traffic Management Set to Meet Cloud Technology," CloudTweaks, September 28, 2012, https://cloudtweaks.com/2012/09/air-traffic-management-technology/

[53] "FAA Selects Leidos to Modernize Safety System," Leidos (news release), December 15, 2020, https://www.leidos.com/insights/faa-selects-leidos-modernize-safety-system.

[54] Susanna Ray, "From Airplane Engines to Streetlights, Transportation Is Becoming ore Intelligent," Microsoft, May 2, 2016, https://news.microsoft.com/transform/from-airplane-engines-to-street-lights-transportation-is-becoming-more-intelligent/.

[55] "American Airlines: The Route to Customer Experience Transformation Is through the Cloud," IBM (case study), March 2018, https://www.ibm.com/case-studies/american-airlines.

[56] Jory Heckman, "USPS Gets Ahead of Missing Packages with AI Edge Computing," *Federal News Network*, May 7, 2021, https://federalnewsnetwork.com/artificial-intelligence/2021/05/usps-rolls-out-edge-ai-tools-at-195-sites-to-track-down-missing-packages-faster/.

ENERGY SECTOR

The energy sector, as Presidential Policy Directive 21 (PPD-21) puts it, is "uniquely critical because it provides an 'enabling function' across all CI sectors."[57] Increasingly, energy has moved away from manual systems to automated ones reliant on the cloud for managing and making use of data.[58] The energy sector looks to the cloud to update aging interfaces and increase data-transmission efficiency.[59]

Smart grids are an example of critical energy-related infrastructure partially or wholly reliant on the cloud. Smart grids increase the resilience and capacity of the grid through activities such as dynamic load balancing and additional visibility into grid operations. The US government continues to fund smart-grid development activities, with up to $3 billion for the task included in the 2021 infrastructure law.[60] Smart grids often rely on the cloud for part or all of their functionality,[61] meaning that, as smart grid projects get underway in more cities, more and more Americans will implicitly rely on the cloud to keep the lights on and to provide power to other CI such as hospitals, financial systems, and, ironically, datacenters hosting some of these same cloud services.

The cloud today appears to already host certain functions critical to energy delivery. Duke Energy, a major US provider, has contracted with IBM to operate its Gas Transportation Management System (GTMS) on cloud infrastructure.[62] Duke Energy provides natural gas distribution to approximately half a million customers in Ohio and Kentucky, and the GTMS is essential for this distribution network's safety and efficiency. Southern Company, the second-largest US power provider, uses Microsoft Azure to analyze real-time data from its energy equipment—and more critically—relies on the cloud's scalability to handle the influx of messages and alerts during storm situations to better marshal its repair crews.[63] Portland General Electric (PGE) serves nine-hundred thousand customers in Oregon and recently transitioned to a hybrid cloud service to store documents and data, as well as run software that assists with energy-loss detection, data analytics, and object storage.[64] Like PGE, General Electric's (GE) Renewable Energy division uses cloud services to analyze performance and maintenance information on a global network of wind turbines.[65] This data will eventually inform machine-learning and artificial intelligence (AI) applications, and but does not currently seem critical to the wind turbines' day-to-day functioning. Southern California Edison, one of the largest utility providers in the United States, uses cloud services to aggregate drone data for fighting wildfires.[66] Again, the functionality appears useful but not critical to keeping the grid running. This is a common trend across sectors—clear examples of increased cloud usage, but with ambiguous degrees of criticality for providing core services.

ENGIE, one of the largest power utilities in France, moved to cloud-based storage to improve business and power-delivery efficiency, forming a company-wide data storage system using cloud services to ingest and

[57] President Barack Obama, "Presidential Policy Directive 21 – Critical Infrastructure Security and Resilience," The White House, February 12, 2013, https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil.

[58] Vince Dawkins, "How the Energy Industry Is Embracing Cloud Computing: Three Key Areas of Success," Cloud Tech, August 7, 2019, https://cloudcomputing-news.net/news/2019/aug/07/how-energy-industry-embracing-cloud-computing-three-key-areas-success/.

[59] Dawkins, "How the Energy Industry Is Embracing Cloud."

[60] Grid Deployment Office, "Smart Grid Grants," US Department of Energy, accessed April 27, 2023, https://www.energy.gov/gdo/smart-grid-grants.

[61] Enterprise.nxt, "How Edge-to-Cloud Computing Powers Smart Grids and Smart Cities," Hewlett Packard Enterprise (HPE), April 12, 2022, https://www.hpe.com/us/en/insights/articles/how-edge-to-cloud-computing-powers-smart-grids-and-smart-cities-2204.html.

[62] "Duke Energy: Keeping Energy Flowing to Hundreds of Thousands of Customers with IBM and Oracle," IBM (case study), January 2015, https://www.ibm.com/case-studies/duke-energy.

[63] "Southern Company," Microsoft, March 19, 2018, https://customers.microsoft.com/en-us/story/southern-company-power-utilities-azure.

[64] "PGE Migrates to AWS, Significantly Improves Energy Loss Detection Performance," Amazon Web Services, 2021, https://aws.amazon.com/solutions/case-studies/portland-general-electric/.

[65] Don McDonnell and Scot Wlodarczak, "AWS Is How: GE Renewable Energy Increases Wind Energy Production," Amazon Web Services, June 21, 2021, https://aws.amazon.com/blogs/industries/aws-is-how-ge-renewable-energy-increases-wind-energy-production/.

[66] Katherine Noyes, "Fighting Fire With Tech at Southern California Edison," *Wall Street Journal*, April 20, 2021, https://deloitte.wsj.com/cio/2021/04/20/fighting-fire-with-tech-at-southern-california-edison/.

store energy consumption data and inputs from a range of small physical sectors.[67] While these functions factor into ENGIE's day-to-day activities, the data stored on the cloud does not appear critical for the actual power delivery. PGE and ENGIE depend on the cloud's elasticity to manage crises. During unprecedented wildfires and catastrophic wind events, PGE relied on the cloud's ability to scale, helping the company mitigate widespread power outages.[68] During a mass service disruption, PGE's communications channels remained online thanks to its cloud infrastructure. ENGIE's growing wind-turbine data collection relies on ever-increasing cloud storage, though, unlike the elasticity use-case for PGE, such data may not be necessary for core operations. The Electric Power Research Institute (EPRI) reported in 2020 that out of the twenty-two US utilities surveyed, half said that they expect to adopt "cloud-hosted transmission and distribution planning applications" within the next five years.[69]

The cloud plays a role not only in energy delivery but also in upstream processes like oil and gas extraction. Some of the largest US and European Union (EU) oil-and-gas firms use the cloud for data storage and processing. ExxonMobil, the largest US publicly traded oil-and-gas company, uses Microsoft Azure to collect and store sensor data from their Permian Basin extraction.[70] Total Energies, a French multinational oil company and the world's fifth largest, works with Nutanix to host and secure large databases in cloud services provided by SAP HANA and Oracle and has said it will move all IT functionality to the cloud eventually.[71] British Petroleum (BP) has gone "all in" on cloud data storage and availability by shutting down two of its largest on-premises data centers in London's Canary Wharf in favor of a package of Amazon services as well as a SAP product to host the oil company's AVEVA Unified Supply Chain decision-making software.[72] Marathon Oil centralized much of its data collection onto the cloud,[73] while Chevron developed a cloud-based tool for oil-well data management.[74] These companies store immense qualities of data in the cloud, and by increasingly shuttering on-premises data centers, will have to spend even more to turn back. Depending on the specific functions for which each relies on the cloud, a cloud outage could have impacts ranging from simply delaying data reporting to shutting down operational facilities or stunning supply chains.

As in healthcare, major energy players—from oil-and-gas companies to electricity-delivery utilities—are adopting the cloud for functions ranging from auxiliary data processing to core operational capabilities. The impacts of potential cloud compromises on energy availability are hard to predict, especially, as the interconnected nature of the energy supply chain and grid could magnify the unavailability of one component or system into widespread cascading effects. While policymakers have begun to grapple with the interconnection of cyber and energy—for example, the recent National Cybersecurity Strategy notes that cybersecurity will grow increasingly important for next-generation energy technologies such as "advanced cloud-based grid management platforms," and pledges to "build in cybersecurity proactively through implementation of the Congressionally-directed National Cyber-Informed Engineering Strategy"[75]—more work is required to fully map out the energy sector's cloud dependence as well as the potential impacts of a devastating cloud compromise for the sector.

[67] "ENGIE Builds the Common Data Hub on AWS, Accelerates Zero-Carbon Transition," Amazon Web Services, 2021, https://aws.amazon.com/solutions/case-studies/engie-aws-analytics-case-study/.

[68] "PGE Migrates to AWS, Significantly Improves Energy Loss Detection Performance," Amazon Web Services, 2021, https://aws.amazon.com/solutions/case-studies/portland-general-electric/.

[69] Michael Matz, "The Grid is Moving to the Cloud," *EPRI Journal*, May 24, 2021, https://eprijournal.com/the-grid-is-moving-to-the-cloud/.

[70] Reuters Staff, "Exxon, Microsoft Strike Cloud Computing Agreement for U.S. Shale," *Reuters*, February 22, 2019, https://www.reuters.com/article/ctech-us-exxon-mobil-microsoft-cloud-idCAKCN1QB1N8-OCATC.

[71] "Total powers Digital Transformation across Energy Production with Nutanix," Nutanix (case study), accessed April 27, 2023, https://www.nutanix.com/company/customers/total.

[72] Computer Business Review (CBR) Staff, "'You've Got to Have Courage!' BP on Going 'All-In' on the Cloud," TechMonitor, June 29, 2022, https://techmonitor.ai/technology/cloud/bp-cloud-migration-interview;
"BP Goes All-in on AWS for Its European Mega Data Centers," British Petroleum (news release), December 4, 2019, https://www.bp.com/en/global/corporate/news-and-insights/press-releases/bp-goes-all-in-on-aws-for-its-european-mega-data-centers.html.

[73] "Helping Marathon Oil Create a Next-Generation Cloud Native Data Platform," EPAM Systems (case study), accessed April 27, 2023, https://www.epam.com/services/client-work/helping-marathon-oil-create-a-next-gen-cloud-native-data-platform;
"Marathon Oil Reduces Intelligent Alert Creation Time from Months to Hours Using AWS Partner Seeq," Amazon Web Services (case study), March 2023, https://aws.amazon.com/partners/success/marathon-oil-seeq/.

[74] Mary Branscombe, "How Microsoft Is Extending Its Cloud to Chevron's Oil Fields," Data Center Knowledge, November 21, 2017, http://www.datacenterknowledge.com/microsoft/how-microsoft-extending-its-cloud-chevron-s-oil-fields.

[75] Biden White House, "National Cybersecurity Strategy."

DEFENSE SECTOR

The defense sector appears to be the slowest in adopting the cloud among the CI sectors surveyed here, perhaps for understandable reasons: defense-related data systems are subject to more stringent and slow-to-change security requirements than any civilian infrastructure sector. However, both the military and large defense contractors have gingerly started placing auxiliary and systems development functions on the cloud. Policymakers have increasingly identified cloud adoption as a linchpin technology for the future of defense information systems: the Acting Chief Information Officer of the Department of Defense, John Sherman, stated in his 2021 Congressional testimony that "[DOD has] made cloud computing a fundamental component of our global IT infrastructure and modernization strategy. With battlefield success increasingly reliant on digital capabilities, cloud computing satisfies the warfighters' requirements for rapid access to data, innovative capabilities, and assured support."[76]

Cloud adoption by the military and defense contractors has been largely facilitated through government-led programs such as Cloud One, which aims to make the cloud accessible across the DOD by acting as a "one-stop-shop" for procuring cloud services from all the hyperscalers.[77] All Cloud One services have been accredited to comply with stringent DOD security requirements, lowering contractual barriers that have traditionally precluded military cloud use. Platform One is a similar initiative providing tooling, development pipelines, and a Kubernetes platform for DOD operators.[78] Platform One aims to help military personnel deploy ready-made, almost-fully-configured cloud products. These initiatives signify an endorsement of cloud in the military, with pathways built out for greater reliance in the future.

While national security considerations often make it impossible to know exactly what kinds of defense data and defense workloads migrate to the cloud, some public information is available. For example, Lockheed Martin has begun moving its test and development instances of SAP HANA, a database used by a variety of applications, onto AWS.[79] Small and medium defense contractors are also transitioning, aided by expertise from the DOD.[80]

In the US military, the Navy stands out as the chief adopter of cloud computing. In 2020, the Navy began moving its planning and tracking tools monitoring hundreds of ships and aircraft, their repair logs, and other operational details, to the cloud.[81] The Naval Information Warfare Center Pacific shifted its DevSecOps environment, a portmanteau of development, security, and operations, called Overmatch Software Armory to the cloud, while other cloud services deliver over-the-air updates to software on some naval vessels, maintain contact between sailors and families onshore, and deliver personnel services and programming to sailors deployed.[82] Apart from the Navy, the US Defense Logistics Agency (DLA) has migrated some of its applications to the cloud, including its procurement management software and a new training suite.[83] In the defense intelligence community, some agencies use cloud services by analyzing satellite imagery and encrypting communications.[84]

[76] House Armed Services Committee, "[H.A.S.C. No. 117-50] Department of Defense Information Technology, Cybersecurity, and Information Assurance for Fiscal Year 2022," 117th Congress House Hearing text, June 29, 2021, https://www.congress.gov/event/117th-congress/house-event/LC67110/text?s=1&r=25.

[77] "Cloud One: Enabling Cloud for Almost Any Department of Defense Use Case," *Air and Space Forces Magazine*, July 2, 2021, https://www.airandspaceforces.com/cloud-one-enabling-cloud-for-almost-any-department-of-defense-use-case/.

[78] "Platform One," US Air Force, accessed April 27, 2023, https://p1.dso.mil.

[79] "Lockheed Martin Case Study," Amazon Web Services, 2017, https://aws.amazon.com/solutions/case-studies/Lockheed-martin/.

[80] Laura Long, "Defense Industrial Base Secure Cloud Managed Services Pilot," EZGSA, April 3, 2019, https://ezgsa.com/tag/defense-industrial-base-secure-cloud-managed-services-pilot/.

[81] AWS Public Sector Blog Team, "Readying the Warfighter: U.S. Navy ERP Migrates to AWS," Amazon Web Services, January 22, 2020, https://aws.amazon.com/blogs/publicsector/readying-warfighter-navy-erp-migrates-aws/.

[82] Liz Martin, "US Navy Deploys DevSecOps Environment in AWS Secret Region to Deliver New Capabilities to Its Sailors," Amazon Web Services, June 29, 2021, https://aws.amazon.com/blogs/publicsector/us-navy-deploys-devsecops-environment-aws-secret-region-deliver-new-capabilities-sailors/.

[83] AWS Public Sector Blog Team, "Defense Logistics Agency Migrates Five Applications to AWS GovCloud (US) Ahead of schedule," Amazon Web Services, January 16, 2020, https://aws.amazon.com/blogs/publicsector/defense-logistics-agency-migrates-five-applications-ahead-schedule/.

[84] "Oracle Cloud for the Defense Department," Oracle, accessed April 27, 2023, https://www.oracle.com/industries/government/us-defense/.

For large defense contractors, cloud deployments have mainly augmented existing on-premises infrastructure—the cloud provides additional computing resources but generally operates alongside on-premises infrastructure rather than as a wholesale replacement. Boeing, in a momentous 2022 decision, decided to use cloud services from multiple CSPs, while in the interim maintaining a mostly on-premises infrastructure.[85] One reason behind Boeing's decision is the cloud's ability to easily scale test environments and store the immense datasets a jet's sensors generate each flight. The company cited a Boeing 787's need to download up to 500 gigabytes (GB) of data per flight, with Raytheon making similar arguments on scalability.[86] Lockheed Martin has recently begun to use cloud's computing capability to help their on-premises capacity for sensitive workloads.[87]

Rates of cloud use in defense seem likely to increase as defense contractors become more acquainted with its risks and benefits. DOD discusses how the "episodic nature" of its mission makes the cloud's scaling capabilities an alluring feature in its 2018 Cloud Strategy.[88] Because the US military and its contractors have been slow to migrate critical systems to the cloud, a cloud compromise would likely not wholly hobble national defense. Less clear is how significant the impacts of such an event would be on important processes such as supply chain and logistics planning. If current defense sector cloud partnerships are successful, then the cloud may grow much more critical to US national defense soon.

FINANCIAL SERVICES INDUSTRY

Financial institutions were among the earliest cloud adopters, but their relatively early experimental use has not yet translated into widespread migration of critical workloads, at least in part, due to the financial industry's substantial data handling and security regulations. Many US-incorporated financial institutions must abide by the requirements of the Basel Accords, Sarbanes-Oxley Act,[89] Payment Card Industry Data Security Standards (PCI DSS), Gramm Leach Bliley Act, bank secrecy acts, and other legal frameworks.[90] The February 2023 report on cloud use in the financial sector from the US Department of the Treasury (or Treasury) said that more than 90 percent of banks had some data or processes in the cloud, but that only 24 percent of North American banks had begun migrating critical workloads to the cloud.[91] The Treasury report suggests that non-bank financial institutions, such as investment advisors and broker-dealers, are also migrating to the cloud relatively cautiously. It noted that cloud adoption has been faster among small institutions, which often rely on third-party service providers that might themselves rely on the cloud. Adoption has also been faster in financial institutions focused on artificial intelligence and machine learning, for which massive computing requirements often functionally require the cloud.[92]

[85] Sebastian Moss, "Boeing Announces Cloud Partnerships with Microsoft, Google, and AWS," Data Center Dynamics, April 6, 2022, https://www.datacenterdynamics.com/en/news/boeing-announces-cloud-partnerships-with-microsoft-and-google/; Aaron Raj, "Boeing Takes to the Cloud with AWS, Google, and Microsoft," TechWire Asia, April 8, 2022, https://techwireasia.com/2022/04/boeing-expands-cloud-services-with-aws-google-and-microsoft/.

[86] Jay Greene and Jon Ostrower, "Boeing Shifts to Microsoft's Azure Cloud Platform," *Wall Street Journal,* July 18, 2016, https://www.wsj.com/articles/boeing-shifts-to-microsofts-azure-cloud-platform-1468861541.

[87] Jeff Morin and Dan Zotter, "Lockheed Martin's Journey to the Cloud," ASUG Annual Conference, May 7, 2019, https://blog.asug.com/hubfs/2019%20AC%20Slide%20Decks%20Thursday/ASUG82404%20-%20Lockheed%20Martin's%20Journey%20to%20the%20Cloud.pdf.

[88] "Department of Defense Cloud Strategy," US Department of Defense, December 2018, https://media.defense.gov/2019/Feb/04/2002085866/-1/-1/1/DOD-CLOUD-STRATEGY.PDF.

[89] Sarbanes-Oxley Act of 2002, Public Law No: 107-204, Government Publishing Office, July 30, 2002, https://www.govinfo.gov/content/pkg/PLAW-107publ204/pdf/PLAW-107publ204.pdf.

[90] See for example, "Cloud Security Implications for Financial Services" Avanade (White paper), 2017, https://www.avanade.com/-/media/asset/white-paper/cloud-security-implications-for-finanical-services.pdf?la=en&ver=1&hash=51A7A54F67E900ADDE743F89AAA96233.

[91] *The Financial Services Sector's Adoption of Cloud Services*, US Department of the Treasury (February, 2023): 27–28, https://home.treasury.gov/system/files/136/Treasury-Cloud-Report.pdf.

[92] US Department of the Treasury, *Financial Services Sector's Adoption of Cloud Services*, 19.

Select financial institutions have more rapidly embraced the cloud as their primary infrastructure for core digital workloads. Capital One, among the largest banks in the United States, announced in 2022 that it had closed all eight of its private data centers and now runs major services entirely in the cloud, including applications working with client data and backup services, an unprecedented move for a financial institution.[93] A catastrophic cloud event leading to the temporary or permanent unavailability of this data would undoubtedly disrupt functionality at Capital One and prove difficult to recover from without cloud-based tools. A massive data breach in 2019 exposed one of the recurring challenges in cloud computing—the trust boundary between a cloud-consuming organization and its CSP. An attacker compromised the firm's AWS-hosted data stores and gained access to personal data for more than 100 million people, in an incident later attributed to a misconfiguration by Capital One that left it vulnerable to a common attack against cloud services.[94]

Nasdaq, the world's largest securities exchange, responsible for matching buyers and sellers across billions of orders, cancellations, and trades each day, moved from on-premises data centers to the cloud in 2014.[95] The stock market's unpredictable trade volume created a need for elasticity best provided by cloud services. Nasdaq's cloud transition proved timely during the COVID-19 pandemic, when the number of transaction records surged to 113 billion a day in March 2020. Since starting up enough on-premises infrastructure for such quantities of data would be virtually impossible in the short-term, losing this capability risk undermining core market functions.[96] A major cloud compromise or disruption could impact the exchange's ability to accurately store the day's transactions, bill customers, and comply with regulatory requirements.

Other major financial institutions have increasingly moved sensitive data to the cloud, though few at the same pace as Capital One or Nasdaq, including Goldman Sachs and the Deutsche Börse Group, which runs the Frankfurt Stock Exchange using a cloud-based tool to analyze investor behavior to offer guidance on better trading strategies rather than any core functions of the exchange.[97]

The financial sector increasingly relies on the ability to rapidly increase and decrease their use of cloud computing resources, to keep pace with unpredictable volumes of financial data. In addition, the increasing complexity of machine learning models that financial institutions use to make decisions about everything from whether a transaction is fraudulent to loan interest rates often necessitates cloud-scale resources. NetApp, a leading cloud data management platform, helped an unnamed "hedge fund division of a major investment bank headquartered in the US" transition their risk modeling functions into Google Cloud to take advantage of its ability to rapidly scale up compute on-demand.[98] Capital One relies on cloud's scalability to manage seasonal transaction surges.[99] Robinhood, a retail investor platform, relied on cloud services to support hundreds of thousands of users at launch.[100] Other banking firms like HSBC and Standard Chartered Standard report using cloud services for customer analytics and even some customer transactions.[101]

[93] Lananh Nguyen, "Banks Tiptoe Toward Their Cloud Based Future," *New York Times*, January 3, 2022, https://www.nytimes.com/2022/01/03/business/wall-street-cloud-computing.html.

[94] Emily Flitter and Karen Weise, "Capital One Data Breach Compromises Data of Over 100 Million," *New York Times*, July 29, 2019, https://www.nytimes.com/2019/07/29/business/capital-one-data-breach-hacked.html; Brian Krebs, "What We Can Learn from the Capital One Hack," Krebs on Security, August 5, 2019, https://krebsonsecurity.com/2019/08/what-we-can-learn-from-the-capital-one-hack/.

[95] Nguyen, "Banks Tiptoe"; "Nasdaq Uses AWS to Pioneer Stock Exchange Data Storage in the Cloud," Amazon Web Services, 2020, https://aws.amazon.com/solutions/case-studies/nasdaq-case-study/; "Trading and Matching Technology," Nasdaq, accessed April 27, 2023, https://www.nasdaq.com/solutions/trading-and-matching-technology.

[96] Nguyen, "Banks Tiptoe"; Amazon Web Services, "Nasdaq Uses AWS to Pioneer"; Nasdaq, "Trading and Matching."

[97] "Goldman Sachs and AWS Collaborate to Create New Data Management and Analytics Solutions for Financial Services Organizations," Goldman Sachs, November 30, 2021, https://www.goldmansachs.com/media-relations/press-releases/2021/goldman-sachs-aws-announcement-30-nov-2021.html; "Deutsche Borse Group Launches Data Analytics Platform in Rapid Time using AWS," Amazon Web Services, 2022, https://aws.amazon.com/solutions/case-studies/deutsche-boerse-case-study1/.

[98] "Cloud Computing in Finance with NetApp Cloud Volumes ONTAP: Case Studies," NetApp, January 18, 2021, https://cloud.netapp.com/blog/cloud-computing-in-financial-services.

[99] David Andrzejek, "Becoming a Fintech: Capital One's Move from Mainframes to the Cloud," CIO, May 17, 2022, https://www.cio.com/article/350288/becoming-a-fintech-capital-ones-move-from-mainframes-to-the-cloud.html.

[100] "Robinhood Case Study," Amazon Web Services, 2016, https://aws.amazon.com/solutions/case-studies/robinhood/.

[101] "Standard Chartered Cuts Risk Grid Costs 60% on Amazon EC2 Spot Instances," Amazon Web Services (YouTube video), November 12, 2019, https://www.youtube.com/watch?v=o-sw9CLY6Go; "HSBC on AWS: Case Studies, Videos and Customer Stories," Amazon Web Services, accessed April 28, 2023, https://aws.amazon.com/solutions/case-studies/innovators/hsbc/.

Other core and critical applications of the financial sector have moved to the cloud, too with Wells Fargo reportedly using Microsoft Azure as the foundation of its "strategic business workloads" and Capital One shifting its disaster-recovery and business-continuity functionality to the cloud.[102], [103] Vanguard, a leading American investment advisor, relies on a similar suite of cloud services as Capital One, reporting near-total adoption of the cloud across more than 850 production software applications in 2021.[104]

The usage of cloud computing services across financial sector firms is growing and attracting notice from policymakers. The Financial Stability Board (FSB), an international body of central bank regulators, expects that there will likely be strong commercial and efficiency incentives for financial institutions to transition to the cloud amongst the financial sector, given its considerable improvements to operational efficiency.[105] In August 2019, Reps. Katie Porter (D-NY) and Nydia Velazquez (D-CA) sent a letter to the Treasury Department which strongly urged naming the leading CSPs (i.e., AWS, Azure, and Google Cloud) "systemically important financial market utilities" (SIFMUs) by the Financial Stability Oversight Council,[106] a designation which would allow the Federal Reserve to more directly examine and regulate CSPs to prevent potential risks to the stability of the financial system.[107] The Dodd-Frank Act created the SIFMU designation in recognition of the fact that the financial sector itself is intricately interconnected and that the availability and functionality of certain components are integral to the continued health and functioning of the financial system as a whole.[108] These systemic dependencies generate additional risk on top of the systemic risks potentially prompted by shared reliance on a handful of CSPs, as an outage at a CSP could lead to a domino effect of cascading failures at other institutions through financial relationships even if they rely on distinct technologies.

Treasury's 2023 report, *The Financial Services Sector's Adoption of Cloud Services*, was a welcome step forward in attempting to map out the complex landscape of cloud service models, noting, "a lack of aggregated data to assess concentration is a key impediment to understanding the potential impact of a severe, but plausible operational incident at a CSP on the financial sector.[109] The report identified as key barriers "(i) the lack of common definitions or identification approaches for critical or material cloud services used by financial institutions, (ii) the lack of a common and reliable method to measure concentration, (iii) different data collection authorities and mandates across FBIIC-member agencies."[110] It further noted the increased difficulties in assessing risk due to "'nth party' dependencies…[as] CSPs provide services to many other third-party service providers that a financial institution may rely on, and also use many sub-contractors, creating indirect dependencies for financial institutions that are more difficult to assess."[111]

The difficulties faced even by the Treasury Department, an experienced sector risk management agency, in assessing the systemic vulnerabilities of the financial sector to cloud incidents are an example of broader measurement challenges common to CI sector regulators attempting to understand the impacts of cloud technology on sector risk.

[102] "Wells Fargo Announces New Digital Infrastructure Strategy and Strategic Partnerships with Microsoft, Google Cloud," Wells Fargo (business wire), September 15, 2021, https://newsroom.wf.com/English/news-releases/news-release-details/2021/Wells-Fargo-Announces-New-Digital-Infrastructure-Strategy-and-Strategic-Partnerships-With-Microsoft-Google-Cloud/default.aspx.

[103] Wells Fargo, "Wells Fargo Announces New Digital Infrastructure Strategy."

[104] Jeff Dowds, "AWS re:Invent 2019 – Jeff Dowds of Vanguard Talks About the Journey to the AWS Cloud," Amazon Web Services (YouTube video), December 12, 2019, https://www.youtube.com/watch?v=8kzOj9cStGo; "Vanguard Increases Investor Value Using Amazon ECS and AWS Fargate," Amazon Web Services, 2021, https://aws.amazon.com/solutions/case-studies/vanguard-ecs-fargate-case-study/.

[105] "FinTech and Market Structure in Financial Services," Financial Stability Board, February 14, 2019, https://www.fsb.org/wp-content/uploads/P140219.pdf.

[106] To the reader, SIFMUs are organizations that, if they fail, would have a catastrophic impact on the stability of financial markets, such as financial clearinghouses. It is important to note that SIFMUs themselves, such as the Options Clearing Corporation, increasingly rely upon cloud services, making cloud providers an even more important part of the financial system. See: "OCC Launches Renaissance Initiative to Modernize Technology Infrastructure," The Foundation for Secure Markets, January 14, 2019, https://www.theocc.com/Newsroom/Press-Releases/2019/01-14-OCC-Launches-Renaissance-Initiative-to-Moder.

[107] Katie Porter and Nydia Velazquez, "Letter to Secretary Mnuchin," US Congress, August 22, 2019, https://velazquez.house.gov/sites/velazquez.house.gov/files/FSOC%20cloud%20.pdf.

[108] "Designations: Financial Market Utility Designations," US Department of the Treasury, accessed May 9, 2023, https://home.treasury.gov/policy-issues/financial-markets-financial-institutions-and-fiscal-service/fsoc/designations.

[109] US Department of the Treasury, *Financial Services Sector's Adoption of Cloud Services*, 57.

[110] US Department of the Treasury, *Financial Services Sector's Adoption of Cloud Services*, 57–58. To the reader, FBIIC is the acronym for the Financial and Banking Information Infrastructure Committee.

[111] US Department of the Treasury, *Financial Services Sector's Adoption of Cloud Services*, 50·

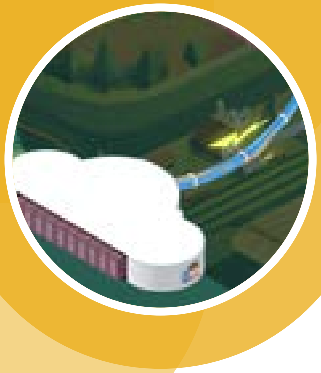

3

THE CLOUD AS CRITICAL INFRASTRUCTURE

The goal of this report is not to evaluate the cloud as a new CI sector. Instead, it addresses the nature of the cloud's criticality on its own merits and to other sectors, in service of specific policy activities which could better handle and govern that criticality. As illustrated above, CI sectors increasingly look to cloud computing to host important workloads. The narrative of the cloud's economic, operational, and security advantages appears uniformly persuasive, even if the rate at which adoption occurs, as well as the operational criticality of the workloads moved to the cloud, varies greatly among sectors. As cloud adoption ramps up, so too will the potential harms to CI from any significant outages, compromises, or cascading failures. Does this mean the cloud services industry, itself, ought to be considered CI?

PPD-21's CI definition is subjective—no quantitative threshold determines criticality to national and economic security, either directly or transitively. Regulatory authorities rely on common knowledge and intuition to make that classification, often focusing on physical, tangible sectors. The majority of the sixteen CI sectors are those with which citizens interact daily (either directly or transitively), such as water, transportation, financial services, and food. Absent a strict methodology, one way to determine criticality is to hypothesize the consequences of an infrastructure's sudden unavailability. The CI sectors of healthcare, transportation and logistics, energy, defense, and financial services increasingly rely on cloud technology for critical workloads. As such, a sudden loss of cloud availability could have cascading consequences of the kind that policymakers sought to avoid by originally designating these sectors as CI. Therefore, the cloud ought itself to be considered as CI, if for no other reason than that it is ever more critical to the operation of already designated CI sectors. The intent of this acknowledgment is not to argue for the addition of cloud as another CI sector—indeed, IT is already a CI sector—but instead to highlight the need for increased scrutiny of cloud computing from existing CI sector regulators and policymakers, given its increasing role as a critical dependency for CI.

Mainstream discussion often glorifies cloud computing as "next-generation" technology, citing cost efficiency, speed, and scalability. Press releases by major CSPs contain myriad references to "transforming" industries,[112] going "all-in" on the cloud,[113] and building next-generation technology.[114] Much policy has sought to speed and streamline government cloud adoption to harness the potential efficiency and cost benefits accordingly.[115] These benefits are real. And this report does not seek to dissuade cloud adoption but instead argues that the real benefits of cloud adoption must also carry measured consideration of the unique risk landscape widespread cloud dependence creates, not just in single cloud services but also in the common infrastructure and architectures that power them.

Despite this, policy discussions about managing cloud risk in the face of critical-infrastructure adoption are lagging. One reason might be that cloud computing remains an opaque topic for many: the technology integrates a mix of old and new computer science concepts, and cloud product offerings are often designed to offload complexity from customers, offering as product a computing paradigm familiar on the surface even if vastly different under the hood.[116] There is then no intuitive designation of cloud computing as critical because its ubiquity and complexity are hidden by design. Where cloud policy discussions are underway, they often focus on the security of specific services rather than the macro interactions in how CSPs design this infrastructure and the emergent properties of widespread adoption. The next section unpacks these properties, specifically **compounded dependence** and **delegated control and visibility**, and how they combine to create systemic risk.

Risk in the Cloud

Cloud computing systems display two properties that create unique risk characteristics: **compounded risk** and **delegated control and visibility**. These properties are the driving cause behind this report, and they arise from cloud architecture and the behavior of cloud infrastructure far more than the security properties of any single cloud service.

Much of the mismatch between risk and policy here is driven by the fact that the cloud, writ large, is a system of systems: ever-growing compute, storage, networking, and management capabilities wired together through the Internet over vast distances and among many jurisdictions, with data, workloads, and capacity in constant flux. Vendors can rip and replace whole generations of infrastructure without notice to customers, so long as they adhere to "uptime" standards. Much of government policy toward cloud risk to date is driven by prescriptive approaches, crafted through the lens of product acquisition and full-system transparency. There are hard limits on the ability of this approach to shape the behavior of these systems, due in equal parts to the complexity and novelty of the cloud writ large and an underdeveloped policymaking toolkit. One might compare the situation to trying to manage an entire economy with policies built for a single commodity or marketplace.

The intuitive approach for cloud risk management is for a single all-seeing entity to observe every system, carefully chart the interdependencies and assigned responsibilities, identify points of failure, and deploy redundancies and fixes to remediate. Besides the obvious technical

[112] "AWS and Atos Strengthen Collaboration with New Strategic Partnership to Transform the Infrastructure Outsourcing Industry," Amazon Press Center, November 30, 2022, https://press.aboutamazon.com/2022/11/aws-and-atos-strengthen-collaboration-with-new-strategic-partnership-to-transform-the-infrastructure-outsourcing-industry; "The University of California, Riverside Enters Into First-of-Its-Kind Subscription-Based Service with Google Cloud to Transform Research and IT," Google Cloud (news release), March 9, 2023, https://www.googlecloudpresscorner.com/2023-03-09-The-University-of-California,-Riverside-Enters-Into-First-of-its-kind-Subscription-based-Service-with-Google-Cloud-to-Transform-Research-and-IT.

[113] "Wallbox Goes All-In on AWS," Amazon Press Center, November 30, 2022, https://press.aboutamazon.com/2022/11/wallbox-goes-all-in-on-aws; "AGL Transforms 200+ Applications, Goes All in on Cloud, and Sets up for Sustained Success," Microsoft News Center, June 9, 2020, https://news.microsoft.com/en-au/features/agl-transforms-200-applications-goes-all-in-on-cloud-and-sets-up-for-sustained-success/.

[114] "AWS and NVIDIA Collaborate on Next-Generation Infrastructure for Training Large Machine Learning Models and Building Generative AI Applications," Amazon Press Center, March 21, 2023, https://press.aboutamazon.com/2023/3/aws-and-nvidia-collaborate-on-next-generation-infrastructure-for-training-large-machine-learning-models-and-building-generative-ai-applications; "Mercedes-Benz and Google Join Forces to Create Next-Generation Navigation Experience," Google Cloud (news release), February 22, 2023, https://www.googlecloudpresscorner.com/2023-02-22-Mercedes-Benz-and-Google-Join-Forces-to-Create-Next-Generation-Navigation-Experience; "Empowering the Future of Financial Markets with London Stock Exchange Group," Official Microsoft Blog, December 12, 2022, https://blogs.microsoft.com/blog/2022/12/11/empowering-the-future-of-financial-markets-with-london-stock-exchange-group/.

[115] Vivek Kundra, "Federal Cloud Computing Strategy", The White House, February 8, 2011, https://obamawhitehouse.archives.gov/sites/default/files/omb/assets/egov_docs/federal-cloud-computing-strategy.pdf; "Strategic Plan to Advance Cloud Computing in the Intelligence Community," Office of the Director of National Intelligence, June 26, 2019, https://www.dni.gov/files/documents/CIO/Cloud_Computing_Strategy.pdf.

[116] Handler, "Dude, Where's My Cloud?"

infeasibility, cloud vendors already build systems beyond the scope of unassisted human management. And unfortunately, current policy is still far from an information environment and regulatory structure where such an approach would be possible.

This information gap is undeniably due in part to internal CSP dynamics—be it a reluctance to share their dependency graphs, algorithms, and infrastructure frameworks or even a lack of internal visibility of the same. Transparency can be perceived as a source of regulatory risk, and there are aforementioned cost advantages of being able to engineer cloud infrastructure to suit CSP needs and cost demands, without prescriptive customer input., Some CSPs may even argue such information constitutes core trade secrets.[117] However, it is worth noting that this relative opacity provides another potential benefit to CSPs: it might allow them to deliver promised services with less accountability and potential legal exposure for the details of—or changes in—how they do so. In the Treasury's cloud report, for example, some financial institutions conveyed that they did not have transparency about how many data centers they were relying on at a CSP until an operational incident occurred.[118]

Of at least equal cause to this information gap is the sheer scale of the cloud—quantifiably the largest-scale information processing systems to date, at the scale of exabytes and growing.[119] Searching for particular nodes of risk and dependence within this sprawling, highly-interconnected system is like searching for needles in an ever-growing haystack. A prospective risk management policy framework for cloud systems and CI will have to interact with these dynamic and complex systems iteratively and with incomplete, often out-of-date information, necessitating flexible, adaptive policy structures designed to intelligently prioritize surface risk, measure behaviors that contribute to or prevent it, and instrument changes in system design or provider behavior to successfully drive it down.

Policy must evolve to match the risk management needs of this system of systems, where many benefits driving adoption also radically transform the relationship between resulting CI customers and their regulators. **Compounded dependence** enables an incredible dynamism to match rapidly shifting cloud workloads while stacking risk on top of key linchpin technologies and common but opaque architectures. Compounded dependence leads to cloud users relying on a few key technologies and designs. Cloud's **delegated control and visibility** means that CI sectors *both* benefit from the engineering prowess of premiere IT firms and struggle to retain full control of and visibility into their systems—an issue magnified by the complexity of the cloud writ large and the corresponding difficulty in fulling understanding the many, constantly shifting dependencies within the cloud. This frustrates the very visibility and capacity for intervention key to managing those dependencies. Cloud systems do not necessarily carry more risk than their on-premises counterparts—they may very well be safer in many respects—but they also create risks not present in on-premises systems and which current policy is ill-equipped to confront.

Compounded Dependence

In cloud systems, many users rely on a small set of similar services and essential technologies. In the pursuit of serving many customers and their varied needs from a single enterprise, cloud systems focus heavily on modular designs and tasking a handful of services with resource orchestration and coordination of hefty workloads. More simply, given the vast burden undertaken by CSPs, redesigning the technologies serving similar purposes across sectors and deployments makes little sense—instead, they reuse technologies across customers. For example, regardless of the sector or organization, CSPs often use the same "identify and access management" (IAM) systems, orchestration systems, and virtualization technologies. Some technologies are even used by multiple CSPs, ranging from hardware components like a handful of Intel CPU SKUs,[120] to operating systems such as Linux, to container management software like Kubernetes. The cloud can therefore make critical sectors operationally dependent on a few linchpin technologies. This compounded dependence is not a unique failure of any particular CSP—it is a fundamental design outcome of the immense complexity of cloud computing infrastructure. It is also a contributing factor to cloud providers' ability to operate thousands upon thousands of massively-scalable environments through automation and standardization.

[117] Paul Teich, "How To Make Public Clouds Transparent," Forbes, October 1, 2019, https://www.forbes.com/sites/paulteich/2019/10/01/how-to-make-public-clouds-transparent/?sh=19f763325746.

[118] US Department of the Treasury, *Financial Services Sector's Adoption of Cloud Services*, 58.

[119] "Will Google Ever Run Out of Storage Space?" MUO, January 19, 2023, https://www.makeuseof.com/will-google-run-out-storage-space.120/.

[120] To the reader, CPU is the acronym for central processing unit, and SKU stands for stock keeping unit.

However, this uniformity also amplifies systemic risk in the event of a failure or outage—risk compounds when much work depends on few services.[121]

Hypervisors are a clear example, sitting at the heart of cloud computing. To provide scalable, economical computing services, CSPs often use one physical server to provide computing resources to multiple customers simultaneously. A hypervisor enables this by partitioning a physical server into several virtual machines, creating a "multi-tenant" environment.[122] To the user, the server appears as a single computer. An attack on a hypervisor could allow a malicious actor to access sensitive data and execute commands in other customers' virtual machines. Multiple CSPs sometimes use the same hypervisor, such as Xen, for both AWS and IBM, further increasing the potential blast radius of a vulnerability.[123]

While every CSP has solved a set of common computing and networking challenges in different ways, they often share a reliance on common or similar technologies and techniques, even at the physical level. The MELTDOWN and SPECTRE vulnerabilities—discovered in Intel processors (most commonly used) and, to a lesser extent AMD processors—illustrate this compounded dependence well. The vulnerabilities posed serious threats to cloud security by enabling attackers to break through the digital walls of multi-tenant environments.[124] Major CSPs made fundamental changes to their infrastructure, in the wake of these vulnerabilities, far beyond patching the affected software.[125] Here, risk blossomed because of the common adoption of similar processors from just two vendors across most CSPs to meet a niche and demanding workload.

Even though on-premises deployments are often more vulnerable individually (and subject to some of the same common-component risks as seen in MELTDOWN and SPECTRE, though eased by avoiding risks specific to multi-tenancy), their diversity and lack of interconnectivity put hard limits on the reach of potential attacks. Malicious actors must conduct reconnaissance and formulate an attack plan for each deployment, slowing their work. In the cloud, however, an adversary could potentially compromise many organizations at once. For example, one disgruntled AWS employee allegedly used a scanner to look for AWS S3 data storage buckets with common misconfiguration patterns, ultimately finding and accessing data belonging to more than thirty organizations (including facilitating the compromise of Capitol One).[126] Uniform solutions can boost the security baseline by bringing cloud customers to a universally higher standard. However, they can also grow the blast radius and magnitude of vulnerabilities and compromises.

Compounded dependence in such large, shared systems also means that, even without malicious interference, one faulty update can cause system-wide failure. For example, on November 25, 2020, AWS cloud services for the eastern United States suffered a severe service outage from an update to a core cloud service relied upon by many AWS systems: Amazon's Kinesis data ingestion engine. The root cause was a capacity addition to front-end servers, which overstressed the operating system on which it was running.[127] The Kinesis outage triggered several downstream issues, including increased errors and latencies for Amazon's CloudWatch monitoring service and its IAM service, Cognito. The cascading chain of failures meant trouble for the availability of

[121] "Recommended Best Practices for Administrators: Identity and Access Management," US National Security Agency (NSA) with the Cybersecurity and Infrastructure Security Agency (CISA), March 2023, https://media.defense.gov/2023/Mar/21/2003183448/-1/-1/0/ESF%20IDENTITY%20AND%20ACCESS%20MANAGEMENT%20RECOMMENDED%20BEST%20PRACTICES%20FOR%20ADMINISTRATORS%20PP-23-0248_508C.PDF.

[122] To the reader, orchestrators play a similar role in deployments using a container-style deployment. See: Handler, "Dude, Where's My Cloud?"

[123] "Xen Vulnerability Allows Hackers to Escape OS VM and Own the Host," Virtalica, accessed April 28, 2023, https://virtalica.com/2016/08/26/xen-vulnerability-allows-hackers-escape-os-vm-host-amazon-aws-rackspace-ibm-affected/.

[124] Jann Horn, "Project Zero: Reading Privileged Memory with a Side-Channel," *Project Zero* (blog), January 3, 2018, https://googleprojectzero.blogspot.com/2018/01/reading-privileged-memory-with-side.html; Birgit Baustädter, "Following on from Meltdown and Spectre: TU Graz Researchers Discover New Security Flaws," Tu Graz News, May 14, 2019, https://www.tugraz.at/en/tu-graz/services/news-stories/tu-graz-news/singleview/article/nach-meltdown-und-spectre-tu-graz-forscher-entdecken-neue-sicherheitsluecken0/.

[125] Daniel Firestone et al., "Azure Accelerated Networking: SmartNICs in the Public Cloud," Usenix, accessed April 28, 2023, https://www.usenix.org/conference/nsdi18/presentation/firestone; Paul McLellan, "HOT CHIPS: The AWS Nitro Project," Cadence, October 2, 2019, https://community.cadence.com/cadence_blogs_8/b/breakfast-bytes/posts/the-aws-nitro-project; Tim Anderson, "Managing the Linux Kernel at AWS: 'A Large Team of Security Experts' Dealing with Fallout from Spectre, Meltdown Flaws," *The Register*, December 10, 2019, https://www.theregister.com/2019/12/10/linux_kernel_aws/.

[126] Michael W. Heiligenstein, "Amazon Web Services (AWS) Data Breaches: Full Timeline Through 2022," *Firewall Times*, April 6, 2023, https://firewalltimes.com/amazon-web-services-data-breach-timeline/; Tara Seals, "Capital One Attacker Exploited Misconfigured AWS Databases," DarkReading, June 20, 2022, https://www.darkreading.com/attacks-breaches/capital-one-attacker-exploited-misconfigured-aws-databases.

[127] "Summary of the Amazon Kinesis Event in the Northern Virginia (US-EAST-1) Region," Amazon Web Services, November 25, 2020, https://aws.amazon.com/message/11201/.

many AWS services and compromised even its ability to provide status updates during the outage because its status page relied on the affected Cognito service. This cascading chain of failures demonstrates the dangers of a tendency to build many systems on top of a few key dependencies. For the public, the prolonged Kinesis failure in Amazon's US-East-1 region (located in Virginia) outed online services "temporarily" for many big-name entities like the New York Metropolitan Transit Authority website, Google's 1Password, Autodesk, Coinbase, and the *Washington Post*.[128]

On the surface, these examples illustrate the complex webbing of cloud systems and networks. Traditional computing networks have similar functions like logging, job assignment, and data analytics. Yet systems, such as activity logging, which tend to be auxiliary in traditional, on-premises deployments, are essential to cloud functioning for workload coordination across thousands of customers and hundreds of data centers. "Smaller" services have a much bigger and more vital role in the cloud because they are critical dependencies of data systems for hundreds of thousands of cloud customers at once. Leveraging a handful of technologies in this way certainly has benefits—it creates efficiencies by re-using the best available solutions and gives CSPs the ability to deploy many resources on securing specific, key dependencies. However, all software, no matter how well-maintained, can have vulnerabilities, outages, or downtime. The cloud is unique in how deeply it concentrates risk for hundreds and thousands of entities by stacking that risk on a few specific technology nodes.

Delegated Control and Visibility

One of the most appealing selling points of cloud systems is delegation—no longer does a CI operator need to shoulder the entire burden of managing the operational and security complexity of managing an entire IT system. Instead, by relying on a cloud offering, they can share that work with a CSP, usually an entity with vast resources, expertise, and experience. This shared handling has been enshrined with respect to security in ideas like the "shared responsibility model," which articulates the components of cloud security for which a CSP is responsible versus those responsibilities which fall on the cloud customer.

Many cloud breaches have originated in breakdowns of the shared responsibility model.[129] While CSPs typically provide tooling, such as access controls, customers must know how to configure and use them correctly, which can pose a challenge for an organization migrating to the cloud for the first time. Often, the issue arises from the cloud user, who might fail to correctly configure restrictions on sensitive resources, or use combinations of permissions with unwanted effects, including paths to improper escalation. Policymakers, especially in CI sectors, must examine the shared responsibility model carefully to ensure that it is also a shared accountability model: CSPs must provide secure-by-default configuration options and adequate support for their customers to ensure that their deployments are configured securely to "shift the burden" for better security onto CSPs,[130] rather than their less-well-resourced customers.

Shared responsibility models only work where technologies, by design, actually allow users to share that responsibility for defense. Risks created by deficiencies in the shared responsibility model—such as confusing or insecure-by-default controls offered by a CSP—are systemic, not in the sense that they all fail as one, but in that a common misconfiguration error at one CSP can impact many of its customers. Attackers will often identify a vulnerable configuration and scan for its presence in as many cloud deployments as they can find.[131] Pushing CSPs to reduce these instances in the design of their services is one step. Addressing how risks in the infrastructure of cloud services amplify these would be even better.

The shared responsibility model does not apply to security alone: users delegate control over the architecture of the computing and storage and networking resources they use to the CSP. This delegated control and visibility model presents challenges because it impairs visibility and control for CI customers *and* regulators. CI customers often lack data or access that enables them to partake in

[128] "Amazon Details Cause of AWS Outage That Hobbled Thousands of Online Sites and Services," GeekWire, November 30, 2020, https://www.geekwire.com/2020/amazon-details-cause-aws-cloud-outage-hobbled-thousands-online-sites-services/.

[129] Drew Wright, "The State of Cloud Security 2020 Report: Understanding Misconfiguration Risk," Cloud Security Alliance (blog), May 5, 2020, https://cloudsecurityalliance.org/blog/2020/05/05/the-state-of-cloud-security-2020-report-understanding-misconfiguration-risk/; Tod Beardsley and Kwan Lin, "Cloud Misconfigurations Report," Rapid 7, April 2022, https://www.rapid7.com/thank-you/2022-cloud-misconfigurations-report/.

[130] Biden White House, "National Cybersecurity Strategy."

[131] Beardsley and Lin, "Cloud Misconfigurations Report"; Guy Alvarenga, "11 AWS Misconfigurations and How to Avoid Them," CrowdStrike, September 12, 2022, https://www.crowdstrike.com/cybersecurity-101/cloud-security/common-aws-misconfigurations/.

risk management, instead requiring trust in the CSP, and regulators often lack direct supervisory or audit access to the entities operating the IT systems that underpin their sector's function.

Though not necessarily common practice, with on-premises data centers, it is theoretically possible for an organization to audit its data system, understand the end-to-end model of where and how their data moved within the system, and identify potential nodes of risk or operationally critical data system components. In contrast, under the delegated control model, contractual negotiations, rather than direct oversight, are often the only opportunities organizations have to understand qualities such as requirements for how much downtime a service may have.[132] Treasury, in its report, found that "some financial institutions conveyed that there were gaps in their ability to assess the resilience of their configuration of a cloud service." The report goes on to identify critical challenges in making these assessments, including:

- "(i) difficulty in understanding their responsibilities or effectiveness of their choices for configuring the cloud services for the appropriate level of resilience;
- (ii) the lack of specific recovery time objectives in some contracts with CSPs;
- (iii) the lack of specific incident notification and response procedures in some contracts with CSPs; and
- (iv) the lack of detail in cloud service documentation regarding resilience dependencies, such as a CSP's reliance on other suppliers of IT services or internal CSP resources (such as other CSP operating regions)."[133]

Cloud platforms have their own complex webs of operational dependencies on other services and resources—from both those provided by the CSPs themselves and some provided by external entities—that add another layer of complexity, particularly when those relationships are opaque to CSP customers. Service-level agreements generally allow a CSP to subcontract systems, creating a multi-layered supply chain. Subcontracting and outsourcing cloud software and hardware occur at virtually every service level, from IaaS to SaaS. For instance, one SaaS provider might be a client of another IaaS provider. The customer would only interact directly with (e.g., pass data to) the SaaS provider, not knowing whether and how their data passes through to a third-party storage service. With every instance of subcontracting and layering, users lose some control, often invisibly.

These layers of dependence and the broader delegated control and visibility model can create visibility challenges for cloud customers when a security incident occurs. Cloud customers can face challenges identifying and acquiring the logs needed to analyze an incident due to the sheer volume of services within the cloud as well as challenges in obtaining log data that might reside with the CSP (perhaps because the data concerns multiple tenants). The sheer scale of the cloud the volume of logs it generates and the lack of standardized logging functionality all compound these issues.[134]

Further complicating the visibility issue is the fact that CSPs themselves can struggle to clearly understand the full set of software, services, and infrastructure their own cloud offerings rely on. Executive Order 14028 tasked the National Telecommunication and Information Administration (NTIA) with standardizing a format for software bills of materials (SBOMs),[135] but it had to defer the question of SBOMs for SaaS and IaaS cloud environments. NTIA explained: "The service provider must not only track metadata from the software supply chain of the software they are responsible for producing, but in the infrastructure stack that supports the application, whether under the direct control of the provider or from some external service provider. [...] Capturing meaningful metadata about the full application stack and third-party services is ongoing work, but not yet standardized or

[132] "Can You Recovering Losses Sustained during a Cloud Outage," *InformationWeek*, August 4, 2022, https://www.informationweek.com/strategic-cio/can-you-recover-losses-sustained-during-a-cloud-outage-; Kharmela Mindanao, "Cloud Downtime Explained (& What You Can Do About It)," Intelligent Technical Solutions (blog), accessed April 28, 2023, https://www.itsasap.com/blog/cloud-downtime-explained; Tobias Mann, "Oracle NetSuite Datacenter Plunges Offline for a Day, Customers Warned of Data Loss," *The Register*, February 15, 2023, https://www.theregister.com/2023/02/15/oracle_netsuite_down/.

[133] US Department of the Treasury, *Financial Services Sector's Adoption of Cloud Services*, 53.

[134] Martin Herman et al., "NIST Cloud Computing Forensic Science Challenges," (Gaithersburg, MD: National Institute of Standards and Technology, August 2020), https://doi.org/10.6028/NIST.IR.8006.

[135] President Biden, Executive Order 14028: "Improving the Nation's Cybersecurity," US General Services Administration, May 12, 2021, https://www.gsa.gov/technology/technology-products-services/it-security/executive-order-14028-improving-the-nations-cybersecurity.

sufficiently mature for cross-organization implementation."[136] If cloud services challenge even their creators' ability to map and understand ever-changing, multi-layered dependencies, then customers and regulators have little hope of understanding them enough to perform their own risk assessment or management without improved standardization of and tooling cloud transparency.

The delegated control and visibility model makes a full assessment or understanding of the risks involved with cloud use difficult for both cloud customers and regulators by obfuscating the possible modes of operational failure and obscuring risks that arise from choices CSPs make in building and operating their infrastructure. CSPs' opacity here protects their ability to make business and technology decisions without user input, a practice that may be hard to change without considerable incentive. Indeed, this arrangement is somewhat necessary, as cloud users lack the resourcing and capacity to manage CSP systems, and in fact, part of the price tag of cloud products is the very offloading of that work from user to provider. It might arise too from the simple fact that CSPs themselves have an incomplete picture of their own technology stack. Whatever the cause, it means that CSPs are the sole entity currently empowered and asked to manage systemic cloud risks, with relatively little oversight of how they choose to do so and how successful their efforts have been. Like security, risk management should be a shared responsibility—among users, regulators, and CSPs rather than entirely the latter. Regulators must ensure CSPs collect the information needed to appraise and manage cloud risks and share it with customers and regulators that also hold deep equities in ensuring cloud security and resilience.

The End Result

Taken together, **compounded dependence** and **delegated control and visibility** create new forms of risk that are both systemic and relatively opaque. To serve large numbers of customers and execute complex, dizzying workloads, CSPs lean on a handful of core technologies, increasing the risks and blast radius of incidents in them. At the same time, while the effective risk management of such an arrangement hinges on visibility into shared infrastructure and linchpin services, the widespread delegation of control from user to CSP diminishes the respecting user-entity capacity. The end result is users—sometimes the majority of entire sectors, including some CI sectors—relying on a few core technologies with little insight into or influence over those technical arrangements.

The cloud is a "system of systems,"[137] in which many components operate independently, but rely on each other, driving more systemic risk than that created by on-premises data centers managed by individual organizations. Customers cannot neatly manage these potential risks, even if they make use of multiple cloud services: disparate cloud systems might appear independent, but that appearance may be illusory, with many systems fundamentally relying on them. For example, the failure of a basic logging service would bring down all of them. The rapid addition of new services—coupled with semi-autonomous cloud business units—increases the likelihood of catastrophic, unsuspected failure cascades.

More and more CI sectors are coming to rely on this complex system of systems, arranged and managed by just a few firms and built around key technical bottlenecks under a shroud of opacity. Many cloud customers currently rely solely on the CSP to effectively govern these risks, both in deployed infrastructure and the architecture of how different systems map to and depend on each other. Users have few ways to hold CSPs accountable to ensure this governance takes place—and policymakers have *no means* to specify or measure—if this governance is effective. Yet cloud adoption continues to grow swiftly, all while essential aspects of the public interest in the security and safety of widely used technology systems remain missing in action. The United States needs to act quickly to improve clarity on this system of systems so crucially depended upon by many of the nation's CI sectors.

[136] National Telecommunications and Information Administration, "The Minimum Elements For a Software Bill of Materials (SBOM)," US Department of Commerce (July 12, 2021): 15, https://www.ntia.doc.gov/files/ntia/publications/sbom_minimum_elements_report.pdf.

[137] John Klein, "Cloud Computing: An Architecture-Centric View," Software Engineering Institute, Carnegie Mellon University, (PowerPoint presentation, 2018): 172–73, https://apps.dtic.mil/sti/pdfs/AD1087041.pdf.

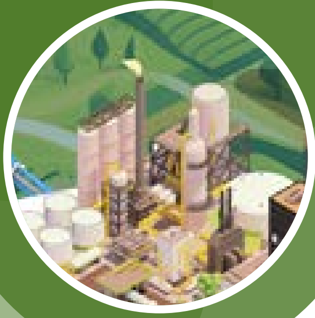

4

POLICY RECOMMENDATIONS

There are several ways to address these two emergent properties of cloud risk. One of the fastest ways, and the one this report focuses on, to improve cloud visibility within CI sectors will be to leverage the existing framework for critical-sector risk management and establish "Cloud Management Offices" (CMOs) within Sector Risk Management Agencies (SRMAs).[138] These CMOs should seek to evaluate sector dependence on cloud computing, understand and outline best practices for cloud adoption and migration based on sector case studies, identify idiosyncratic points of risks (especially those that might be made unintentionally worse by sector-specific regulations), and highlight sector-specific needs such as continuous availability or security for highly sensitive data. By fusing the sector-specific risk management expertise of the SRMA with an entity specifically tasked to understand the paradigmatic changes of cloud computing, SRMA CMOs could assess risks and establish policies to help its sector balance the new risks and benefits of cloud computing. These offices would also provide a useful point of focus for new budget and hiring authorities to build cloud security competency in each SRMA while reducing start-up time relative to trying to establish wholly new entities which would need to build sector expertise and relationships from scratch. For smaller SRMAs, the Cybersecurity and Infrastructure Security Agency (CISA) could assist in the establishment of a CMO with sufficient technical expertise through close partnership with and support for that SRMA. Regardless of size, CISA, through its established cloud expertise, should assist all SRMA CMOs in setting up cloud risk-assessment capabilities.

The first task for each newly created CMO should be a **survey of cloud usage within the respective sector** to understand the degree of dependence that sector actors have on the cloud for critical functionality and the barriers to their current ability to understand and manage associated risks. Second, the CMO should **leverage this information to assess existing sector risk reporting and management requirements** and determine adequacies for capturing and gauging cloud risk, as well as use this information to **outline sector-specific best practices in cloud migration and adoption** based on successful case studies and identified best practices. Where gaps exist that impact regulator visibility into cloud-associated risks—such as a lack of requirements to report on which CSPs are used, or the specifics of CSP contracts—the CMO should recommend updates or new reporting requirements to address such gaps.

[138] To the reader, only a few critical sectors currently have guidelines or programs that encompass cloud computing.
The financial sector has the PCI DSS Cloud Computing
Guidelines, and the healthcare sector has HIPPA.

Throughout, **CISA, in its capacity as the "quarterback" for federal cyber risk management**,[139] **should play a facilitating role**. CISA can offer resources such as reports or frameworks to assist CMOs in surveying cloud usage. It might also work with SRMAs as they construct sector-specific reporting requirements for CSPs to harmonize reporting requirements across sectors and minimize burdens on CSPs while ensuring each SRMA has the information it needs to manage risk.

Separate from creating these CMOs to manage CI-sector cloud oversight, need for **a new entity or new authorities** might surface to more directly oversee the cloud sector itself, **focusing on cross-sector risks to cloud stability and security**. Such an entity could center efforts on strengthening transparency, resilience, and security at major CSPs, benefitting whole-of-system integrity while avoiding possible duplicative or incomplete efforts from the sole reliance on CI sector-specific lenses on the cloud. Such a new structure will need sufficient authorities to obtain data to enable visibility into cloud software stacks and associated risks, as well as effective tools to translate its insights into policy interventions.

This report recommends, therefore, makes three interconnected policy recommendations.

1. **Systematically evaluate cloud computing use in critical sectors:** CMOs should conduct formal surveys with key actors in their critical sectors to determine:
   - **A.** The overall dependence on the cloud by that sector, with special attention to features and functions that are essential to ongoing operations;
   - **B.** The presence and adequacy of fallback systems for workloads essential to ongoing operations;
   - **C.** An overview of models of hybrid cloud adoption, including the distribution of workloads, the interaction between on-cloud and on-premises data systems, and any risks specific to hybrid cloud deployments;
   - **D.** The potential impacts of different cloud incidents on the sector and follow-on effects;
   - **E.** The distribution of CSPs and cloud services used;
   - **F.** The distribution of contractual terms with CSPs pertaining to sector-specific requirements such as continuous availability or data security; and
   - **G.** Lessons learned from existing examples of sector cloud migrations and configurations, including sector-specific best practices.
2. **Survey and update cloud policies and resources:** CMOs should survey existing reporting requirements for data systems and evaluate how well they capture cloud-related risk factors. Where required, CMOs should make policy recommendations to extend regulations to ensure their SRMA has adequate information to manage sector risk. Additionally, CMOs should evaluate whether they have provided adequate resources such as best practice guides and runbooks for sector participants migrating to the cloud and seek to create and provide such resources based on their survey of sector cloud use where appropriate.[140]
3. **Develop a structure for cross-sector cloud risk oversight:** Congress should identify a policy structure for an entity to oversee and manage systemic, cross-sector cloud risk.

#1: Systematically Evaluate Cloud Computing Use in Critical Sectors

There is little information available for regulators to assess cloud adoption within CI sectors, and often even less available for them to assess the degree to which critical workloads depend on the cloud and how operationally resilient these cloud workloads are. General market surveys on cloud adoption tend to miss the distinction between the migration of mission-critical versus auxiliary workloads and data. They lack granular information such as particular CSPs, contractual terms, and configuration options essential for regulators to appraise sector risk. The lack of transparency has made determining whether a critical sector—or at least some of its major players—relies on cloud computing for daily operations difficult, making risk management by regulators challenging in turn.

To achieve appropriate operational awareness, each CMO should survey its critical sector to determine its degree of cloud dependence and maturity. CMOs should conduct surveys to capture information most essential to

[139] Jen Easterly, "Written Statement: [as] Nominee for Director of the Cybersecurity and Infrastructure Security Agency Before the U.S. Senate Committee on Homeland Security and Governmental Affairs," June 10, 2021, https://www.hsgac.senate.gov/wp-content/uploads/imo/media/doc/Prepared%20Statement-Easterly-2021-06-10.pdf.

[140] To the reader, the term "runbook" refers to a set of standard written procedures for completing repetitive IT steps or processes.

assessing the sector's cloud dependence and its degree of systemic cloud-related risk, including: which workloads have transitioned to the cloud and how necessary those workloads are for core operations; the presence or availability of non-cloud data systems and whether such systems can substitute for the cloud in the event of an outage; which CSPs the sector entities use; relevant contract terms such as uptime expectations and responsibilities allotted under the shared responsibility model; and any barriers reported by sector regarding their ability to use the cloud in a risk-aware way, such as inability to obtain desired contractual terms or CSP unwillingness to disclose information about its software supply chain or stack. To assess cloud maturity, CMOs should also seek to understand whether the subject organization has plans and processes in place for management of cloud incidents—including cybersecurity incidents and continuity outages—and to understand other challenges the organization has faced in their adoption of cloud computing.

Such a survey will not immediately reduce cloud adoption risks, but it will provide a critical window into the nature of sector cloud dependence to aid in developing subsequent coordination, guidance, and possibly regulation. The survey will help regulators understand variance in the types of shared responsibility models adopted throughout the sector, as well as the degree of control and insight sector actors typically have into cloud management and risks. Survey results could be public with sufficient aggregation and anonymization. Much like Treasury's report on the cloud in the financial sector, access to the results helps policymakers and decision-makers better understand the emerging cloud risk landscape.

#2: Survey and Update Cloud Policies and Resources

Cloud deployments in some CI sectors are already subject to regulations for data systems security and, more infrequently, resilience and operational continuity. For example, the Treasury report details existing financial sector regulations related to data system security (generally issued by different regulators under the authority granted in the Gramm-Leach-Bliley Act) and data system resilience and availability [under regulations promulgated by the US Securities and Exchange Commission (SEC) and the Consumer Financial Protection Bureau] that often propagate via contract to third-party service providers to financial institutions.[141] In healthcare, CSPs are subject to requirements by virtue of sector data system regulations, most typically HIPAA, which requires a particular standard of confidentiality for patient data and requires providers to enter into "business associate agreements" that govern the security behavior of any third party contracted for data processing.

In some cases, existing sector oversight authorities may allow regulators to collect information about sector cloud use, examine CSPs that provide services to their sector, or require that sector participants thoroughly vet and contractually enforce protections against cloud risks, including security breaches and outages. For example, entities such as the Federal Reserve have conducted some limited direct examinations of CSPs under their existing authority within the Bank Secrecy Act.[142] However, many existing regulations may not specify the ability for regulators to directly examine CSPs providing core functionality for their sector or often have ill-designed language that fails in requiring sector cloud customers to obtain necessary system assurances or information through the process of contracting with a CSP. CMOs should survey existing sector information system requirements to understand their adequacy in an era when many—perhaps most—sector participants rely on the cloud.

CMOs can later synthesize the policy survey results with those from the cloud-usage survey to inform potential policy recommendations. Such policy recommendations might include updating existing sector reporting requirements to capture essential data on risks associated with cloud computing, granting regulators the ability to directly examine or interact with CSPs, or requiring sector participants to obtain contractual security and resilience guarantees from their CSP and prescribing means by which sector participants must verify such properties.

Regulations will need sufficiently nuanced updates to account for the wide variance in cloud usage and should reserve stringent requirements for the most critical operational systems. For example, it may not be worth the oversight burden to ask sector participants to compile reports on their use of the cloud for auxiliary data storage or on their use of common enterprise SaaS tools, such as Microsoft Outlook. The table below illustrates some examples of labeling particular cloud services relative to their criticality, but SRMA CMOs should adapt workloads

[141] US Department of the Treasury, *Financial Services Sector's Adoption of Cloud Services*, 32.

[142] Porter and Velazquez, "Letter to Secretary Mnuchin."

and reporting timelines to match the relative criticality of different functions within their sector.

Additionally, CMOs can provide helpful resources beyond regulation to drive down risk, including best-practices guides and cloud-migration runbooks adapted for their sector's specific needs. Such documents should reflect their experience talking to sector participants about their cloud migration experiences, as well as potentially through interviews with CSPs and other stakeholders. In this process, CMOs might consider collaborating with other CMOs to identify cross-sector best practices for cloud migration and to leverage insights or guidance from cloud-mature sectors and organizations.

#3: Develop a Structure for Cross-sector Cloud Risk Oversight

Regulators need better data to understand complex webs of interdependencies that create risks in the cloud environment. As this report indicates, many other CI sectors rely on the cloud, meaning that the potential impacts of cloud risks span across sectors, and improvements to the security or resilience of cloud offerings can have multi-sector benefits. Examples of cloud failures and outages from the sections above illustrate that particular components, such as identity and access management or resource allocation services, can be failure points for the function of a CSP's entire system. Understanding and mapping system risk is a necessary first step to determine which systemically important cloud components might need greater scrutiny, testing, and support. This report, therefore, recommends Congress consider potential designs for an oversight entity to collect data related to cloud and cloud risks and suggest interventions to strengthen cloud security and resilience to benefit the multiple other CI sectors that rely on the cloud. The key will be creating reporting structures that capture dependencies and relationships to inform regulators about vital potential failure points, while becoming standardized enough to facilitate easy reporting and analysis.

Open questions remain about the structure, information requirements, and required authorities for a centralized cloud risk oversight agency. Congress could take the first steps towards its creation by convening a task force or other intragovernmental body to inform its design and implementation. Like the Cyberspace Solarium Commission (CSC), such a task force could include membership from bipartisan, Congressional cloud leaders, heads of stakeholder agencies—in this case, SRMAs and other cloud regulators such as the Federal Trade Commission—alongside operational experts in cloud security, resilience, and risk from outside of government. Like the CSC, such a task force could issue recommendations for the creation of a cloud oversight entity and other legislative measures to improve cloud resilience and then work with Congress to implement those recommendations in the form of legislation.

The structure of the cloud risk oversight entity must contend with at least four basic questions.

1. Where in government should a cross-sector cloud risk oversight entity reside?

2. What data will it need from CSPs to allow mapping complex webs of risk and software dependence?

   A. What information needs to be included in the reported data, and in what format?

   B. Which CSPs must report?

      i. At what point does a CSP become systemically important enough to merit reporting requirements of this kind?

      ii. Should reporting requirements apply only to "hyperscaler" IaaS providers, or to key SaaS providers as well?

   C. How frequently does data need to be reported?

3. How can such an entity obtain the data it needs?

4. What should the entity do with the information it obtains about cloud risk?

   A. How can it propagate information to SRMA CMOs and other organizations tasked with managing sector-specific cloud risks?

   B. How can it share useful risk-related information back to the CSPs themselves?

In its report, the Treasury Department acknowledged the need for interagency coordination to manage cloud risks, for which it created a new Cloud Services Steering Group to "promote coordination and collaboration among US financial regulators on these challenges."[143] This is an essential first step towards much-needed collaboration between sector regulators to manage cloud risks, but the Treasury Department alone cannot create a body that reaches beyond the financial sector to collaborate with

[143] US Department of the Treasury, *Financial Services Sector's Adoption of Cloud Services*, 8.

other critical sectors facing the same questions and challenges. An intragovernmental task force could jumpstart such a process.

3.1: Where in government?

This recommendation intends to lead to the creation of a new entity or the identification of an existing, capable one to task with studying and managing cloud risks. An obvious home for such an entity would be within CISA, given the agency's existing mandates around coordination for CI risk management and the Department of Homeland Security's status as the SRMA for the IT CI sector. Given CI risk management is an exercise in managing the behavior of diverse private sector actors, CISA, with its emphasis on private-public partnership, is a more fitting home than agencies focused on standards for government cloud use such as the General Services Administration. This effort would have more in common with CISA's existing work around strengthening private sector adoption of secure by design and default practices than with programs such as the Federal Risk and Authorization Management Program (FedRAMP) that focus on individual cloud product security.

A task force designing the structure of such an entity might consider whether CISA already has the authorities required to establish such an office or initiative, or whether legislative action must assign CISA the appropriate coordinating authority with respect to the cloud. Potential redistribution of responsibilities under a PPD-21 update could also alter considerations about where and how to create such an entity. Additionally, the task force might consider whether CISA requires additional funding or hiring authorities to ensure that such an effort is appropriately supported to pursue an ambitious, cross-sector cloud risk-management mission.

3.2: What data is needed?

This question requires a potential task force to consider the type and format of data reporting needed for an oversight agency to effectively evaluate and identify systemic points of vulnerability for important CSPs. A worthwhile starting point for such an initiative would be the creation and collection of Software Bills of Materials (SBOMs) from cloud providers. SBOMs for CSPs have been a topic of much discussion, though coalescing on a standard methodology and format has been a persistent challenge.[144] A first step for this task force on the issue of data reporting might involve selecting a standard methodology and format for cloud SBOMs. Additionally, any separate work toward standardizing cloud SBOMs will be a meaningful step forward.

However, the task force should also consider whether other forms of information beyond SBOMs would be necessary or helpful to this effort, including system-architecture diagrams (with third-party dependencies) to identify technical nodes with over-leveraged risk, documentation of CSP security and resilience processes and behaviors, high-level architectural documentation, and information about contingency and resilience plans detailing potential catastrophic failure scenarios, information about physical data center locations, the relative distribution of data between them, and more. The reporting data format will need to balance the need to be granular enough to allow an oversight entity to identify specific nodes of risk, without creating intractable and burdensome volumes of data about the bespoke configurations of countless different cloud customers.[145]

Policymakers should consider whether third parties that provide services critical to cloud functionality are necessary reporting requirements to better understand potential chains of transitive dependence. The task force should also address the frequency of reporting parameters, balancing the need for up-to-date risk analyses with the administrative burdens on both CSPs and the receiving entity of each data update cycle.

3.3: How can this data be obtained?

Another open question is how a centralized oversight entity could obtain the data needed from CSPs. Understanding the intricacies of cloud systems requires access to details of a CSP's software stack and supply chain, which it would likely guard as a business secret. A structure for systemic cloud-risk oversight will need to avoid requiring cloud customers themselves to attempt to map cloud dependencies, which would likely lead to mistakes and delays. Instead, any structure would need to shift the responsibility onto those best equipped to share information through the creation of SBOMs and similar artifacts: the CSPs themselves.

There are multiple avenues for designing reporting that meet these requirements. One involves leveraging

[144] Beth Pariseau, "CNCF, CISA Address Hurdles to SBOM for Cloud Security," TechTarget IT Operations, accessed May 18, 2023, https://www.techtarget.com/searchitoperations/news/252522983/CNCF-CISA-address-hurdles-to-SBOM-for-cloud-security.

[145] To the reader, the term "bespoke" refers to custom or tailored specifications.

existing SRMAs with regulatory authority to oversee and examine third-party CSPs that provide integral services for their sector by having these SRMAs request SBOMs and similar information from CSPs that provide services to their sector and then having them aggregate this information up to the central oversight entity. For example, regulators for designated financial market utilities (DFMUs)—such as the Federal Reserve Board, SEC, and Commodity Futures Trade Commission—can examine a service provided by another entity when that "service is 'integral' to the operation of the DFMU,"[146] which might already apply to certain CSPs. Many CMOs/SRMAs could require additional legislative authority to directly examine CSPs serving their sector, yet this whole-of-government approach to authorities for third-party review appropriately models against the complexity and scale of this industry.

In lieu of new authorities, sector regulators could also require that regulated CI entities obtain information such as SBOMs directly from CSPs as part of their contractual agreement and then aggregate them up to a central oversight authority. However, this system has downsides in that the presence of an intermediary (the sector cloud customer) might increase CSP fears about revealing proprietary business information through such a process, in addition to the overhead created through the need to cross-reference and aggregate distinct system architecture diagrams to gain a whole-of-cloud perspective.

A third potential approach could involve using legislation to grant authority to a centralized entity to directly obtain SBOMs and other information from major cloud services to aid risk monitoring. This approach would help provide the centralized entity with a holistic overview of the cloud risk and reduce regulatory reporting burdens on CSPs, but it would be difficult to implement without legislative action.

3.4 How should data be used to reduce cloud risk?

An important advantage of designating an entity for cloud risk oversight is its increased ability to review and strengthen critical dependencies and shared points of failure across industries. Such a structure also avoids the laborious duplication and potential piecemeal effects of making each SRMA CMO solely responsible for functionally managing risks associated with the cloud through the lens of its own industry. However, SRMA CMOs will still be the best-positioned to measure and manage the specific impacts of cloud risks on their sector and, therefore, will need robust and open channels of communication with any centralized entity to request and receive information about risk-related findings that may be relevant to their own sector risk management.

With data in hand, a centralized oversight entity would have several potential options for interventions to reduce risk. Potentially, it could work directly with CSPs to identify systemically important components and dependencies for testing, auditing, and hardening. It could develop a certification scheme for cloud resilience commensurate with CI needs, which CSPs could obtain to foster trust with CI customers. While the task force need not delineate every possible way an oversight entity would use cloud data to reduce risk, identifying broad categories of viable functions will be helpful for shaping considerations of what data is needed, as well as suggesting how the body will achieve longer-term objectives. The task force should also consider guidelines around the reporting and use of CSP-provided data to protect its confidentiality.

[146] US Department of the Treasury, *Financial Services Sector's Adoption of Cloud Services*, 36.

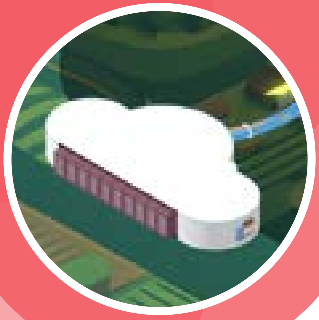

5

CONCLUSION

Cloud computing morphs and shifts through countless configurations and designs while retaining certain core characteristics and risks. This flexibility, coupled with seemingly endless elasticity, has enticed governments and firms to make obsolete many on-premises data centers. Among an ever-larger user base for cloud services, one can find a growing number of CI owners and operators. Companies from the healthcare, transportation and logistics, energy, defense, and financial services sectors are moving some of their core functions to the cloud. Cloud services come with considerable benefits but obscure much of their complexity from users and concomitantly imperil the "shared responsibility" model putatively at the heart of secure cloud adoption.

The results are shifts in the risk landscape for data systems that policymakers have yet to fully reckon with. As critical sectors grow more dependent on cloud computing overall, they also increasingly use similar data-storage systems, job schedulers, and orchestrators to serve their data systems. This consolidation—often cited when deciding to transition—does deliver security and innovation benefits.[147] However, the discussion of these benefits has dominated the mainstream at the expense of a more thorough accounting of the many ways that centralization and standardization could impact the resilience of the entire ecosystem.

This report seeks to bring attention to the increasing reliance of CI on cloud services and to the unique dynamics of cloud risk that policymakers must grapple with in thinking through managing the new forms of risk created by this transition. As more entities adopt the cloud, and as more of the core infrastructure of systems like the Internet come to rely on it, this dependence and the systemic nature of its attendant risks will only compound. Risk management must have visibility. The thrust, therefore, of this report's recommendations is towards increased fact-finding and awareness as a key first step for policy.

It is time to address the fact that the cloud may have already become critical by the metrics policymakers use when considering whether a system needs oversight to ensure its resilience. Thankfully, it is not too late for judicious policymaking to earnestly engage with cloud risks to create a more robust regulatory framework suited for the cloud's present and increasing role as a linchpin of critical national functions.

[147] Daniel Geer, et al., "Cyber/nsecurity: The Cost of Monopoly – How the Dominance of Microsoft's Products Poses a Risk to Security," Cryptome, September 27, 2003, https://cryptome.org/cyberinsecurity.htm.

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.

1030 15th Street, NW, 12th Floor,
Washington, DC 20005
(202) 778-4952
www.AtlanticCouncil.org