

SCHOLARLY ARTICLE

With Great (Computing) Power Comes Great (Human Rights) Responsibility: Cloud Computing and Human Rights

Vivek Krishnamurthy* 

Samuelson–Glushko Professor of Law and Director of the Samuelson–Glushko Canadian Internet Policy and Public Interest Clinic (CIPPIC), University of Ottawa, Canada; Fellow, Carr Center for Human Rights Policy, Harvard Kennedy School, USA.

*Corresponding author. Email: Vivek.Krishnamurthy@uottawa.ca

Abstract

Cloud computing underlies many of the most powerful and controversial technologies of our day – from large-scale data processing to face recognition systems – yet the human rights impacts of cloud computing have received little scrutiny. This article explains what cloud computing is, how it works, and what some of its most significant human rights impacts are through three case studies. It offers some initial thoughts on the human rights responsibilities of cloud computing providers under the United Nations Guiding Principles on Business and Human Rights and concludes by suggesting directions for future research.

1. Introduction

In December 2020, Google announced it would be opening a new cloud computing data centre in Saudi Arabia to serve the needs of customers ranging from Saudi Aramco – the world’s largest oil producer – to Snapchat, the most popular social network in the Middle East.¹ Google’s announcement raised serious concerns in the human rights community about building and operating such powerful data processing and storage capabilities in a country with a deeply troubling human rights record.²

The ensuing controversy cast some much-needed light on the human rights implications of cloud computing. This family of technologies, which allows users to access incredibly powerful computing and data storage capabilities on a pay-as-you-go basis, garnered considerable media and popular attention about a decade ago – before big data, artificial intelligence (AI), and the Internet of Things emerged as the next big things in technology. Cloud computing capabilities underlie all of these technologies, yet the human rights impacts of cloud computing have not been studied to nearly the same extent as newer technologies that are powered by the cloud.

¹ Dave Stiver, ‘Google Cloud Announces New Regions’, *Google Cloud Blog* (21 December 2020), <https://cloud.google.com/blog/products/infrastructure/google-cloud-announces-new-regions/> (accessed 21 January 2022).

² Access Now, ‘No to the New Google Cloud in Saudi Arabia: Access Now, CIPPIC Flag Serious Human Rights Concerns’ (3 February 2021), <https://www.accessnow.org/google-cloud-in-saudi-arabia-human-rights-concerns/> (accessed 21 January 2022).

This article hopes to close this gap by exploring some of the human rights impacts of cloud computing. The cheap and easy availability of enormous amounts of computing power has been vital to the development of transformative technologies like AI. Yet the availability of so much power also poses unique human rights challenges, as it allows cloud computing customers with risky business models to scale up rapidly and cause adverse human rights impacts that are disproportionate to their revenue and headcount. Furthermore, the technical need for cloud computing capabilities to be located in close physical proximity to one's customers is contributing to jurisdictional risks, while reinforcing concerning trends towards data localization.

This article will begin by explaining how cloud computing works, prior to examining the technology's human rights impacts through three case studies. It will then offer some preliminary thoughts on how cloud computing providers should meet their responsibility to respect human rights, as recognized by the United Nations Guiding Principles on Business and Human Rights (UNGPs).³

The analysis below will focus on two forms of cloud computing known as Platform as a Service (PaaS) and Infrastructure as a Service (IaaS), rather than on a third form known as Software as a Service (SaaS). As will be explained in Part I, much has been written about the human rights implications of SaaS technologies, but very little work has been done on PaaS or IaaS. This lack of attention is surprising, as PaaS and IaaS services are major profit centres for technology giants such as Amazon, Microsoft and Google, who are among the leading global providers of these technologies. My hope is that this article will begin to correct this imbalance.

II. An Introduction to Cloud Computing

Computer scientists define cloud computing as a model of computing that 'enables ubiquitous, convenient, on-demand access to a shared pool of configurable computing resources that can rapidly be provisioned at any time and from any location via the Internet or a network.'⁴ This section will seek to explain cloud computing in non-technical terms and lay the groundwork for the human rights-centric discussion that follows in Parts II and III.

Historical Development of Cloud Computing

Cloud computing is a cutting-edge technology, but in many ways it is an old idea. In the early days of computing, processing and storage operations occurred on enormous mainframe computers that were connected to 'terminals' from which the mainframe could be accessed and controlled.⁵ Such terminals were known as 'dumb' terminals because they did not have any significant processing or storage capabilities; rather, they simply provided portals by which to access a mainframe.⁶

Technology took a different direction starting in the 1980s with the rise of the personal computer (PC). PCs incorporated powerful processing and storage capabilities for their time,

³ Human Rights Council, 'Guiding Principles on Business and Human Rights: Implementing the United Nations "Protect, Respect and Remedy" Framework', A/HRC/17/31 (21 March 2011) Principle 11 (UNGPs).

⁴ Peter Mell and Timothy Grance, 'The NIST Definition of Cloud Computing', National Institute of Standards and Technology (September 2011) 2, <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf> (accessed 24 January 2022).

⁵ Benj Edwards, 'The Forgotten World of Dumb Terminals', *PC Magazine* (13 October 2016), <https://www.pcmag.com/news/the-forgotten-world-of-dumb-terminals> (accessed 24 January 2022).

⁶ *Ibid.*

but due to the rudimentary networking technology of the era, PCs mostly operated as standalone machines.

The rise of the internet and the development of broadband networks in the first decade of this century changed the direction of technology again. Faster internet connections allowed data to be more easily stored or archived on distant servers.⁷ The shift towards mobile computing heralded by the launch of the iPhone in 2007 created further incentives to move computing processes off the devices we own, and into centralized data centres.

The smartphone revolution reinforced a trend among large enterprises of moving away from maintaining their own on-site computer servers, and moving them into dedicated data centres operated by someone else for cost and efficiency reasons.⁸ It was, however, the development of a technology known as *virtualization* that led enterprises to get out of the business of operating their own computing infrastructures, and purchase computing and storage capabilities from third parties as they needed them.

Virtualization is a remarkable technology that allows a single physical computer to run multiple ‘virtual machines’ on its hardware.⁹ Some readers may be familiar with software that allows users to run one operating system – say Microsoft Windows – on a different computing platform (like the Apple Macintosh) as if the former were an ordinary application on the latter.¹⁰ Likewise, virtualization software allows many different ‘virtual machines’ to operate on one or more interconnected servers within a data centre. This permits many different customers to share the same physical servers to perform their computing and storage tasks without compromising data security.¹¹

Virtualization, along with cheap and reliable broadband internet connections, are the key enabling technologies for cloud computing.¹² Virtualization permits companies like Amazon or Microsoft to build enormous data centres to service the computing and storage needs of a vast array of customers. As a given customer needs more or less computing or storage power, a cloud service provider (CSP) can scale up or down the number of virtual machines it supplies to meet the customer’s demands.¹³ Customers therefore need not maintain enough computing infrastructure to meet their peak demand, thanks to the ‘elastic’ nature of cloud providers’ services.

These trends have resulted in cloud computing growing into an enormous industry. The global cloud computing market was valued at nearly US\$275 billion in 2020 and is expected to grow at an annual rate of 19 per cent to 2028.¹⁴ Amazon Web Services (AWS), an Amazon subsidiary that is the world’s largest cloud computing company, reported US\$16.1 billion in revenues in the third quarter of 2021, and US\$4.88 billion in operating

⁷ Michael Armbrust et al, ‘Above the Clouds: A Berkeley View of Cloud Computing’, EECS Department, University of California, Berkeley (10 February 2009), <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html>, 6–7.

⁸ Ibid.

⁹ IBM, ‘Virtualization’, *IBM Cloud Learn Hub* (30 March 2021), <https://www.ibm.com/cloud/learn/virtualization-a-complete-guide> (accessed 24 January 2022).

¹⁰ See, e.g., Edward Mendelson, ‘VMware Fusion Review’, *PC Magazine* (20 September 2021), <https://www.pcmag.com/reviews/vmware-fusion> (accessed 24 January 2022).

¹¹ IBM, note 9.

¹² Ali Sunyaev, *Internet Computing: Principles of Distributed Systems and Emerging Internet-Based Technologies* (Cham, Switzerland: Springer, 2020) 197.

¹³ Ibid.

¹⁴ Grandview Research, ‘Cloud Computing Market Size & Share Report, 2021–2028’ (July 2021), <https://www.grandviewresearch.com/industry-analysis/cloud-computing-industry> (accessed 24 January 2022).

income. This represents more than 100 per cent of Amazon's operating income for that quarter, as Amazon's other divisions posted losses.¹⁵

An industry of this size, impact and profitability deserves serious study by scholars of all kinds, especially for as we will see below, cloud computing is the key technology underlying the transformative impact of AI on our society.

Taxonomy of Cloud Computing

Technologists have devised a taxonomy of cloud computing that encompasses three different 'service models', each of which is based on 'the abstraction levels of the [cloud computing] capabilities provided'.¹⁶ We can think of these 'abstraction levels' in terms of the level of value-added services that cloud companies provide to their customers, or to the level of technical expertise required to use them.

The three basic cloud computing service models are Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS).¹⁷ The three service models can be conceptualized as a pyramid, with SaaS at the top, IaaS at the bottom, and PaaS in the middle. Much like a pyramid, the upper levels can be 'layered' on top of the lower levels, as will be explained below.¹⁸

Software as a Service

SaaS is the most familiar kind of cloud computing to non-technical audiences. As the name implies, SaaS involves providing software one might run as an application on one's own computing device as a 'service' that is provided remotely instead.¹⁹

Consider a web-based email service such as Gmail, for example. Twenty years ago, individuals would install email software on their PC to download their messages from an email server, from which they would then be deleted due to storage constraints. Today, services such as Gmail and [Outlook.com](https://www.outlook.com) offer email on a host of internet-connected devices through web browsers or via specialized applications. Most users of such services never download their messages onto their own computing device. Rather, messages are stored in the cloud forever, and they are retrieved for display on the browser or in an app upon the user's request.

The range of SaaS services is vast, and includes everything from file storage services such as Dropbox and [Box.com](https://www.box.com), to social media services such as Facebook and YouTube, to enterprise software ranging from Microsoft's Office 365 productivity suite to the customer relationship management applications of [Salesforce.com](https://www.salesforce.com).²⁰

The human rights impacts of SaaS services are well-known and well-studied, at least when compared with the other cloud service models. Much has been written about how SaaS companies should protect the privacy and free expression rights of their users against unlawful government demands, and there are multi-stakeholder

¹⁵ Jordan Novet, 'Amazon Web Services Tops Analysts' Estimates on Profit and Revenue', *CNBC* (28 October 2021), <https://www.cnn.com/2021/10/28/aws-earnings-q3-2021.html> (accessed 24 January 2022).

¹⁶ Sunyaev, note 12, 203.

¹⁷ Mell and Grance, note 4, 2–3.

¹⁸ Each of these three 'service models' can be paired with one of three dominant cloud 'deployment models' (namely public, private and hybrid cloud). The three deployment models are not detailed in this article because they do not contribute significantly to the human rights analysis developed below. See *ibid.*

¹⁹ Sunyaev, note 12, 204–205.

²⁰ *Ibid.*

organizations like the Global Network Initiative that provide guidance and oversight on such topics.²¹

Likewise, there is a robust and ongoing public discussion on how SaaS companies can use the data that their customers entrusted to them for their own commercial purposes, such as the targeting of advertisements or the use of analytics to uncover facets of their customers' behaviour or personality.²² Data protection laws such as the European Union's General Data Protection Regulation (GDPR) already significantly constrain what SaaS companies can do with their users' data in many jurisdictions,²³ and even the United States is considering enacting a comprehensive federal data privacy law of its own.²⁴ Considering the attention that is being paid to the impacts of SaaS, this article will focus its analysis on the IaaS and PaaS service models, which are described below.

Infrastructure as a Service

IaaS forms the base of the cloud computing service model pyramid. This service model involves providers making vast amounts of computing and storage infrastructure available to their customers to configure and use as they please.

The seeds for IaaS were planted in the early 2000s when large enterprises began to decommission their on-site servers for applications ranging from email to document management, in favour of renting server capacity from large commercial providers.²⁵ At this point, specific servers in a data centre could be identified as hosting the data and computing operations of a particular enterprise customer, such as a bank or insurance company.

With the development of virtualization, however, the same physical server infrastructure could be used to service multiple 'tenants' – as cloud computing customers are sometimes called. As explained above, virtualization allows CSPs to provide their customers with processing and storage capabilities on an as-needed, or 'elastic' basis.

Computing power and storage capabilities are two of the three core IaaS services. The third is networking: cloud computing only works when the servers in a provider's data centre are reliably interconnected with each other, and with the customer wherever they might be.²⁶ To that end, many large CSPs maintain their own private network of fibre-optic

²¹ See, e.g., Michael A Samway and Warren Ryan, 'The Internet, Human Rights, and the Private Sector' (2014) 15:1 *Georgetown Journal of International Affairs* 25; Vivek Krishnamurthy, 'Cloudy with a Conflict of Laws', Berkman Center Research Publication (February 2016), <https://ssrn.com/abstract=2733350> (accessed 24 January 2022).

²² See, e.g., Amnesty International, 'Surveillance Giants: How the Business Model of Google and Facebook Threatens Human Rights', <https://www.amnesty.org/en/documents/pol30/1404/2019/en/> (accessed 24 January 2022); Human Rights High Commissioner, 'Addressing Business Model Related Human Rights Risks', https://www.ohchr.org/Documents/Issues/Business/B-Tech/B_Tech_Foundational_Paper.pdf (accessed 24 January 2022); Hannah Darnton, Dunstan Allison-Hope and Gina Lau, 'Responsible Product Use in the SaaS Sector' (15 July 2021), <https://www.bsr.org/reports/BSR-Responsible-Product-Use-SaaS-Sector.pdf> (accessed 24 January 2022).

²³ A Code of Conduct has recently been developed to help facilitate compliance by CSPs with the complex requirements of the GDPR. 'European Union Data Protection Code of Conduct for Cloud Services Providers', October 2020, https://euoc.cloud/fileadmin/cloud-coc/files/former-versions/European_Cloud_Code_of_Conduct_2.10.pdf (accessed 24 January 2022).

²⁴ Jessica Rich, 'After 20 Years of Debate, It's Time for Congress to Finally Pass a Baseline Privacy Law', *Brookings Blog* (14 January 2021), <https://www.brookings.edu/blog/techtank/2021/01/14/after-20-years-of-debate-its-time-for-congress-to-finally-pass-a-baseline-privacy-law/> (accessed 26 January 2022).

²⁵ Rich Miller, 'Which is Cheaper: In-House or Colo?', *Data Center Knowledge* (16 January 2007), <https://www.datacenterknowledge.com/archives/2007/01/16/which-is-cheaper-in-house-or-colo/> (accessed 26 January 2022).

²⁶ Sunyaev, note 12, 203.

cables to interconnect data centres in various parts of the world.²⁷ Such high-quality interconnections between data centres also permit the customer's data to be stored in multiple, redundant locations for disaster recovery purposes.

The IaaS service model provides customers with enormous flexibility in using the massive computing and storage resources that cloud providers make available. The trade-off is that customers must possess the technical expertise to configure the services to meet their needs and decide what applications to install and run in the cloud.²⁸

Platform as a Service

PaaS makes up the middle level of the cloud computing service model pyramid. In this service model, the CSP not only supplies and manages the underlying computing infrastructure, but it also provides advanced software tools that customers can use to build powerful applications tailored to their needs.²⁹

Amazon, Microsoft and Google are the three leading platform providers. Each provides its customers with a PaaS offering that includes dozens of advanced computing capabilities that are designed to be interoperable. These capabilities include standard database software and analytics tools, but also powerful machine learning models for applications ranging from translation to voice recognition to face recognition.³⁰

The power of PaaS is best illustrated through an example. Consider a global retailer that requires its customers to purchase memberships to access its stores. Suppose the membership cards include a digital photo of the member that is stored, along with other customer details, in a PaaS provider's cloud. Using PaaS tools, the company could do away with membership cards and instead identify its customers at the point of entry using face recognition. The retailer could feed its members' photos into the trained face recognition modules provided by leading PaaS providers, and then deploy cameras at its store entrances that could be connected to the cloud to identify customers as they enter.

Now suppose the retailer in question also provides its customers with feedback cards. The company could use tools provided by each of the leading PaaS providers to build a workflow that takes scanned images of the cards, recognizes the handwriting on them, translates the data into a common language, and then perform detailed statistical analysis to identify trends in customer satisfaction.

A large multinational with a sophisticated IT department could purchase its own software for each of these applications, and combine them into a customized workflow to meet its needs. PaaS services make it orders of magnitude easier, however, for someone to combine powerful capabilities such as face recognition and data analytics into a customized workflow – with less time, labour and cost than just a few years ago.

Layering Service Models

The three cloud computing service models are described as a pyramid because one can be 'layered' on top of the other.³¹ For example, all but the largest providers of SaaS services use storage and processing infrastructure provided by IaaS companies to build their services. Consider Dropbox, the leading cloud file storage service, which is best described as a SaaS

²⁷ Dan Dziedzic, 'Google Expands Network with New Data Centers, Subsea Cables', *CNET* (16 January 2018), <https://www.cnet.com/news/google-expands-network-with-new-data-centers-subsea-cables/> (accessed 26 January 2022).

²⁸ Sunyaev, note 12, 204.

²⁹ *Ibid.*

³⁰ Maria Yatsenko, 'AI Platform as a Service: Definition, Architecture, Vendors', *Apriorit* (5 September 2019), <https://www.apriorit.com/dev-blog/635-ai-ai-paas> (accessed 26 January 2022).

³¹ Sunyaev, note 12, 212–214.

service. Until 2016, Dropbox relied on Amazon to provide its underlying computing and storage infrastructure, before it decided to build out its own data centres with its own servers.³² As will be described in Part II, Snapchat – the leading social networking service in much of the Middle East – uses cloud infrastructure operated by Google for certain aspects of its service.

Service Models: A Food Service Analogy

We can use the analogy of a university seeking to hire a contractor to open up a new cafeteria on campus to help clarify the differences between SaaS, PaaS and IaaS.

One option for the university would be to hire a food service company to build and operate a cafeteria. The contractor would decide how to configure the cafeteria, and what food to sell. The university would have no control over what food the contractor sells, although customers could make some customizations to their food – such as applying condiments or asking for extra cheese. This scenario is analogous to the SaaS service model. SaaS customers can choose products from a provider’s menu, but their ability to customize the software is limited to those options made available by the provider.

A second option would be for the university to contract with a company to build a cafeteria and maintain the equipment and appliances. The university would have to hire its own chefs, who would decide what ingredients to order, and what food to cook. The quality of the cafeteria’s output would depend greatly on the skill of the chefs and the quality of the ingredients procured by the university. The scenario is analogous to IaaS, whereby the CSP builds and maintains the computing infrastructure, but the client is responsible for almost everything else.

A third option would be for the university to contract with a company that does everything outlined in the second option, above, but also to supply the cafeteria with high-quality, easy-to-use meal preparation items that complement each other, such as pre-made sauces and pre-sliced vegetables. The university would be responsible for staffing the cafeteria, and the chefs would be responsible for deciding what exactly to cook. Even so, this model allows the university to hire relatively low skilled cooks and operate a cafeteria that produces high-quality meals. This scenario is analogous to PaaS, whereby the cloud provider gives customers powerful, easy-to-use building blocks for customized workflows.

Cloud Computing Contracts

Complex contracts govern the relationship between CSPs and their customers. The contracts are made up of multiple documents, many of which are offered on a ‘take it or leave it’ basis to all but the very largest customers.³³

Each of the three leading CSPs (Amazon, Google and Microsoft) structure their contracts somewhat differently, but they all have key features in common. These include (1) commercial terms, which govern the pricing of the cloud computing services; (2) service level terms, which pertain to the availability and reliability of the services, and liability for service interruptions; and (3) clauses allowing cloud providers to take necessary actions to comply with legal requirements or law enforcement demands. The contracts also impose limitations

³² Cade Metz, ‘The Epic Story of Dropbox’s Exodus from the Amazon Cloud Empire’, *Wired* (14 March 2016), <https://www.wired.com/2016/03/epic-story-dropboxs-exodus-amazon-cloud-empire/> (accessed 26 January 2022).

³³ W Kuan Hon, Christopher Millard and Ian Walden, ‘Negotiating Cloud Contracts: Looking at Clouds from Both Sides Now’ (2012) 16 *Stanford Technology Law Review* 79–128.

on how the services can be used, terms that govern compliance with data protection laws, and service-specific terms governing the use of certain functionalities.

Use Limitations

Each cloud provider's standard-form contracts include provisions prohibiting the use of their services in certain circumstances. For example, Google Cloud Platform's Terms of Service prohibit the use of its services in situations where they 'would reasonably be expected to lead to death, personal injury, or environmental or property damage (such as the creation or operation of nuclear facilities, traffic control, life-support systems, or weaponry)'.³⁴ Interestingly, Amazon's terms include a provision prohibiting the use of any Microsoft software on its cloud infrastructure in 'high risk' applications including civil aviation, nuclear and chemical facilities, life support systems, and so on.³⁵

Amazon and Google's contractual documents also include free-standing Acceptable Use Policies which include additional restrictions on customers' uses of the companies' cloud services.³⁶ These short, terse documents prohibit the use of each company's cloud services in furtherance of illegal activities, or to violate the legal rights of others. Such clauses could be viewed as referencing human rights, but they appear to be directed at the kinds of private legal rights that are readily enforceable in most domestic legal systems – such as intellectual property or contractual rights.³⁷

Data Protection Terms

Each cloud provider's standard-form contracts include provisions regarding compliance with relevant data protection laws. Data protection laws are an important means of protecting the right to privacy, hence these provisions have significant human rights implications.³⁸ Most providers' contracts assign the primary legal responsibility for complying with relevant data protection laws to the cloud computing customer.³⁹ Even so, the provisions give CSPs the right not to carry out data processing tasks they believe contravene applicable data protection laws.⁴⁰

Service-Specific Terms

Standard-form cloud computing contracts also include special terms that govern the use of specific services. Providers such as Microsoft, Google and Amazon offer their cloud

³⁴ Google Cloud, 'Google Cloud Platform Terms of Service', <https://cloud.google.com/terms> ss 3.3, 14.19 (accessed 26 January 2022).

³⁵ Amazon Web Services, 'AWS Service Terms', <https://aws.amazon.com/service-terms/> s 5.1.1 (accessed 26 January 2022) (AWS Service Terms).

³⁶ Amazon Web Services, 'AWS Acceptable Use Policy', <https://aws.amazon.com/aup/> (accessed 26 January 2022). Google Cloud, 'Google Cloud Platform Acceptable Use Policy', <https://cloud.google.com/terms/aup> (accessed 26 January 2022). For its part, Microsoft includes similar provisions within the high-level Microsoft Services Agreement. Microsoft, 'Microsoft Services Agreement', <https://www.microsoft.com/en-ca/servicesagreement/> s. 3 (accessed 26 January 2022).

³⁷ For example, Google Cloud's Acceptable Use Policy prohibits customers from using its services 'to violate, or encourage the violation of, the legal rights of others (for example, this may include allowing Customer End Users to infringe or misappropriate the intellectual property rights of others in violation of the Digital Millennium Copyright Act) ...'.

³⁸ J Kokott and C Sobotta, 'The Distinction between Privacy and Data Protection in the Jurisprudence of the CJEU and the ECtHR' (2013) 3:4 *International Data Privacy Law* 222–28.

³⁹ AWS Service Terms, note 35, ss 1.11 and 1.14.

⁴⁰ See, e.g., Google Cloud, 'Data Processing and Security Terms (Customers)', <https://cloud.google.com/terms/data-processing-terms> s 5.2.3 (accessed 26 January 2022).

customers dozens of different tools as part of an integrated service.⁴¹ While the standard terms are sufficient to govern most aspects of most services, special clauses are required to govern how some services are used. For example, Amazon's contracts include service-specific terms governing Rekognition, the company's controversial face recognition system. These terms prohibit Rekognition's use in certain law enforcement contexts pursuant to a moratorium announced by Amazon,⁴² and also mandate appropriate training for law enforcement officials using Rekognition 'in making decisions that could impact civil liberties or equivalent human rights ...'.⁴³ The possibility of developing service-specific terms to address human rights concerns with particular cloud computing functionalities is explored in Part III.

III. Human Rights Impacts of Cloud Computing: Three Case Studies

This section describes three case studies illustrating the range of human rights impacts related to the provision and use of PaaS and IaaS cloud computing services. Three key take-home messages will emerge from this discussion:

1. The PaaS service model makes it easy for small organizations and even individuals to build powerful applications that would otherwise be beyond their capabilities, and to operate these applications at enormous scale;
2. The IaaS service model places enormous amounts of raw computing power in the hands of organizations and even individuals who would otherwise face insurmountable barriers to entry to computation on such a scale;
3. The technical need to locate cloud computing infrastructure in close physical proximity to their end users creates significant new forms of geographic and jurisdictional risk, especially in an era where governments are demanding that data be stored within their borders.

These case studies focus on the adverse human rights impacts of cloud computing, both potential and actual. Of course, cloud computing has many pro-social uses – and on balance, the technology probably represents a net positive for the full range of human rights.⁴⁴ The focus of these case studies is nevertheless justified by the need to illuminate what CSPs should be doing to meet their responsibility to respect human rights.

Telephone Surveillance in American Prisons

The first case study illustrates how PaaS technologies can be used by small entities to build powerful workflows and applications that can have deleterious human rights impacts, and to operate these at massive scale. In November 2021, investigative journalists with the Thomson Reuters Foundation exposed how several companies – including California-based

⁴¹ For example, Google offers more than 150 distinct services as part of its Google Cloud family of cloud computing services. The distinct services are helpfully listed in a document on the Google Cloud website. Google Cloud, 'Google Cloud Platform Services Summary', <https://cloud.google.com/terms/services> (accessed 26 January 2022).

⁴² AWS Service Terms, note 35, s 50.9.

⁴³ *Ibid*, s 50.8.2.

⁴⁴ Pursuant to the UNGPs, the positive human rights impacts of a business activity cannot be used to 'offset' its negative impacts; the latter must be remediated when they occur. See UNGPs, note 3, Principle 11.

LEO Technologies – were leveraging Amazon’s PaaS offerings to build and deploy powerful tools to surveil the telephone conversations of inmates at several American prisons.⁴⁵

The Verus system built by LEO Technologies appears to combine several different AWS PaaS tools to store and analyse vast quantities of inmate telephone conversations. At just one American county jail system, Verus is used to analyse more than 10,000 hours of inmate conversations every month.⁴⁶ Specifically, once the recordings of prisoners’ conversations are uploaded into the Amazon Cloud, the Verus system uses AWS’s Transcribe service to generate transcripts of the inmate conversations.⁴⁷

The Thomson Reuters story reports that prison officials perform keyword searches on the transcripts to identify conversations of interest to them. Some of these searches may well be connected to legitimate penological purposes, such as identifying and disrupting gang activity within the prison. Other uses are dubious, such as the use of the system by an Alabama prison embroiled in litigation over incarceration conditions to find evidence useful to its defence.⁴⁸

As of 19 January 2022, the LEO Technologies website reports that its Verus system has been used to transcribe over 107 million inmate conversations, representing nearly 564 years of audio.⁴⁹ The product page for Verus strongly suggests the company leverages some of AWS’s machine learning tools (another PaaS service), inasmuch as the application is capable of ‘proactively flagging’ inmate calls for review by the prison authorities based on ‘keywords and phrases’.⁵⁰ This can be surmised from the AWS website for its Transcribe PaaS transcription service, which includes information about a related Transcribe Call Analytics service, which has the capability to ‘analyse call recordings ... [for] actionable insights’.⁵¹

This case study arises in the prison context, where the law recognizes that inmates have a reduced expectation of privacy in their communications.⁵² Even so, the capabilities provided by Verus represent a quantum leap in the ability of the prison authorities to surveil the telephone conversations of their inmates. What LEO Technologies has managed to build for its prison customers using PaaS services would have required the capabilities of a sophisticated intelligence service just a few years ago.⁵³

Furthermore, the underlying technologies explored in this case study could very easily be deployed in other contexts where they are much more concerning. It is conceivable someone could use the same AWS PaaS modules to build something similar to Verus, but deploy it in an environment that raises even more significant privacy concerns. The adverse human rights impacts of deploying Verus in the prisons of a country with a poor human rights record are not hard to imagine.

⁴⁵ Avi Asher-Schapiro and David Sherfinski, ‘“Scary and Chilling”: AI Surveillance Takes U.S. Prisons by Storm’, *Thomson Reuters Foundation News* (16 November 2021), <https://news.trust.org/item/20211115095808-kq7gx/> (accessed 24 January 2022).

⁴⁶ *Ibid.*

⁴⁷ *Ibid.* Amazon Web Services, ‘Amazon Transcribe – Speech to Text – AWS’, <https://aws.amazon.com/transcribe/> (accessed 21 January 2022).

⁴⁸ Asher-Schapiro and Sherfinski, *note 45*.

⁴⁹ LEO Technologies, ‘Home Page’, <https://leotechnologies.com/> (accessed 21 January 2022).

⁵⁰ LEO Technologies, ‘Verus’, <https://leotechnologies.com/services/verus/> (accessed 21 January 2022).

⁵¹ Amazon Web Services, ‘Amazon Transcribe Call Analytics | Transcripts & Insights | AWS’, <https://aws.amazon.com/transcribe/call-analytics/> (accessed 21 January 2022).

⁵² See, e.g., Sydney Ingel et al., ‘Privacy Violations and Procedural Justice in the United States Prisons and Jails’ (2020) 15:2 *Sociology Compass* 1.

⁵³ See, e.g., Eyder Peralta, ‘Report: NSA Can Record, Store Phone Conversations Of Whole Countries’, *NPR* (18 March 2014), <https://www.npr.org/sections/thetwo-way/2014/03/18/291165247/report-nsa-can-record-store-phone-conversations-of-whole-countries> (accessed 21 January 2022).

Amazon has not responded to the Thomson Reuters story, but LEO Technologies remains one of just 35 companies worldwide to be selected and endorsed by Amazon for their ‘Public Safety and Disaster Response Competency’.⁵⁴ This raises questions of the nature of Amazon’s responsibility under the UNGPs for the potential and actual adverse human rights impacts arising from the use of its PaaS services by companies such as LEO Technologies.

IaaS, Machine Learning and Face Recognition

The second case study illustrates how IaaS capabilities can be misused by individuals and small-scale organizations to engage in conduct that can have large-scale human rights implications.

Face recognition (FR) is a controversial technology whose application has been the subject of considerable human rights scrutiny in recent years.⁵⁵ Most of this scrutiny has been focused on either the accuracy of FR – especially when it is applied to women and members of minority communities,⁵⁶ or on the myriad privacy implications of the use of this technology – especially in public places.⁵⁷ Efforts are under way in many jurisdictions to regulate the use of FR.⁵⁸ For example, Washington State has enacted comprehensive legislation to govern the use of FR by state government entities (including law enforcement agencies),⁵⁹ while the European Parliament has called for a ban on law enforcement use of FR in public places.⁶⁰

Each of the three leading PaaS providers (Google, Microsoft and Amazon) provides FR capabilities as part of their cloud service offerings.⁶¹ Following the Black Lives Matter protests in 2020, several large technology companies imposed moratoria on the sale and use of FR by law enforcement agencies,⁶² while others (notably IBM) have abandoned their FR research altogether.⁶³

⁵⁴ Amazon Web Services, ‘AWS Public Safety and Disaster Response Competency’, <https://aws.amazon.com/stateandlocal/justice-and-public-safety/partner-solutions/> (accessed 21 January 2022).

⁵⁵ See, e.g., Clare Garvie et al, ‘The Perpetual Line-up: Unregulated Police Face Recognition in America’, *Georgetown Law Center on Privacy & Technology* (28 October 2016), <https://www.perpetuallineup.org/sites/default/files/2016-12/The%20Perpetual%20Line-Up%20-%20Center%20on%20Privacy%20and%20Technology%20at%20Georgetown%20Law%20-%2020121616.pdf> (accessed 21 January 2022). Davide Castelvecchi, ‘Is Facial Recognition Too Biased to Be Let Loose?’ (2020) 587:7834 *Nature* 347–350.

⁵⁶ See, e.g., Joy Buolamwini and Timnit Gebru, ‘Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification’, paper presented at conference on ‘Fairness, Accountability and Transparency’ on 1 January 2018, <https://proceedings.mlr.press/v81/buolamwini18a.html> (accessed 22 January 2022).

⁵⁷ Tom Simonite, ‘How Face Recognition Can Destroy Anonymity’, *Wired* (20 April 2021), <https://www.wired.com/story/how-face-recognition-destroy-anonymity/> (accessed 21 January 2022).

⁵⁸ Rashida Richardson, ‘Facial Recognition in the Public Sector: The Policy Landscape’, *German Marshall Fund of the United States* (1 February 2021), <https://www.gmfus.org/sites/default/files/Richardson%20-%20Facial%20recognition.pdf> (accessed 22 January 2022).

⁵⁹ Alexander Berengault and Jadzia Pierce, ‘Washington State Passes Bill Limiting Government Use of Facial Recognition’, *Inside Privacy Blog* (23 March 2020), <https://www.insideprivacy.com/united-states/state-legislatures/washington-state-passes-bill-limiting-government-use-of-facial-recognition/> (accessed 22 January 2022).

⁶⁰ Eileen Li, ‘Europe’s Next Steps in Regulating Facial Recognition Technology’, *Columbia Journal of Transnational Law Bulletin* (7 November 2021), <https://www.jtl.columbia.edu/bulletin-blog/europes-next-steps-in-regulating-facial-recognition-technology> (accessed 22 January 2022).

⁶¹ See, e.g., Makena Kelly, ‘Big Tech Faces New Pressure over Facial Recognition Contracts’, *The Verge* (15 January 2019), <https://www.theverge.com/2019/1/15/18183789/google-amazon-microsoft-pressure-facial-recognition-jedi-pentagon-defense-government> (accessed 21 January 2022).

⁶² Lauren Feiner and Annie Palmer, ‘Rules Around Facial Recognition and Policing Remain Blurry’, *CNBC* (12 June 2021), <https://www.cnb.com/2021/06/12/a-year-later-tech-companies-calls-to-regulate-facial-recognition-met-with-little-progress.html> (accessed 21 January 2022).

⁶³ *Ibid.*

What has escaped notice in the public debate over FR, however, is the role IaaS providers play in enabling such systems. Without the massive amounts of computing and storage infrastructure IaaS providers make available, small-scale entities simply would not be able to develop FR technology or deploy them on a large scale.

The controversy surrounding Clearview AI (Clearview) is a case in point. Clearview is a small, privately held FR company based in Manhattan that provides advanced FR capabilities to law enforcement agencies and other organizations in countries around the world.⁶⁴ The company is under investigation by numerous privacy regulators for the manner in which it developed its FR system, and for its sales practices.⁶⁵ Specifically, Clearview used automated tools to download ('scrape') more than 10 billion images of faces from social media to 'train' its FR software to identify individuals from photographs.⁶⁶ Once these images have been processed by Clearview, they can be used to identify unknown individuals from photographs fed into the system.

Clearview's data gathering practices violate the privacy and data protection laws of many countries,⁶⁷ as well as the terms of service which govern the use of social media websites.⁶⁸ We do not know if Clearview developed its FR technologies in-house, or whether it used FR capabilities developed by other companies (including the three leading PaaS companies) to power its service. There is evidence that Clearview has used AWS for some storage needs, but little else is known about Clearview's business relationship with these PaaS and IaaS providers.⁶⁹

We can surmise that Clearview had access to IaaS services to develop and deploy its FR system, however. Using machine learning techniques for tasks such as FR, voice recognition or automatic translation requires enormous amounts of computing power. For example, Google's AlphaGoZero system for playing Go (an East Asian board game) required 1800 petaflop-days of computing power to train its machine learning algorithms.⁷⁰ This is an amount of computing power equivalent to the world's most powerful supercomputer, which features more than seven million processors, running at full tilt for the better part of five days.⁷¹

⁶⁴ Kashmir Hill, 'The Secretive Company That Might End Privacy as We Know It', *The New York Times* (18 January 2020), <https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html> (accessed 21 January 2022).

⁶⁵ Privacy International, 'The ICO's Announcement about Clearview AI Is a Lot More than Just a £17 Million Fine' (6 December 2021), <http://www.privacyinternational.org/news-analysis/4714/icos-announcement-about-clearview-ai-lot-more-just-ps17-million-fine> (accessed 21 January 2022).

⁶⁶ James Vincent, 'Clearview AI Ordered to Delete All Facial Recognition Data Belonging to Australians', *The Verge* (3 November 2021), <https://www.theverge.com/2021/11/3/22761001/clearview-ai-facial-recognition-australia-breach-data-delete> (accessed 21 January 2022).

⁶⁷ Privacy International, 'The ICO's Announcement about Clearview AI Is a Lot More than Just a £17 Million Fine' (6 December 2021), <http://www.privacyinternational.org/news-analysis/4714/icos-announcement-about-clearview-ai-lot-more-just-ps17-million-fine> (accessed 21 January 2022). See also Vincent, note 66.

⁶⁸ Jon Porter, 'Facebook and LinkedIn Are Latest to Demand Clearview Stop Scraping Images for Facial Recognition Tech', *The Verge* (6 February 2020), <https://www.theverge.com/2020/2/6/21126063/facebook-clearview-ai-image-scraping-facial-recognition-database-terms-of-service-twitter-youtube> (accessed 21 January 2022).

⁶⁹ Dell Cameron, Dhurv Malhotra and Shoshana Wodinsky, 'We Found Clearview AI's Shady Face Recognition App', *Gizmodo* (27 February 2020), <https://gizmodo.com/we-found-clearview-ais-shady-face-recognition-app-1841961772> (accessed 20 January 2022).

⁷⁰ Dario Amodei and Danny Hernandez, 'AI and Compute', *OpenAI* (16 May 2018), <https://openai.com/blog/ai-and-compute/> (accessed 20 January 2022).

⁷¹ Author calculations based on supercomputing performance data obtained from TOP500, 'November 2021 | TOP500', <https://www.top500.org/lists/top500/2021/11/> (accessed 21 January 2022). Running the same calculations on Amazon's most powerful cloud computing array would require nearly 200 days.

We do not know how much computing power Clearview required to process 10 billion images, but we can assume it is within an order of magnitude of AlphaGoZero.⁷² Only the largest and most well-resourced institutions, such as governments, large multinationals and research universities, have the financial wherewithal to purchase components and build large-scale computers with this level of processing power. The only option available to everyone else is to purchase such power from an IaaS provider.

Correspondingly, when individuals or smaller-scale organizations are developing and deploying FR, or other technologies incorporating AI, it is very likely that they are relying on the capabilities of IaaS providers to do so. Given the wide range of human rights impacts that such technologies can have, the question arises of what responsibility IaaS providers bear for the adverse human rights impacts of their customers' uses of the capabilities they provide.⁷³

Google Cloud in Saudi Arabia

The third case study explores the human rights risks arising from building and operating cloud computing infrastructure in jurisdictions around the world. It does so by examining the implications of a decision by Google to build a cloud computing data centre in Saudi Arabia.⁷⁴

Like all of the other major PaaS and IaaS providers, Google operates a global network of cloud computing data centres, with facilities in 29 locations in 23 countries on five continents.⁷⁵ Growing demand for cloud computing services of all kinds, combined with fierce commercial competition among CSPs, is leading all of them to build more data centres in more locations around the world. Rather than simply expanding existing data centres to meet growing needs, however, the cloud computing industry is building more data centres in more locations around the world. There are at least three reasons why this is so.

The first is *latency*, a term which refers to the time it takes for information to travel through a network from a sender to a recipient.⁷⁶ Although the speed and bandwidth of internet connections continue to improve every year, latency remains a major challenge when moving large volumes of data across the internet. At the consumer scale, readers might have experienced unstable connections when using video-conferencing software with someone far away – especially when one party is located in a place with relatively poor communications infrastructure. At the enterprise scale, when terabytes of data are being exchanged with CSPs, there are enormous performance benefits to be accrued from having the relevant cloud provider's infrastructure being located as close as possible to the customer.⁷⁷ This in turn is leading to more cloud infrastructure being built in a greater variety of locations around the world.

⁷² Author correspondence with Dr Jason Millar, Canada Research Chair in the Ethical Engineering of Robotics and Artificial Intelligence, Faculty of Engineering, University of Ottawa, 4 December 2021.

⁷³ See, e.g., Filippo Raso et al, 'Artificial Intelligence & Human Rights: Opportunities & Risks', Berkman Klein Center for Internet & Society at Harvard University (25 October 2018), <https://ssrn.com/abstract=3259344> (accessed 21 January 2022).

⁷⁴ Stiver, note 1.

⁷⁵ Google Cloud, 'Global Locations – Regions & Zones', <https://cloud.google.com/about/locations> (accessed 21 January 2022).

⁷⁶ Cloudflare, 'What Is Latency? | How to Fix Latency', <https://www.cloudflare.com/learning/performance/glossary/what-is-latency/> (accessed 21 January 2022).

⁷⁷ Latency is such a problem in transferring very large volumes of data over the internet that Amazon has developed a special shipping container filled with high-performance servers to help large enterprises transfer their data on the Amazon Cloud. Amazon Web Services, 'AWS Snowmobile | Exabyte-Scale Data Transfer | Amazon Web Services', <https://aws.amazon.com/snowmobile/> (accessed 21 January 2022).

A second factor driving the growing geographic diversification of cloud computing infrastructure is resilience.⁷⁸ Data held by CSPs on behalf of their customers is almost always stored in multiple locations to provide continuity of access and service in the event of short-term failures (due to power interruptions, for example) or longer-term outages caused by natural disasters and the like.

A third factor driving this trend is the growing number of data localization laws and regulations being enacted by governments around the world.⁷⁹ Such laws require certain kinds of data held by certain entities to be stored within the jurisdiction. For example, the Canadian province of British Columbia requires all public-sector entities to store personal information on servers located in Canada,⁸⁰ while India requires all companies doing business in India to store their financial records on computer systems located within the country.⁸¹

Different rationales underlie different data localization laws. Some countries have enacted data localization laws to protect against foreign surveillance,⁸² while others demand that sensitive personal data be stored on computer equipment within their borders to ensure compliance with privacy and data protection laws.⁸³ Some jurisdictions are motivated by economic considerations, such as incentivizing the construction of new data centres within their borders and stimulating the growth of data-intensive industries, such as AI research.⁸⁴

A more sinister motivation, however, is to make data of interest to law enforcement and intelligence agencies more easily accessible than it would be were it stored abroad.⁸⁵ Serious concerns arise when authoritarian governments impose data localization requirements on companies handling various forms of sensitive personal data – from cloud-based email providers to social media network operators – that are susceptible to abuse by the country's security apparatus.

These concerns were front and centre when Google announced in December 2020 that it would be building a new cloud data centre in Saudi Arabia.⁸⁶ A Google blog post suggested that Snapchat – the most popular social networking application in the Middle East⁸⁷ – would be among the tenants of Google's new Saudi data centre.⁸⁸

The idea that a social networking provider like Snapchat – which stores and handles sensitive personal information on behalf of its users – would use infrastructure located in Saudi Arabia raised alarm bells in the human rights community,⁸⁹ given the country's poor

⁷⁸ Google Cloud, 'Architecting Disaster Recovery for Cloud Infrastructure Outages | Cloud Architecture Center', <https://cloud.google.com/architecture/disaster-recovery> (21 January 2022).

⁷⁹ Anupam Chander and Uyên P Lê, 'Data Nationalism' (2015) 64:3 *Emory Law Journal* 677.

⁸⁰ Kirsten Thompson, 'British Columbia Modifies Data Residency Requirements in Response to COVID-19', *Dentons Data Blog* (4 April 2020), <https://www.dentonsdata.com/british-columbia-modifies-data-residency-requirements-in-response-to-covid-19/> (accessed 21 January 2022).

⁸¹ Anirudh Sharma and Upasana Burman, 'How Would Data Localization Benefit India?', Carnegie Endowment for International Peace Working Papers 04/2021, <https://carnegieindia.org/2021/04/14/how-would-data-localization-benefit-india-pub-84291> (accessed 21 January 2022).

⁸² Chander and Lê, note 79, 714.

⁸³ *Ibid.*, 718–719.

⁸⁴ *Ibid.*, 721–723.

⁸⁵ *Ibid.*, 730.

⁸⁶ Stiver, note 1.

⁸⁷ Damian Radcliffe, 'Snapchat's Middle East Success Story', *Damian Radcliffe Blog* (6 August 2021), <https://medium.com/damian-radcliffe/snapchats-middle-east-success-story-49b932b67357> (accessed 21 January 2022).

⁸⁸ Stiver, note 1.

⁸⁹ Access Now, note 2.

human rights record.⁹⁰ The continuing furore over Saudi Arabia's use of cyber-espionage software in its plot to assassinate the journalist Jamal Khashoggi exemplifies the concerns associated with locating social networking data on servers located in Saudi Arabia – and therefore within the easy reach of the Saudi government.⁹¹

In response to open letters issued by two digital rights organizations,⁹² Snapchat's parent company explained that its use of Google's Saudi cloud infrastructure was limited to 'public content – such as our Discover content from major news organizations, online publications and public influencers – so that it can be more efficiently and quickly served to the large numbers of people who view it'.⁹³ Snapchat's clarifications alleviated some concerns, but there remain many others with Google continuing to operate cloud infrastructure in a country with an abysmal human rights record.⁹⁴

Google later issued a statement explaining that it engages in human rights due diligence before building cloud computing infrastructure in new jurisdictions.⁹⁵ Yet the result of these processes did not preclude the construction of cloud computing facilities in a country with a troubling human rights record.

Two implications follow. First, we cannot know what factors Google considered in deciding to build a cloud data centre in Saudi Arabia, but a new Saudi law requiring companies to store the personal data of Saudis within the Kingdom's borders might have influenced the decision.⁹⁶ This points to the likely importance of data localization laws leading to CSPs expanding their physical footprint in jurisdictions that may present significant human rights risks.

Second, we might never have known Snapchat was using infrastructure in Saudi Arabia were it not for Google's press release touting the fact. It is technically difficult for outside observers to determine if an entity is a customer of a given CSP, or if it is using a particular data centre – especially when IaaS or PaaS services are involved.⁹⁷ This has ramifications for the human rights responsibilities of CSPs relating to adverse impacts caused by their customers, as discussed below.

IV. The Human Rights Responsibilities of Cloud Service Providers

The powerful capabilities CSPs place in the hands of their customers give rise to questions about their responsibility to identify, prevent and mitigate adverse human rights impacts connected to them. This section will provide some initial thoughts on how CSPs should

⁹⁰ Saudi Arabia received a score of just 7/100 in Freedom House's 2021 Freedom in the World report. Freedom House, 'Saudi Arabia: Freedom in the World 2021 Country Report', <https://freedomhouse.org/country/saudi-arabia/freedom-world/2021> (accessed 21 January 2022).

⁹¹ Freedom House, 'Saudi Arabia: Freedom on the Net 2021 Country Report', <https://freedomhouse.org/country/saudi-arabia/freedom-net/2021> (accessed 21 January 2022).

⁹² Access Now, note 2.

⁹³ Access Now, 'Letter from Jennifer Park Stout, VP Global Public Policy, Snap Inc., to Peter Micek, General Counsel, Access Now, and Vivek Krishnamurthy, Director of CIPPIC' (2 February 2021), https://www.accessnow.org/cms/assets/uploads/2021/02/Response-to-Access-Now_CIPPIC-letter.pdf (accessed 21 January 2022).

⁹⁴ Human Rights Watch, 'Saudi Arabia: Google Should Halt "Cloud Region"', <https://www.hrw.org/news/2021/05/26/saudi-arabia-google-should-halt-cloud-region> (accessed 25 January 2022).

⁹⁵ Access Now, 'Letter from Pablo Chavez, Vice President, Global Government Affairs and Public Policy, Google Cloud, to Peter Micek, General Counsel, Access Now, and Vivek Krishnamurthy, Director, CIPPIC' (12 February 2021), <https://www.accessnow.org/cms/assets/uploads/2021/02/Google-Cloud-Response-to-Access-Now-and-CIPPIC.pdf> (accessed 21 January 2022).

⁹⁶ Nick Simpson, Simon Topping and John Balouziyeh, 'New Personal Data Protection Law in Saudi Arabia', *Dentons Blog* (14 October 2021), <https://www.dentons.com/en/insights/articles/2021/october/14/new-personal-data-protection-law-in-saudi-arabia> (accessed 21 January 2022).

⁹⁷ Author correspondence with computer security researcher (name withheld), 28 January 2021.

prioritize their human rights due diligence efforts, and on what steps such companies should take when their services are implicated in adverse human rights impacts. This analysis is grounded in the UNGPs, which recognize the normative responsibility of businesses to respect human rights, parallel to the legal duty of states to protect against human rights abuse by state and non-state actors.⁹⁸

Prior to doing so, this section will first revisit some of the case studies presented in Part II to determine whether the CSPs implicated in the cases are best viewed as causing, contributing, or directly linked to adverse human rights impacts.

Causation, Contribution or Direct Linkage?

What companies must do to prevent, identify and respond to adverse human rights impacts under the UNGPs depends on whether they are causing or contributing to such impacts, or whether they are merely directly linked to them through their business relationships. Specifically, UNGP 13 states that business respect for human rights requires companies to:

- (a) Avoid causing or contributing to adverse human rights impacts through their own activities, and address such impacts when they occur;
- (b) Seek to prevent or mitigate adverse human rights impacts that are directly linked to their operations, products or services by their business relationships, even if they have not contributed to those impacts.⁹⁹

There are certainly scenarios where a cloud provider causes adverse human rights impacts ‘through their own activities’ – such as by misusing data that has been entrusted to them in ways that violate privacy laws and norms. The more difficult question is how we conceptualize the connection between CSPs and adverse human rights impacts that are caused by their customers, or their customers’ business partners.

Consider, for example, the relationship between Amazon and the adverse human rights impacts caused by LEO Technologies’ prison telephone surveillance system. Is Amazon *contributing* to these impacts, such that it is responsible for addressing them under UNGP 13(a)? Or is Amazon merely *directly linked* to them, meaning that it should ‘seek to prevent or mitigate’ these adverse impacts as envisioned by UNGP 13(b)?

This question is hard to answer definitively given the lack of agreement on the meaning of causation, contribution and direct linkage as these terms are used in the UNGPs.¹⁰⁰ In his later writings, John Ruggie explained there is a ‘continuum’ between the concepts of contribution and direct linkage, and that:

A variety of factors can determine where on that continuum a particular instance may sit. They include the extent to which a business enabled, encouraged, or motivated human rights harm by another; the extent to which it could or should have known about such harm; and the quality of any mitigating steps it has taken to address it.¹⁰¹

⁹⁸ UNGPs, note 3, Principles 1 and 11.

⁹⁹ UNGPs, note 3, Principle 13.

¹⁰⁰ Debevoise & Plimpton LLP and Enodo Rights, ‘Practical Definitions of Cause, Contribute, and Directly Linked to Inform Business Respect for Human Rights’, 9 February 2017, <https://media.business-humanrights.org/media/documents/files/documents/Debevoise-Enodo-Practical-Meaning-of-Involvement-Draft-2017-02-09.pdf> (accessed 26 January 2022).

¹⁰¹ John G Ruggie, ‘Comments on Thun Group of Banks Discussion Paper on the Implications of UN Guiding Principles 13 & 17 in a Corporate and Investment Banking Context’, 21 February 2017, https://media.business-humanrights.org/media/documents/files/documents/Thun_Final.pdf (accessed 26 January 2022).

A recent report builds on Ruggie's insight by defining 'enabling' as 'contributing to an environment where harm is more likely to occur, e.g., by looking the other way in the face of *known risks* or allowing known or suspected bad actors to continue working unimpeded.'¹⁰² Correspondingly:

The closer the connection between the *company's core business operations, specific products, or specific purchasing activities* and the resulting harm – balanced with other factors – the greater likelihood that the company contributed to the harm, and vice versa.¹⁰³ [emphasis added]

This concept of *specificity* suggests, all things being equal, that PaaS providers are more likely than IaaS providers to be *contributing* to adverse human rights impacts caused by customers, because the former services are more *specific*. Whereas IaaS operators provide large quantities of raw computing power for their customers to use, PaaS providers deliver to their customers a range of highly sophisticated computing capabilities (such as voice and face recognition) that can easily be combined and integrated to execute complex workflows. Correspondingly, PaaS providers should have more reason to know the risks associated with their products and services than IaaS providers.

Human Rights Due Diligence Responsibilities of Cloud Service Providers

The greater specificity of PaaS versus IaaS services impacts what should be required in terms of human rights due diligence (HRDD) from providers of each kind of service. UNGP 17(a) specifies that HRDD should cover situations where an enterprise causes or contributes to adverse human rights impacts through its own business activities, as well as those situations where the enterprise may be 'directly linked' to such impacts through its business relationships. The commentary to UNGP 17 recognizes that businesses should prioritize their HRDD efforts 'in areas where the risk of adverse human rights impacts is most significant, whether due to certain suppliers' or clients' operating context, the particular operations, products or services involved, or other relevant considerations ...'. Prioritization is also warranted in situations where it may be 'unreasonably difficult' for enterprises with large numbers of entities in their value chains to inquire into the human rights impacts of each of them.

The discussion below presents some initial thoughts on how PaaS and IaaS companies should evaluate the human rights-related risks of their services and of providing them to particular customers. Given that all three of the leading CSPs provide PaaS services that are built on top of infrastructure that they also use to offer IaaS services, these companies should be carrying out all forms of HRDD specified below.

Due Diligence by PaaS Providers

The greater specificity of PaaS services suggests that PaaS providers should conduct HRDD that examines the human rights risks relating to the various specific functionalities that they make available to their customers.

¹⁰² Jonathan Drimmer and Peter Nestor, 'Seven Questions to Help Determine When a Company Should Remedy Human Rights Harm under the UNGPs', BSR (January 2021), https://www.bsr.org/reports/Seven_Questions_to_Help_Determine_When_a_Company_Should_Remedy_Human_Rights_Harm_under_the_UNGPs.pdf 5 (accessed 26 January 2022).

¹⁰³ Ibid, 8.

As noted, PaaS providers such as Amazon, Google and Microsoft provide their customers with literally dozens of different capabilities that can be seamlessly integrated to perform incredibly complex computing tasks. Therefore, HRDD by PaaS providers should consider the human rights risks posed by each of these capabilities standing alone, and the additional risks that arise from the ability of customers to combine them to perform functions that would otherwise be well beyond their reach.

Such due diligence should consider whether all functionalities should be made available to any customer that signs up for an account, or whether certain high-risk functionalities (such as face recognition) should be limited to vetted customers located in jurisdictions where the rule of law and respect for human rights are well established. Given the rapid rate of innovation in the cloud computing sector, it is especially important for PaaS providers to conduct ongoing HRDD to identify new sources of human rights risk. Such risks may emerge from the interaction of multiple PaaS functionalities provided by one company, or through interactions between the company's services and other vendors' products.

Google is notable among the leading PaaS companies for publicly disclosing that it has conducted HRDD into specific PaaS capabilities. In 2019, Google engaged a leading consultancy to conduct a human rights assessment of a new face recognition system that would allow its cloud customers to recognize images of 'celebrities' for customers in the media and entertainment sector.¹⁰⁴ The public summary of the assessment notes several human rights risks arising from the proposed product, as well as mitigations for Google to consider implementing. For its part, Microsoft conducted a human rights impact assessment in 2017–2018 into the full range of AI technologies that the company offers.¹⁰⁵ No information was found during this project, however, to confirm that Microsoft or Amazon have conducted HRDD to evaluate the human rights implications of specific functionalities that are offered as part of their PaaS services.

Due Diligence by IaaS Providers

The less specific nature of IaaS services also has implications for how companies providing such services should carry out HRDD. Given that IaaS companies provide as a service the same computing capabilities that equipment manufacturers provide in tangible products, one might expect HRDD in the IaaS sector to resemble the approaches taken by equipment vendors to evaluating their human rights risks.

There is, of course, much that IaaS companies can learn from how equipment vendors evaluate and manage their human rights risks, but IaaS companies are distinguishable from equipment vendors in that they build and operate enormous data centres at locations around the world. IaaS companies are therefore in ongoing business relationships with their customers, whereas equipment vendors only have service and support obligations to their customers once title to the equipment being purchased passes to the buyer.

Furthermore, as the case study involving Google's decision to build a cloud data centre in Saudi Arabia indicates, there are considerable human rights risks that arise from a company's decision to build computing infrastructure and locate customer data in jurisdictions with varying human rights records. Correspondingly, an evaluation of these jurisdictional risks should be at the heart of an IaaS provider's HRDD processes.

¹⁰⁴ BSR, 'Google Celebrity Recognition API Human Rights Assessment | Executive Summary' (October 2019), <https://services.google.com/fh/files/blogs/bsr-google-cr-api-hria-executive-summary.pdf> (accessed 26 January 2022).

¹⁰⁵ Article One Advisors, 'Microsoft HRIA', <https://www.articleoneadvisors.com/case-studies-microsoft> (accessed 25 January 2022)

This notion is hardly a revelation: indeed, the risks associated with CSPs building infrastructure and storing customer data in ‘difficult jurisdictions’ is what catalysed the original internet and human rights conversation in the mid-2000s.¹⁰⁶ The scrutiny that Yahoo! attracted following its Chinese subsidiary’s compliance with a law enforcement demand seeking a dissident journalist’s email records gave rise to the formation of the Global Network Initiative in 2010,¹⁰⁷ and to a wave of scholarship and analysis on online jurisdictional challenges.¹⁰⁸ Similar concerns appear to have spurred the creation of a new set of ‘Trusted Cloud Principles’ which were released by a coalition of CSPs in September 2021.¹⁰⁹ These Principles, which have been endorsed by Amazon, Google, Microsoft, SAP and Salesforce.com, among others, calls upon governments to request data from cloud customers rather than providers in the first instance, and to work together to resolve conflicts of law regarding access to cloud-stored data for investigative purposes.¹¹⁰

The Saudi case study establishes that Google conducts HRDD before locating cloud infrastructure in new jurisdictions, and there is also publicly available information that suggests that Microsoft engages in similar due diligence whenever it expands its cloud service offerings to a new market.¹¹¹ No such information regarding Amazon’s approach to HRDD in relation to its IaaS offerings was uncovered while researching this article, however.

HRDD is meaningful only if its results are integrated into an enterprise’s operations – as contemplated by UNGP 19. We cannot know whether Google or Microsoft have adequately integrated their HRDD findings into their operations, but it is encouraging that both have HRDD processes in place to capture and evaluate human rights-related risks arising from their cloud computing offerings.

Customer-Focused Human Rights Due Diligence

Many human rights risks across the ICT sector arise not from product functionalities in the abstract, but rather from the concrete misuse of these functionalities by end-users. While there are certainly steps that IaaS and PaaS providers should take to mitigate the human rights risks posed by their products and services, there is also a pressing need for companies to evaluate the risks posed by particular customers using their offerings.¹¹²

As explained in Part I, CSPs make their services available to anyone who creates an account with them and provides a valid method of payment. No information was found while researching this article to suggest that the three industry leaders conduct due diligence to

¹⁰⁶ See, e.g., Surya Deva, ‘Corporate Complicity in Internet Censorship in China: Who Cares for the Global Compact or the Global Online Freedom Act?’ (2007) 39:2 *George Washington International Law Review* 255–320.

¹⁰⁷ Dunstan Hope, ‘Global Network Initiative: An Ethical Compass for Information and Communications Firms in the Internet Age’ (2011) 30:3 *Global Business and Organizational Excellence* 7–14.

¹⁰⁸ See, e.g., Jack L Goldsmith and Tim Wu, *Who Controls the Internet? Illusions of a Borderless World* (New York: Oxford University Press, 2006); Teresa Scassa and Robert J Currie, ‘New First Principles – Assessing the Internet’s Challenges to Jurisdiction’ (2011) 42:4 *Georgetown Journal of International Law* 1017–1082; Dan J B Svantesson, *Solving the Internet Jurisdiction Puzzle*, 1st edn (Oxford: Oxford University Press, 2017); Andrew Keane Woods, ‘Litigating Data Sovereignty’ (2018) 128:2 *Yale Law Journal* 328.

¹⁰⁹ Aimee Chanthadavong, ‘Amazon, Google, Microsoft and Other Tech Giants Establish Trusted Cloud Principles’, *ZDNet* (30 September 2021), <https://www.zdnet.com/article/amazon-google-microsoft-and-other-tech-giants-establish-trusted-cloud-principles/> (accessed 27 January 2022).

¹¹⁰ Trusted Cloud Principles, ‘Principles’, <https://trustedcloudprinciples.com/principles/> (accessed 27 January 2022).

¹¹¹ Global Network Initiative, ‘The GNI Principles at Work: Public Report on the Third Cycle of Independent Assessments of GNI Company Members 2018/2019’ (22 April 2020), <https://globalnetworkinitiative.org/wp-content/uploads/2020/04/2018-2019-PAR.pdf> 60 (accessed 25 January 2022).

¹¹² BSR, *Human Rights Due Diligence of Products and Services* (July 2021), <https://www.bsr.org/en/our-insights/report-view/human-rights-due-diligence-of-products-and-services/> (accessed 11 March 2022).

identify human rights risks posed by particular end-users of their services. That said, CSPs likely have customer-focused due diligence systems in place for other reasons, such as to comply with U.S. sanctions and export control laws which prohibit the sale of cloud services to certain countries and entities worldwide.¹¹³

It may not be feasible for companies to review the human rights risks posed by each of their customers, but it is essential that companies have processes in place to identify high-risk users and review their conduct.

CSPs can look to the Know Your Customer (KYC) regime that has developed to combat financial crimes for inspiration in designing HRDD processes focused on customer-related risks. KYC requirements, which are specified in the laws of many jurisdictions, require financial services providers to conduct various forms of due diligence regarding their customers throughout the lifecycle of their business relationship.¹¹⁴ For example, U.S. KYC regulations require banks to identify beneficial owners who possess a 25 per cent or greater interest in a customer, although this beneficial ownership threshold is lower for organizations possessing certain characteristics that suggest that they are risky.¹¹⁵

CSPs could emulate the KYC regime and expand on their export control and sanctions-related due diligence processes to consider human rights-related risks as well. Such systems would evaluate the risks customers pose based on their geography, their industry, and other relevant factors – such as the magnitude of their usage, or their deployment of PaaS applications that pose a high degree of human rights risk – such as face recognition or machine learning involving sensitive data.

As cloud services are billed on a per-use basis,¹¹⁶ companies might also consider monitoring customers' usage patterns for indications of human rights risk. CSPs could even deploy their machine learning expertise to develop automated systems to detect patterns of behaviour that are correlated with human rights risk. These means can be used to prioritize particular customers and use-cases for more detailed due diligence.

Beyond Due Diligence

There is more to business respect for human rights than due diligence. As UNGP 15 makes clear, companies must also have policy commitments in place to respect human rights, and processes to address and adverse human rights impacts when they are connected to a company's business activities.

When a business is contributing or directly linked to adverse human rights impacts, *leverage* is a key factor in determining how it should respond. Leverage refers to the ability of an enterprise to 'effect change in the wrongful practices of an entity that causes a harm'.¹¹⁷ The commentary to UNGP 19 implores businesses to exercise 'leverage to prevent or mitigate' adverse impacts in their value chain when they possess it, and to increase it when they do not.

¹¹³ John P Barker et al, 'Sanctions-as-a-Service: US Regulators Escalate Sanctions Enforcement Priorities Into the Cloud with SAP Settlement | Publications and Presentations', *Arnold & Porter* (7 May 2021), <https://www.arnoldporter.com/en/perspectives/publications/2021/05/us-regulators-escalate-sanctions-enforcement> (accessed 25 January 2022).

¹¹⁴ Armen Meyer et al, 'FinCEN: Know Your Customer Requirements', *The Harvard Law School Forum on Corporate Governance Blog* (7 February 2016), <https://corpgov.law.harvard.edu/2016/02/07/fincen-know-your-customer-requirements/> (accessed 25 January 2022).

¹¹⁵ *Ibid.*

¹¹⁶ Google Cloud, 'Overview of Cloud Billing Concepts', <https://cloud.google.com/billing/docs/concepts> (accessed 25 January 2022).

¹¹⁷ UNGPs, note 3, Principle 19.

Leverage is a key concept in thinking of how CSPs should respond to the human rights risks posed by their customers' misuse of their services. Given the market concentration in the cloud sector – especially among PaaS providers – CSPs possess considerable leverage over their customers.¹¹⁸ This is especially so given the very high costs associated with switching from one PaaS provider to another.¹¹⁹ Consider the example of LEO Technologies, discussed in Part II, above. If the company were to shift to a competitor's PaaS offerings, it would have to reprogram many key aspects of its flagship surveillance product (Verus) to operate on those other platforms. Correspondingly, Amazon and other PaaS providers can use this leverage to foster respect for human rights in their customers' business practices.

One way CSPs can use their leverage is by incorporating human rights conditionality into the standard-form contracts that govern most aspects of their service provision to most customers. The incorporation of such clauses is becoming more and more common in the supply chain context, especially with the development of model contract clauses that make them easy for transactional lawyers to insert into contracts.¹²⁰ A similar approach could be beneficial in the cloud computing context as well, accompanied by adequate transparency measures to gauge their effectiveness.

As explained in Part I, the acceptable use policies of all three companies require customers to agree that they will use cloud services for lawful purposes only, and that they will respect 'the rights of others' in doing so. Neither of these expressions refers specifically to human rights, although an enterprising lawyer could argue that these provisions contemplate human rights considerations as human rights are a branch of law. Some companies, notably Amazon and Google, have gone a step further by including express references to human rights in provisions of their standard-form contracts governing the use of face recognition technology.

One way that providers could expand on these efforts is to apply the human rights policies they have developed to govern their own operations to their customers as well. Microsoft and Google have been leaders within the ICT sector in developing company-wide human rights policies, and in promulgating detailed principles to govern the development and use of technologies incorporating AI. Google's AI Principles require such technologies to be used for socially beneficial purposes only.¹²¹ They also commit Google not to design or deploy AI technologies that 'cause or are likely to cause overall harm', or whose 'purpose contravenes widely accepted principles of international law and human rights'.¹²² Meanwhile, Microsoft's Global Human Rights Statement commits the company to 'helping people use technology ... to protect advanced privacy, security, safety, freedoms of opinions, expression, association, peaceful assembly, and other human rights'.¹²³ It also commits Microsoft to conducting 'due diligence to assess the impact of (its) technologies on human rights'.

CSPs could mandate compliance with such human rights principles in their contracts to reduce the adverse human rights impacts of these services. Or they could develop tailored

¹¹⁸ Ariel Levite and Gaurav Kalwani, 'Cloud Governance Challenges: A Survey of Policy and Regulatory Issues', Carnegie Endowment for International Peace Working Paper 11/2020, https://carnegieendowment.org/files/Levite_Kalwani_Cloud_Governance.pdf 11 (accessed 25 January 2022).

¹¹⁹ *Ibid.*, 14.

¹²⁰ Sarah Dadush, 'Contracting for Human Rights: Looking to Version 2.0 of the ABA Model Contract Clauses' (2019) 68:5 *American University Law Review* 1519–1554.

¹²¹ Google AI, 'Artificial Intelligence at Google: Our Principles', <https://ai.google/principles/> (accessed 25 January 2022).

¹²² *Ibid.*

¹²³ Microsoft CSR, 'Human Rights Statement', <https://www.microsoft.com/en-us/corporate-responsibility/human-rights-statement> (accessed 25 January 2022).

contractual causes to address particular human rights risks, including by requiring their customers to conduct HRDD, or by mandating remedies when harms occur.

V. Conclusion

It is an enduring irony that a technology named after a natural phenomenon that is so clearly visible in the sky has escaped serious human rights analysis for so long. As this article has shown, cloud computing underlies so many of the most powerful yet controversial technologies of our age, yet cloud computing has escaped serious human rights analysis and scrutiny thus far.

This situation is similar to the conversation around human rights in global supply chains a quarter of a century ago. Human rights abuses in distant places could easily be ignored when the nature of our connection to them was not widely known. But greater public awareness of abuses carried out in the fields and factories supplying us with the necessities of life gave rise to scrutiny, analysis, and eventually change as well. My hope is that this article sparks a similar conversation about the human rights impacts of cloud computing, by showing the key role the cloud plays in powering so many modern technologies.

There is a vast research agenda to be pursued on the human rights impacts of the cloud requiring contributions from many different disciplines. Technical research is required to determine whether human rights protections can be designed into cloud computing technologies,¹²⁴ and to determine the feasibility of CSPs using AI techniques to detect customer usage patterns that might be indicative of human rights abuse. Legal research is needed into the effectiveness of incorporating human rights conditionality into contracts for services such as cloud computing, in view of the current focus of such efforts on contracts for the production of goods through global supply chains.¹²⁵

There is also a need for interdisciplinary research into how Chinese cloud services providers are approaching their human rights responsibilities. This article has examined the human rights impacts of the cloud computing industry by focusing on the operations of Amazon, Microsoft and Google, but it is important to recognize that China's Alibaba and Huawei both have significant cloud computing businesses¹²⁶ that are now beginning to grow internationally.¹²⁷ As with Huawei's telecommunications equipment business,¹²⁸ concerns have been raised by the U.S. government and others regarding the security of Chinese cloud providers in view of the close connections between such companies and the Chinese government.¹²⁹ Furthermore, the historic willingness of Chinese technology firms to do

¹²⁴ See, e.g., Vivek Krishnamurthy, 'Are Internet Protocols the New Human Rights Protocols? Understanding "RFC 8280 – Research into Human Rights Protocol Considerations"' (2019) 4:1 *Business and Human Rights Journal* 163.

¹²⁵ See, e.g., Martijn Scheltema, 'The Mismatch Between Human Rights Policies and Contract Law: Improving Contractual Mechanisms to Advance Human Rights Compliance in Supply Chains' in Lisbeth Enneking, Ivo Giesen, Anne-Jetske Schaap, et al (eds.), *Accountability, International Business Operations, and the Law*, 1st edn (Abingdon, UK; New York, USA: Routledge, 2020) 259.

¹²⁶ Meghan Rimol, 'Gartner Says Worldwide IaaS Public Cloud Services Market Grew 40.7% in 2020', *Gartner* (28 June 2021), <https://www.gartner.com/en/newsroom/press-releases/2021-06-28-gartner-says-worldwide-iaas-public-cloud-services-market-grew-40-7-percent-in-2020> (accessed 11 March 2022).

¹²⁷ Arjun Kharpal, 'Alibaba Expands Cloud Business Abroad with New Data Centers in Asia, Stepping up Rivalry with Amazon', *CNBC* (20 October 2021) <https://www.cnn.com/2021/10/20/alibaba-expands-cloud-business-abroad-with-new-data-centers-in-asia.html> (accessed 11 March 2022).

¹²⁸ 'US Designates Huawei, Four Other Chinese Tech Firms National Security Threats', *Deutsche Welle* (13 March 2021), <https://www.dw.com/en/us-designates-huawei-four-other-chinese-tech-firms-national-security-threats/a-56860474> (accessed 11 March 2022).

¹²⁹ Donna Goodison, 'China's Alibaba Cloud Focus of U.S. Government Probe: Report', *CRN* (18 January 2022) <https://www.crn.com/news/cloud/china-s-alibaba-cloud-focus-of-u-s-government-probe-report> (accessed 11 March 2022).

business in jurisdictions with difficult human rights records¹³⁰ raises the age-old question of whether rights-respecting companies should operate in such places or avoid them instead.

Lastly, there are unanswered questions about who is responsible for remedy when cloud computing technologies give rise to adverse human rights impacts. The 'Remedy' pillar of the Protect, Respect and Remedy framework has long proven difficult to implement in the ICT sector,¹³¹ but it is of key importance to ensure that new technologies remain a positive force for human rights. Correspondingly, work is needed to develop guidance on what remedy means in the technology sector generally, and for CSPs specifically, for the three-pillar framework underlying the UNGPs to be effective in our increasingly digital age.

Acknowledgements. I would like to thank Dunstan Allison-Hope, Steve Crown, Hannah Darnton, Marwa Fatafta, Peter Micek, Jason Millar, Shont Miller, Lale Tekisalp, Alexander Voss and Alex Walden for their helpful insights on various aspects of this project, and the participants at a workshop hosted by the Carr Center for Human Rights Policy at the Harvard Kennedy School for their feedback on an earlier version of this paper. I am grateful to Laura Reed for her outstanding work on the advocacy campaign described in the 'Conflicts of interest' declaration below, and to Rachel Kuchma for her extraordinary research and editorial assistance.

Conflicts of interest. As the Director of CIPPIC, I was engaged in advocacy relating to Google's decision to build a new cloud data center in Saudi Arabia. Google's decision to do so is examined as a case study in this article, and published correspondence related to this advocacy campaign is cited throughout.

¹³⁰ Samuel Woodhams, 'Huawei, Africa and the Global Reach of Surveillance Technology', *Deutsche Welle* (12 September 2019), <https://www.dw.com/en/huawei-africa-and-the-global-reach-of-surveillance-technology/a-50398869> (accessed 11 March 2022).

¹³¹ See, e.g., Access Now, *Forgotten Pillar: The Telco Remedy Plan* (Access Now, 2013), https://www.accessnow.org/cms/assets/uploads/archive/docs/Telco_Remedy_Plan.pdf (accessed 11 March 2022).