RE: Comments to the Federal Trade Commission Regarding Advanced Notice of Proposed Rulemaking on Commercial Surveillance and Data Security ("Commercial Surveillance ANPR, R111004")

Dear Commissioners,

Palantir Technologies ("Palantir") is a US-based software company with a global presence. We build data platforms that enable public, private, and non-governmental organizations to integrate, analyze, and collaborate on their data in a secure and privacy-protective way. Our vision is a future in which public institutions, commercial enterprises, and non-profit organizations are fully equipped to more effectively and responsibly use their data to carry out their mandates, to deliver value to their customers and constituencies, and to provide critical services to those most in need.

Palantir operates as a data processor with respect to our clients' data. Unlike many other commercial and technology companies, our business model is not based on the collection, storage, dissemination, or monetization of consumer or citizen data. By contrast, our business involves building and deploying software to help some of the most critical organizations around the world make better use of the data they already lawfully possess or access.

As we build and implement technology, we believe that protecting privacy and other civil liberties is essential to that mission. We therefore welcome efforts by the Federal Trade Commission ("FTC" or "Commission") to establish new trade regulation directed at reducing harmful commercial surveillance and lax data security practices. We seek to contribute to these efforts by sharing some of the lessons that we have learned in our nearly 20 year history about effective Privacy by Design and Default (PbDD) technology practices and about the ways that regulation — either self- or externally-imposed — can help to establish institutional governance and cultures of responsibility around data security, data protection, and informational privacy.

Specifically, we aim to describe some of the principles that we have developed internally as our guidance for building technology that promotes the responsible and value-enhancing use of information assets, including and especially as they relate to advanced data science and analytics techniques. Our experience of working with a broad range of private and public organizations has allowed us to gather unique, pragmatically oriented insights into the challenges of responsible data use and analytics.

We believe that some of these lessons can be a helpful resource for the Commission as it works towards a Trade Regulation Rule on Commercial Surveillance and Data Security. We also believe that our insights may help point the direction to ways that technological innovation can proceed responsibly within the bounds of well-constructed ethical and regulatory constraints. In fact, it may be the case that thoughtfully establishing additional rules directing commercial organizations to optimize their technology and data practices for *both* business outcomes *and* consumer privacy

interests may be a pathway to freeing industry from zero-sum proclivities (that data utility and commercial value creation must necessarily come at the expense of consumer rights, fair practices, and data privacy).

Given Palantir's strategic and operational focus on responsible data integration and analytics, our following response to the Advance Notice of Proposed Rulemaking is limited to Questions 24, 26, 31, 32, 35, 45, 46, 47, 48, 49, and 51. The Executive Summary immediately following provides a brief overview of the insights we aim to share throughout our detailed responses to these questions.

We are thankful to the Commission for this opportunity to contribute to the ANPR and we welcome any requests for clarification, as well as further occasions to contribute as the rulemaking process continues.

Sincerely,

Courtney Bowman, Global Director of Privacy and Civil Liberties Engineering, Palantir Technologies
Arnav Jagasia, Privacy and Civil Liberties Engineering Lead, Palantir Technologies
Helena Vrabec, Data Protection and Privacy Lead, Palantir Technologies

Table of Contents

Executive Summary

To briefly summarize the responses to follow, we encourage the Commission to consider that:

- Cost-benefit analysis should privilege as a benefit the long-term goals of preempting embedded surveillance practices.
- Cost-benefit analysis should consider as a benefit ways that new capabilities and new forms of accountability can be innovated as a result of well-developed privacy rules.
- Rules such as data minimization could dramatically improve user privacy without requiring action on consumers' part. We detail several reasons why we think minimization is among the most effective consumer privacy mechanisms, including how minimization contributes to more effective oversight, how minimization supplements privacy by design, and how minimization prevents data spills.
- Rules such as data minimization are simpler to implement and are less reliant on ex post, privacy-enhancing technologies that may be promising in experimental environments but remain unproven in live business environments.
- Rules should specify goals rather than technical means to achieve those goals. Technology specificity can result in choosing poorly, and it can overlook the layering of different approaches tailored to firm's specific privacy threat models.
- The Commission should start a Section 18 process to establish baseline security rules because security incidents are widespread, because consumers cannot negotiate security protections, and because security incidents cause substantial consumer injury.
- Specifically, the Commission could mandate data deletion schedules, the minimization of all sensitive data, the requirement to keep provenance of data, and a requirement for data security governance.
- It is technologically possible to implement purpose specification and purpose-based access rules in managing and using commercial datasets.
- Concerns that data minimization will adversely affect machine learning ("artificial intelligence") are unfounded because data quality is important to ML efficacy. Responsible artificial intelligence frameworks converge upon values that are harmonious with privacy and product safety.

Balancing Costs and Benefits

In responses to Questions 24 and 26, we provide supporting arguments for framing the balancing of costs and benefits on a longer time horizon that encourages the weighting of more lasting societal outcomes over near-term gains. In a related vein, we advocate that new trade regulation rules on data security and commercial surveillance, if thoughtfully constructed, both enable and even enhance innovation by directing it towards the most socially valuable outcomes — including the protection of informational privacy.

Response to Question 24

Question 24: The Commission invites comment on the relative costs and benefits of any current practice, as well as those for any responsive regulation. How should the Commission engage in this balancing in the context of commercial surveillance and data security? Which variables or outcomes should it consider in such an accounting? Which variables or outcomes are salient but hard to quantify as a material cost or benefit? How should the Commission ensure adequate weight is given to costs and benefits that are hard to quantify?

Costs and benefits cannot be adequately assessed without first framing the time horizon of the assessment. Assessments of commercial business practices should, therefore, treat as a first order concern the potential for negative, long-term effects of surveillance practices as these are the hardest to assess and control, but also likely the most harmful and lasting. By the same token, the rulemaking should be oriented towards structural effects as opposed to providing quick fixes that may work in the short-term but do not also ensure consumer protections in the long term. The long view is essential for focusing companies on business outcomes that are more likely to avoid accretive practices that entrench persistent, embedded surveillance.

With this in mind, we believe that the FTC rulemaking should **incentivize companies to prize long-term social value of data protection and privacy over short-term financial gain (although that gain might be easier to quantify).** We expect that companies will assign more value to privacy and security as a consequence of this approach.

Companies with a “move fast and break things” ethos may operate under a limited perception of harms – e.g., harms that can be resolved with a “bug fix” in a future software release. Such an approach, however, can lead to short-sighted decisions if used to address the impacts of unbridled commercial surveillance. The new rule framework should encourage a (re-)orientation towards values and broader social impact. Specifically, the incremental cost of additional planning cycles, engineering hours, and subject matter expert input in order to, for example, ensure Privacy by Design and Default (PbDD) principles are embedded in a new commercial offering may pale in comparison to the benefits of reinforcing tenets central to a free and open

society. The more immediate risk of regulatory fines might serve as a powerful accelerant to align this view of cost-benefit assessments.

Secondly, we believe that the FTC rulemaking should **help stir innovation towards providing long-lasting consumer and business benefits.** Although the short-term effects might indicate potential losses for business, the Commission's focus should be on achieving long-term gains for both consumers and businesses that handle consumer data.

Experiences from other jurisdictions suggest to us that privacy regulations, when carefully constructed, need not stifle innovation and instead may even contribute to or nudge technological developments that might not have otherwise been possible. For example, consider the data subject rights frameworks under the General Data Protection Regulation ("GDPR")[1] and the Brazilian General Personal Data Protection Law ("LGPD").[2] After the laws were enacted (along with supporting enforcement efforts), companies managed to move from the old, analog method of handling data requests to more sophisticated, scalable, and responsive digital handling methods.[3] Some regulatory mandates created entirely new capabilities that exposed irresponsible practices. For instance, HIPAA required the engineering of accounting of disclosures that in turn revealed "browsing" of celebrities and others' medical records.[4] Similarly, the GDPR requires controllers to keep a record of processing personal data.[5] While creating a record requires a short-term investment, such records have long-term gains for businesses and consumers alike. Specifically, they contribute to more accurate, easily available and searchable data records, as well as documentation of access and use violations.

---

[1] *See* Regulation (EU) 2016/679, of the European Parliament and the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), art. 5, 2016 O.J. (L 119) 1 [hereinafter GDPR].

[2] *See* Lei No. 13.709, de 14 de Agosto de 2018, DIÁRIO OFICIAL DA UNIÃO [D.O.U.] de 15.8.2018 (Braz.) [hereinafter LGPD].

[3] *See Commission Staff Working Document accompanying the document Communication from the Commission to the European Parliament and the Council: Data protection as a pillar of citizens' empowerment and the EU's approach to the digital transition - two years of application of the General Data Protection Regulation* (Eur. Comm'n. SWD/2020/115 final), https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52020SC0115&from=EN.

[4] Health Insurance Portability and Accountability Act of 1996 (HIPAA), Pub. L. No. 104-191, 110 Stat. 1936 (codified as amended in scattered sections of 18, 26, 29, and 42 U.S.C.) [hereinafter HIPAA].

[5] GDPR, *supra* note 1, art. 30.

Response to Question 26

Question 26: To what extent would any given new trade regulation rule on data security or commercial surveillance impede or enhance innovation? To what extent would such rules enhance or impede the development of certain kinds of products, services, and applications over others?

Contrary to the predisposition of many other information technology firms, we believe there is good reason to question the presumption that trade regulation is prone to constrain product and business development and will necessarily impede the capacity of firms to innovate. Some industry advocates will go so far as to suggest that the very construction of rules — independent of their structure and content — impose a looming specter of regulatory backlash that prohibits talented creators from exploring novel ideas, even when those ideas may be wholly defensible. Proponents of this view will point to highly regulated sectors and jurisdictions as areas of limited innovation, while elevating examples like the Silicon Valley of the 2000s — an epoch and place with nary a regulatory concern — as the locus of technology and cultural renaissance. While lax regulation may have allowed for innovation in these instances, this simplified narrative glosses over the collateral damage: for every successful innovator there are hordes of bad idea engines operating with venal intent and net destructive force. And even the once-lauded internet behemoths, many of whom rapidly developed in an era of minimal regulation, have now come under intense and growing scrutiny for the societal collateral impact of their massive economic success.

Indeed, in contrast with this prevailing industry view, it should *not* be taken as a foregone conclusion that innovation and regulation are zero sum. On the contrary, we observe that careful regulation can establish meaningful rails for encouraging targeted innovation that promotes, rather than undermines, social outcomes. Regulation with real enforcement teeth may, however, be needed to establish clear lanes of innovation.

It is our belief that trade regulation rules reinforcing data security and data protection principles, rather than impeding innovation, can instead help to establish safer lanes for innovative product and business developments. Take for example a rule that might establish a baseline or default standard for mandatory data minimization by default in the absence of a compelling exception. The rule would encourage interventions such as obscuring, masking, or encrypting of sensitive data fields *at the point of collection* or, barring that, *at the point of ingestion into the systems using the data*. Such a rule framework might further articulate an onus on potential downstream use cases and users to innovate ways of working with minimized or deidentified data carrying lower privacy risks, or at least to think twice about carrying out frivolous or blatantly exploitative data practices that would require a clear justification in order to reidentify or unmask the data.

There are, however, at least two critical elements to establishing truly effective and lasting trade regulation rules on data security or commercial surveillance. The first is to **avoid overly specific technologies or technical terminology in crafting rules**. Technology changes rapidly and rules that are pegged to the specific technology

implementations at the time of rule-setting are likely to rapidly become obsolete or irrelevant. The second is to **focus on rules that favor operational utility as the standard of excellence rather than pinning hopes on seemingly promising but often over-hyped research and development projects**.

We respect and participate in academic dialogues on privacy enhancing technologies (PETs). This work is important to develop the field, yet few academics studies implement solutions at business-scale. Thus, regulations that establish impractical standards based on limited results that have only been shown in experimental settings may create unreasonable expectations for practicable innovations in live production. The landscape of PETs is increasingly littered with examples that bear out this point. For example, several years of touting specific PETs such as Differential Privacy and Fully Homomorphic Encryption have produced much hype, but arguably limited commercial value. The reason is that these technologies, while technologically impressive in certain controlled settings, are often far from comprehensive solutions to all of the commercial sector's privacy woes. In most practical settings, some of these tools may have very narrow applicability, and even slimmer successful implementations to point to. This is not to suggest that all PETs are doomed to rapid obsolescence or that all such technologies must achieve a mark of perfection to justify deployment, but rather to emphasize that the value of specific PETs must be proven in real-world settings rather than assumed based on marketed messaging.

By contrast and returning to the earlier *obfuscation by default* example, a corresponding rule that seeks to affirm through regulatory rulemaking the necessity of minimization by default lends itself to perhaps less exotic but more reliable, proven, and easy to implement privacy-enhancing innovations. For example, such a rule may be transposed into practices of masking data at the point of data collection or ingestion using a number of flexible and adaptable encryption tools to obfuscate sensitive data using irreversible or reversible tokenization (as the specific needs dictate). The point being that rules that seem to place an emphasis on wedging in the use of exotic, well-hyped, but poorly proven and narrowly applicable PETs will often lead to practical failure and confused notions about how technologies should be appropriately developed and adapted to address privacy risks. And innovation may truly be inhibited in the sense that firms are then forced to make outsized investment on implementing techniques that may never work, and at the expense of pursuing other more value-generative development projects.

Data Security

In the responses to questions 31, 32, and 35 we argue for enhanced data security and suggest avenues which can be used to incentivize security-friendly business practices. We express support for strict enforcement of security standards, particularly those technical and organizational measures that we have seen successfully implemented in practice. Finally, we call for the Commission to consider rules enforced by foreign jurisdictions to avoid unnecessary regulatory burden and open up international markets.

Response to Question 31

Question 31: Should the Commission commence a Section 18 rulemaking on data security? The Commission specifically seeks comment on how potential new trade regulation rules could require or help incentivize reasonable data security.

Because security incidents are widespread and because they cause consumer injury, we believe the Commission should commence a Section 18 rulemaking on data security. In our experience building technology that rigorously enforces and upholds data protection standards, we have found that the data management, retention, and minimization capabilities of a robust data security architecture are critical for protecting against commercial surveillance risk and harms. These capabilities are often available to commercial organizations, but they may lack sufficient motivation to adopt and use them.

Trade regulation rules not only can enforce organizations' use of robust data security practices but can also point organizations toward value-generating business practices that align with consumer expectations of data security. Pursuant to the FTC Act, the Commission can seek civil penalties for violations of trade regulation rules. These actions can compel compliance with strong security standards and enforce baseline requirements across all organizations handling consumer data. Enforcement actions, however, can also direct businesses towards value creation that aligns with consumer expectations and benefits regarding data security. In our experience, consumers value companies that can provide both business value and strong data security capabilities. However, many companies are failing deliver on consumer preferences leading to growing consumer distrust. A regulated realignment toward more privacy-protective practices can therefore carry the benefit of both enhancing consumer trust domestically while also making businesses more competitive in regulated international markets.

The Commission's rulemaking can further promote guiding principles and best practices for data security and privacy. A trade regulation on data security would give the Commission the opportunity to establish general baseline requirements to address some of the most common customer harms, without proposing overly prescriptive rules that might require further industry- or domain-specific context. For example, the Commission can enforce that organizations handling sensitive data implement some

form of access controls that prevent the indiscriminate or arbitrary use of that data. Such an approach can guide organizations to practices that significantly reduce common commercial surveillance harms. Our responses to Questions 32, 45, and 47 propose further tactical and technology-facilitated approaches both for the Commission to consider when enforcing strong data security practices and for regulated organizations to consider when adhering to those data security rules.

Response to Question 32

Question 32: Should, for example, new rules require businesses to implement administrative, technical, and physical data security measures, including encryption techniques, to protect against risks to the security, confidentiality, or integrity of covered data? If so, which measures? How granular should such measures be? Is there evidence of any impediments to implementing such measures?

The Commission's new trade regulations should require organizations to implement technical and administrative measures for data security. Specifically, we believe that proposing the following four general data protection requirements would significantly curb lax data security practices and reduce consumer harms:

1. **Scheduled Deletion.** Consumer data should always be stored with scheduled deletion dates by default. If scheduled deletion is not appropriate, the onus should be placed on the firm to explain why data should be held indefinitely.
2. **Data Minimization.** Sensitive consumer data should always be minimized by default. If sensitive data needs to be preserved in its raw form, the onus should again be placed on the firm to explain why the intended purposes of use preclude data minimization at the outset.
3. **Robust Provenance.** The provenance of data derived from consumers should always be maintained to help ensure that data usage is consistent with the purposes of its collection and to provide a clear accounting of any applicable restrictions as datasets travel through an organization and are processed over time.
4. **Oversight and Governance.** Organizations handling consumer data should institute a framework for oversight and data governance to ensure that consumer data is handled responsibly and securely.

We propose that all organizations managing consumer information should implement the technical and institutional capabilities to carry out these basic data security measures. Each of the four requirements should apply universally to all such organizations as a baseline requirement, and the specific parameters of each requirement (retention windows, minimization strategies, etc.) should be determined in an industry-specific, contextual manner by the appropriate regulatory or industry body.

For example, considering the scheduled deletion of data, the retention window and method of deletion for a piece of data are parameters that will depend on many factors such as the type of data, the intended use of the data, or specific regulation regarding

data retention. As these parameters are contextually driven, the Commission or other regulatory bodies can propose more domain-specific guidance on how firms in a particular industry should meet the baseline requirements outlined above. Regardless of those specific data retention parameters, however, we propose that the Commission should enforce that every system that handles consumer data should be able to implement some form of scheduled deletion. Establishing such a baseline capability will minimize the harms that stem from the indefinite retention and aggregation of sensitive data.

Technical measures for deletion, data minimization, and provenance are complemented by institutional capacities for oversight and governance. From our experience, we have found that the adoption of data security measures is best put into practice by individuals within an organization who have both domain expertise and a clear oversight responsibility. Aligning with increasingly recognized international approaches as in the European GDPR and Brazilian LGPD, the draft American Data Privacy and Protection Act, the currently proposed federal privacy legislation in the United States, and the well-established HIPAA require organizations to have a Data Privacy and Security Officer to ensure compliance with new privacy requirements.[6] In order to ensure that organizations can successfully comply with technical requirements for data security, the Commission should consider similar administrative requirements for organizations to clearly identify individuals with a mandate to carry out these crucial governance responsibilities. Moreover, such individuals performing governance or compliance roles should not be siloed from the technology or data they need to oversee. Instead, we have found success building tools for compliance users right into our software platforms. This reduces friction, improves compliance, and allows both business and oversight functions to work toward the same long-term goals.[7]

Response to Question 35

Question 35: Should the Commission take into account other governments' requirements as to data security (e.g., GDPR). If so, how?[8]

The Commission should take into account other governments' data security requirements, and we advocate that there are at least two (groups of) reasons why this is important.

First, the requirements that other governments have put in place may indicate the right balance between under- and over-regulating. To take the European example,

[6] GDPR, *supra* note 1, art. 37; LGPD, *supra* note 2, art. 5 (VII); 45 C.F.R. §164.308 (a)(2); American Privacy and Data Protection Act, H.R. 8152, 117th Cong. §208(b)(6) (2022).
[7] Paula Cipierre & Yeong Wei Wee, *Data Protection in Palantir Foundry*, PALANTIR BLOG (2020), https://blog.palantir.com/data-protection-in-palantir-foundry-5ab9f346195 ("To facilitate communication and collaboration between these teams, we have developed a data governance infrastructure that allows business and compliance users to collaborate in Palantir Foundry itself").
[8] Our response only addresses the latter part of Question 35, which relates closer to our experience as a global company.

the GDPR is not overly prescriptive regarding security requirements. Instead, it includes a fairly high-level and risk-based language on organizational, contractual, and technical measures that companies should adopt to ensure sufficient data security. Ultimately, each company has to decide for itself what needs to be done to comply with the data protection standards. This makes sense, as not all companies that handle personal and other sensitive data lack security protections, and so the one-size-fits-all approach is inappropriate. For instance, on the market for B2B software where customers' security expectations have been increasingly high, companies can often only survive with strong security practices in place. In this market, standards like ISO, SOC2 and NIST represent key competitive advantages and have enabled the businesses to effectively address security risk even where no governments' requirements are in place. B2B entities have the institutional capacities and expertise to negotiate on privacy and security, and to demand adequate terms. In the other hand, security practices of B2C service providers are often more relaxed. Contrary to the business customers, individual consumers tend to lack the understanding and leverage to achieve better security protections. Given these strong information asymmetries, governments' requirements as to data security appear to be more essential on the B2C market.

The Commission should take this diversity into account and avoid adopting a rule that provides a detailed security roadmap to the companies. Instead, the Commission should follow the EU lead and

- set minimum standards that can stand the test of time,
- ensure exceptions are available for specific situations such as research, and
- continue to promote other well-established external standards (e.g., NIST)

Second, there is value in understanding the requirements of foreign regulations on security because the more the regulations are aligned, the easier it is for the companies to carry out cross-border work. To start with, harmonizing with international standards enables cross-border work and opens up international markets. Currently, the cross-border transfers of data between US and EU represent a major compliance burden for many data-dependent companies as there is a misalignment in the US and EU privacy frameworks. While we do not think the Commission should tailor rules to fit a certain international or regional framework, we do see value in paying attention to data security benchmarks in other jurisdictions. Furthermore, harmonized security standards enable US corporations to build on multinational investments that many have already made and would, for the most part, like to see standardized across the board. Alignment with foreign regulations can help minimize the cost of maintaining substantially different security regimes with separate infrastructure in different jurisdictions. Finally, harmonized security standards can help promote a sense of a level playing field internationally, which alleviates suspicions about American firms and minimizes the ground for complaints that US-based firms are privacy-adverse or anti-competitive.

Purpose Limitation and Data Minimization

In responses to Questions 45 through 49, we provide arguments for the Commission to pursue data minimization and purpose justification requirements as part of this data security rulemaking. We argue that such requirements can significantly improve consumer data security and will neither hamper innovation nor pose undue administrative burdens. Moreover, we advocate for a tiered regulatory approach: the Commission should enforce general baseline data security requirements in this rulemaking, and delegate any more specific requirements to more domain-specific regulations that can better assess the contextual needs of each domain.

Response to Question 45

Question 45: Pursuant to a purpose limitation rule, how, if at all, should the Commission discern whether data that consumers give for one purpose has been only used for that specified purpose? To what extent, moreover, should the Commission permit use of consumer data that is compatible with, but distinct from, the purpose for which consumers explicitly give their data?

Supporting purpose specification and limitation has been a central focus in our practice of building software platforms for our customers. From this nearly 20 year experience of building software for purpose specification and enabling organizations to enforce purpose limitation principles, we have refined approaches for implementing purpose limitation capabilities that more effectively promote need-based data access and responsible information use. There are several techniques that organizations can use to implement purpose specification and limitation capabilities in their software systems, including adopting frameworks for purpose specification or enforcing purpose-based access controls for purpose limitation. Such techniques for purpose limitation and specification not only improve direct institutional data security practices, but also could be used to provide the Commission with the ability to conduct external oversight and ensure regulatory compliance.

We detail two such approaches that we have enabled for our customers to comply with purpose specification and purpose limitation rules. We also comment as to how each of these approaches might be further extended to the Commission's external regulatory work.

Purpose Specification Frameworks: We have built our software platforms with frameworks for configurable purpose specification to address regulatory requirements and best practices across jurisdictions. Purpose specification is the first step of full purpose limitation and asks users to specify their purpose for accessing or performing some operation on data. Even without full purpose *limitation*, purpose specification alone can provide regulators with a better understanding of why individuals have taken certain actions. In practice, all user-specified purposes should be captured in a software

system's audit logs, and this audit trail can be reviewed in real-time or retroactively to assess whether users intend to use data for appropriate purposes.

Purpose-Based Access Controls: Access control mechanisms allow software platforms to define which users can perform which operations on which pieces of data. Granular access control systems can guard against the misuse and repurposing of data, while still allowing for legitimate uses to proceed. Typically, access controls schemes are governed by users' roles within a software system (RBAC - Role-Based Access Controls), based on their attributes (ABAC - Attribute Based Access Controls), or even set and updated dynamically within a system (DAC - Discretionary Access Controls). In supporting organizations responding to the COVID-19 crisis, we implemented Purpose-Based Access Controls to help governance teams enforce purpose limitation in practice.

In organizations using a Purpose-Based Access Controls (PBAC) approach, governance or compliance teams manage a standard, controlled set of legitimate data processing purposes. In order to use data, users first apply for access to a purpose, and each purpose contains the necessary and minimal set of data needed to carry out the task. In this way, PBAC can enforce purpose limitation requirements required by GDPR[9] or other use limitation standards. Ensuring that each use of data gets assigned a purpose and further assigning each purpose an owner via PBAC can strengthen accountability and data security. This approach allows organizations to comprehensively and exhaustively review *all* data processing activity. Moreover, PBAC clarifies why someone has access to data, what they intend to do with that data, and when access to data might no longer be needed – all of which are insights necessary for oversight into the appropriate use of data.

Together, purpose specification frameworks and purpose-based access controls can ensure both that users specify for what purpose they need data and that data is only ever used for that purpose. Techniques in privacy engineering for purpose justification provide organizations with simple yet robust ways to comply with purpose specification rules. We have publicly presented our framework for purpose justification as a demonstration of the feasibility of these approaches in commercial information technologies and in the hope that illuminating the "art of the possible'" will encourage more a commonplace expectation for adopting these capabilities across technology platforms.[10] Moreover, in responding to the COVID-19 crisis, organizations were able to take advantage of PBAC in Palantir Foundry for strict purpose limitation guarantees when handling sensitive data.[11]

---

[9] GDPR, *supra* note 1, art. 5(b).
[10] Future of Privacy Forum, *PEPR 2021: Session 8.2 - Lightweight Purpose Justification Service for Embedded Accountability,* YOUTUBE (Jun. 16, 2021), https://www.youtube.com/watch?v=T3aRNTa2Bwg. This presentation of a framework for purpose justification was given at the 2021 Conference on Privacy Engineering Practice & Respect (PEPR).
[11] Basil Jennings, *Purpose-based Access Controls at Palantir (Palantir Explained, #2)*, PALANTIR BLOG (2020), https://blog.palantir.com/purpose-based-access-controls-at-palantir-f419faa400b3.

The adoption of such techniques also may constitute a critical groundwork for facilitating regulatory oversight by competent authorities. If purpose justifications or requests for processing purposes are recorded in well-structured and accessible audit logs, for example, organizations will be better positioned to provide requisite information to the Commission that can then be used to validate legitimate and appropriate uses of consumer data. Moreover, the Commission could propose and continually revise an authorized set of data processing purposes for consumer data to standardize such audit reviews.

Response to Question 46

Question 46: Or should new rules impose data minimization or purpose limitations only for certain designated practices or services? Should, for example, the Commission impose limits on data use for essential services such as finance, healthcare, or search—that is, should it restrict companies that provide these services from using, retaining, or transferring consumer data for any other service or commercial endeavor? If so, how?

While there are specific sectors, practices, and services that tend to inherently present greater privacy risks, the general availability of data, processing capacity, and growing temptation to exploit accumulated data assets to the hilt leads to privacy harms arising in unlikely places.

We therefore are inclined to advocate for a flexible, risk-based assessment framework that is both broadly applicable across all practices and services and that establishes a tiered framework for the subsequent imposition of data minimization or purpose limitation rules. Such an assessment framework would be aimed at determining the aggregate or overall risk profile of a project or system, which at minimum would evaluate dimensions of risk such as sensitivity of the domain, sensitivity of the data, sensitivity of intended uses and workflows, and sensitivities of reasonably anticipated uses and workflows. The resulting aggregate assessment would than translate into no-, low-, medium-, and high-risk classifications, each with corresponding data minimization or purpose limitation requirements.

On the no-risk end of the spectrum, for example, data and processing tools tracking and analyzing manufacturing sensors might be deemed free of minimization or purpose limitations requirements because the data and workflows carry no direct or indirect connections to natural persons. However, if that same manufacturing line then incorporated data inputs related to worker activities, the risk profile may be ratcheted up to low- or medium-risk if the added worker information is deemed a critical input for understanding the manufacturing line throughput. In this instance, rules imposing minimization on worker data might allow it to be safely used for the intended purpose while removing the residual risk of privacy infringing uses such as worker mobility tracking or performance/disciplinary determinations.

Now, it may be the case that some uses of consumer data may always be sensitive and inherently high-risk. Such a framework would still accommodate those designations and could be elaborated to identify, recommend, or require specific baseline data minimization and purpose limitation standards. One example for the Commission to consider is backstop retention policies for private sector retention of health-related consumer data. Outside of the scope of medical practices (covered by separate regulation), an expectation may be established that such data in non-minimized form (i.e., anything more granular than high-level statistical aggregates) should not be allowed to persist for greater than a predetermined, specific time duration under any circumstances.

The general implication of these referrals is that the best and most effective privacy-protective practices are necessarily contextual and should be drawn from the details of the specific environment, data, and workflow in question. Using these dimensions of evaluation to make contextually appropriate risk assessments ultimately allows for the proportionate and necessary level of tuning of specific data minimization and purpose limitation practices to the identified privacy concerns. A flexible, tiered approach helps ensure that the imposition of rules is not overly onerous, but instead are responsive to the demands of particular practices and the attendant privacy harms that need to be guarded against.

Response to Question 47

Question 47: To what extent would data minimization requirements or purpose limitations protect consumer data security?

Our work across both commercial and government organizations in a variety of domains and jurisdictions has clearly demonstrated that data minimization and purpose limitation requirements can greatly protect consumer data security and privacy. We present four reasons why data minimization and purpose limitation can protect consumer data security drawn from our own experience building technologies that uphold and enforce these crucial data protection principles.

These requirements encourage well-scoped, need-based, and proportionate use of data. Both data minimization and purpose limitation guard against arbitrary and indiscriminate access to data. As we observed during our work helping governments and businesses respond to the COVID-19 crisis, data is not a panacea.[12] In fact, the continuous aggregation of data, especially in the absence of granular access controls, can lead to significant privacy harms if the availability of excess data exposes temptations to misuse it for purposes outside of previously authorized, legitimate workflows. Data minimization and purpose limitation intentionally constrain the way organizations use data to ensure that data is accessed only when necessary. In our

[12] Courtney Bowman, *Best Practices for Using Data During a Crisis,* PALANTIR BLOG (2020), https://blog.palantir.com/best-practices-for-using-data-during-a-crisis-f2639d5eeea4. *See also* Courtney Bowman, *Reflections and Lessons from the COVID-19 Crisis*, PALANTIR BLOG (2022), https://blog.palantir.com/reflections-and-lessons-from-the-covid-19-crisis-b406c03fbb4e.

experience, we have found that data minimization and purpose limitation are essential for ensuring that any sensitive data is used with necessity and proportionality. Moreover, access controls, privacy engineering techniques for data minimization, and frameworks for purpose justification can derive new value for consumers by allowing organizations to provide benefits to consumers without violating their privacy or exploiting their data for other purposes.

These requirements facilitate oversight. Purpose limitation and data minimization rules provide compliance and privacy officers or auditors with a clearer understanding of why and how data is used. As mentioned above, purpose specification – as a component of purpose limitation – helps auditors performing an incident analysis understand not only what happened, but also how a user justified that action.[13] This can help distinguish malicious activity from misunderstandings of data processing rules, allowing for further refinement and continuous improvement of data protection policies. Similarly, default data minimization rules ensures that auditors and compliance officers know specifically when data needs to be accessed in its raw, non-minimized form. By making data minimization the default, each access to granular sensitive data – if any at all – becomes a more intentional action and easier to regulate.

These requirements enforce key principles of Privacy by Design and Default (PbDD). "Privacy by Design and Default" focuses on building privacy-protective paradigms into systems from inception, rather than being incorporated as an afterthought. At Palantir, we build our products with PbDD principles based on best practices from the privacy engineering industry and data protection standards.[14] Data minimization and purpose limitation rules encourage organizations to adopt PbDD approaches from the earliest stages of architecting systems. Data minimization encourages developers to consider the default exposure of any piece of data and make sure there are robust capabilities for data pseudonymization, aggregation, or obfuscation. Purpose limitation encourages developers to build access controls, granular permissioning capabilities, and frameworks for capturing users' intentions when taking sensitive actions to ensure that all data access is need-based and tied to a legitimate processing purpose.

These requirements better prevent accidental or malicious sharing of sensitive data. The repurposing or transfer of consumer data from one system to another represents one of the most serious challenges to enforcing robust data security. Whether accidental or malicious, sharing sensitive data can greatly harm consumers if it is not performed for a legitimate purpose. In our experience, both data minimization and purpose limitation can better prevent the sharing or transfer of data or encourage that such data sharing happens in a need-based, controlled manner that can best protect consumer privacy. Purpose limitations enforce that data can only be used for a pre-specified use, and approaches like Purpose-Based Access Controls can help ensure

[13] *See supra* Response to Question 45.
[14] GDPR, *supra* note 1, art. 25; Anne Cavoukian, *Privacy by Design: The 7 Foundational Principles*, INFO. & PRIVACY COMM'R OF ONTARIO (2009), https://www.ipc.on.ca/wp-content/uploads/resources/7foundationalprinciples.pdf.

that data is not casually repurposed for unauthorized uses.[15] Data minimization techniques can be used to obfuscate sensitive identifiers or aggregate data by default, which would reduce the harm to consumer privacy if accidentally shared.

By enforcing data minimization and purpose limitation rules, the Commission can use this rulemaking to improve the default privacy posture of organizations that handle consumer data.

Response to Question 48

Question 48: To what extent would data minimization requirements or purpose limitations unduly hamper algorithmic decision-making or other algorithmic learning-based processes or techniques? To what extent would the benefits of a data minimization or purpose limitation rule be out of proportion to the potential harms to consumers and companies of such a rule?

It is our position that sensible data minimization, purpose limitation, and — more generally — strong data governance practices need not be treated as fundamentally at odds with developments and innovations in the domain of algorithmic decision-making or other algorithmic learning-based processes or techniques (often referred to as artificial intelligence (AI) and machine learning (ML)). On the contrary, we observe that these principles are often critical for grounding AI/ML processes and techniques in real-world conditions and for providing clear rails for researchers, engineers, and entrepreneurs to focus their efforts.

Data quality matters as much as quantity. As a starting point, we ask the Commission to consider that behind many algorithmic decision-making projects lies the presumption that unbridled data (often referred to as "Big Data") is an essential ingredient to training and building AI/ML models. This fixation on data scale (Volume, Velocity, Variety) is then taken as the singular consideration for unlocking the potential of algorithmic decision-making/AI/ML and that therefore any data minimization or purpose limitation constraint will unduly hinder progress in this domain. What this position neglects, however, is the fact that the qualitative characteristics of data are just as — and arguably more — important as the quantitative dimensions of the data. Insofar as data minimization principles translate into reductions in the scale of data but also serve to enhance qualitative characteristics, including accuracy, precision, representativeness, lineage, etc., AI/ML projects tend to be placed on firmer footing for dealing with both efficacy and ethics considerations when such data security principles are enforced.

Rulemaking that enforces data minimization, purpose limitations, and other data governance principles should emphasize that, especially for critical and consequential data-driven decision making, data is useful not simply by virtue of its size, but because it can be trusted and used effectively. Drawing upon tools and practices that help to

[15] *See supra* Response to Question 45.

manage, validate, and apply data in meaningful ways and towards useful outcomes serves to both minimize risks of undue privacy harms *and* provide a sounder methodological framework for helping to ensure better algorithmic design and deployment outcomes.

Responsible AI Principles reinforce and support the Commission's rulemaking intent. We wish to further draw the Commission's attention to a set of data governance practices and principles, increasingly referred to as Responsible AI Principles, that help to ensure the protection of consumer rights and interests, while enabling responsible use of data for algorithmic decision making. These concepts highlight a need for responsible data handling throughout the full data lifecycle, and not just at the moment of model design and development.

It is worth noting that many of these concepts are already being promoted and enshrined in domain-specific guidance by Federal government entities.[16] The emerging body of Responsible AI frameworks, however, can be further refined by emphasizing an operational focus (as distinct from more theoretical or academic evaluations) that better covers the full lifespan of data and algorithmic use.

In our own business practices, we have established and promote a Responsible AI framework that outlines the key considerations for keeping algorithmic decision-making and other AI projects on operational and ethically sound rails. Here, we wish to highlight a subset of those principles that most closely address the present question of data minimization and purpose limitation measures and accentuate how upholding these principles can be additive to rather than a hindrance of algorithmic decision-making or other algorithmic learning-based processes or techniques.

- **Reliable (Safe, Secure, Resilient, Robust)**: AI systems should be built with capabilities for assessing the safety, security, and effectiveness of models throughout their entire lifecycles. AI systems should also be designed to mitigate or reduce the potential impact of accidents and other unintended harmful behavior[17] and provide capabilities for assessing and eventually minimizing adversarial attempts to either degrade models or undermine the privacy of individuals whose data might have been used to train the models.
- **Traceable (Auditable, Governable)**: AI systems should provide the capabilities to understand relevant development processes, data sources, and the provenance of all data used for model development. AI systems should also provide transparent access to auditable standard operating procedures, design guidelines, and appropriate documentation.[18]

[16] *See generally* Jared Dunnmon et al., *Responsible AI Guidelines in Practice*, DEF. INNOV. UNIT (2021), https://www.diu.mil/responsible-ai-guidelines.
[17] Dario Amodei et al., *Concrete Problems in AI Safety* (Jul. 25, 2016), https://arxiv.org/abs/1606.06565.
[18] Margaret Mitchell et al., PROC. CONF. FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY 220 (Jan. 2019), https://dl.acm.org/doi/abs/10.1145/3287560.3287596.

- **Accountable (Liable, Responsible)**: Accountability has widely been cited as an important consideration for the development of algorithms and models.[19] In order to put accountability into practice, there should be a clear definition of the roles and workflows for people responsible for the different parts of the AI system. Moreover, such systems should allow for both third-party oversight and internal audits.[20]
- **Human-centered (Participatory, Socially Beneficial):** AI systems should benefit individuals, society and the environment overall. They should not erode trust, and should augment, not replace, human decision making. Particularly for uses of automation that impact individuals' privacy and civil liberties, the goal of AI systems should be to enhance the context and quality of human judgment.
- **Scoped (Problem-driven, Reproducible, Rigorous):** AI systems should be built for a well-defined and appropriately scoped purpose. It should be expected that models powering the system are useful within that scope, but outside of that scope, no such guarantee holds. The steps of the model lifecycle must be performed with scientific rigor, so that model results can be reproduced for a given modeling problem.[21]

The Commission should examine rules that encourage organizations to invest in data management, security, and governance tools that facilitate the above (and similar) principles. Strong data management, security, and governance not only contributes to safety of AI, but also promotes privacy and transparency interests, while focusing innovative algorithmic decision-making developments along reasonable, defensible, and more socially beneficial trajectories.

Response to Question 49

Question 49: How administrable are data minimization requirements or purpose limitations given the scale of commercial surveillance practices, information asymmetries, and the institutional resources such rules would require the Commission to deploy to ensure compliance? What do other jurisdictions have to teach about their relative effectiveness?

Based on our experience, both technical controls and organizational procedures can be used to support the administration of data minimization and purpose limitation requirements within organizations regardless of their size and the scale of "surveillance" practices. In responses to the prior questions, we described various examples of the

[19] Maranke Wieringa, *What to account for when accounting for algorithms: a systematic literature review on algorithmic accountability*, PROC. CONF. FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY 1 (Jan. 2020), https://dl.acm.org/doi/abs/10.1145/3351095.3372833.
[20] Raji et al., *Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing*, PROC. CONF. FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY 33 (Jan. 2020), https://dl.acm.org/doi/10.1145/3351095.3372873.
[21] Sayash Kapoor and Arvind Narayanan, *Leakage and the Reproducibility Crisis in ML-Based Science* (Jul. 14, 2022), https://arxiv.org/pdf/2207.07048.pdf.

controls that our customers have successfully deployed as part of our products,[22] and we believe such approaches can be generally used at organizations that process large amounts of data.

With regards to the resources necessary for the Commission to ensure compliance, we do not think that new rules would make enforcement unmanageable or require an unrealistic workforce. To the contrary, we think, for various reasons, that the Commission is in a good position to set ambitious rulemaking goals and avoid making any undesirable compromises. Moreover, the Commission can heavily lean on compliance frameworks, audit formats, evaluations, assessments that are generally accepted across the industry and have introduced standardized frameworks to verify compliance.[23]

Furthermore, the Commission can learn from the experience in other jurisdictions. We think the EU example is particularly illustrative. Although the principle of purpose limitation and the principle of amount limitation are fairly inflexible in the GDPR, the EU authorities have been able to successfully oversee compliance and ensure effective enforcement.[24]

Under the principle of purpose limitation as set forth in the GDPR, the purposes for which personal data is collected must be specified and the data must only be used for those purposes. In turn this means that any secondary data use, unless stipulated at the moment of data collection, is in principle prohibited. More specifically, the data can only be used for a secondary purpose which is compatible with the one for which it was collected. Such prior determination of purposes creates a sense of certainty and transparency and enhances data subjects' control over their personal data.

As the Federal Trade Commission notes, the principle of purpose limitation is difficult to enforce in practice. Most obviously, it is unlikely that all possible reuses can be defined or predicted in advance. This can be frustrating for data economy actors, as they might feel that the possibilities to exploit the collected data have been disproportionately restricted. As a response to the restraining provision, controllers have started using an open and indefinite language that lacks specificity, with regard to both the data the networks collect and how they use this data. However, through the GDPR lens, this may be seen as sidestepping the intention of the legislator, and the processing based on such a policy may be considered illegitimate.[25]

To help data users assess whether reusing the data in another context is legitimate, the GDPR provides detailed guidance on what sort of data processing should be considered compatible. The following criteria are key: (a) links between the purposes for

---

[22] *See supra* Response to Question 45.
[23] *See infra* Response to Question 51.
[24] A search of the *gdprhub.eu* database reveals over 70 decisions recently taken by the EU data protection authorities that rely on the principle of purpose limitation.
[25] Lokke Moerel and Corien Prins, *On the Death of Purpose Limitation*, IAPP (Jun. 2, 2015) https://iapp.org/news/a/on-the-death-of-purpose-limitation/#.

which the data has been collected and the purposes of the intended further processing; (b) the context in which the data has been collected; (c) the nature of the personal data; (d) the possible consequences of the intended further processing for data subjects; and (e) the existence of appropriate safeguards.[26] The European Data Protection Board (a data protection advisory body composed of EU member states' representatives, formerly known as the Article 29 Working Party) proposed a similar test composed of a formal and substantive assessment.[27] The formal assessment is focused on the comparison between the purposes provided by the controller and actual data reuse, and the subjective assessment on the context and the way in which the purposes can be understood.

With regards to the principle of data minimization, the GDPR requires those who control data to observe that data remains relevant, not excessive in relation to the purpose, and kept no longer than necessary for processing.[28] The GDPR stipulates that personal data should be kept in a form that permits identification of data subjects for no longer than is necessary for the purposes for which the personal data are processed. Both requirements are of course at odds with the information-rich society, which collects vast amounts of data because it might prove useful in the future. To overcome this challenge, the GDPR foresees some exceptions to data minimization when data is processed solely for archiving purposes in the public interest, or for scientific and historical research purposes or statistical purposes.[29] Moreover, subject to implementation of appropriate technical and organizational measures, the storage time may be longer. However, authorities have already made clear that the data minimization principle, just like the principle of purpose limitation, should in its essence remain unchanged despite the growing big data sector.[30]

To summarize, the case of GPDR shows that European regulators have not needed to make data security requirements more lax in order to make the regulation more administrable. The GDPR proved that there are a variety of regulatory techniques – substantial monetary fines, ex-post EDPB guidance, carveouts for specific areas that may be disproportionately affected by data minimization and purpose specification such as scientific research – that help make the regulation more administrable while ensuring that data security is strictly controlled.

While compliance remains challenging, the gains of adhering to the principle should not be overlooked. Based on our experience with European customers, the restrictive requirements of the data minimization and purpose limitation principle have incentivized companies to better articulate their data use, align on the goals and implement

---

[26] GDPR, *supra* note 1, art. 6(4).
[27] Article 29 Data Protection Working Party, *Opinion 03/2013 on Purpose Limitation*, at 20, 00569/13/EN, WP 203 (Apr. 2, 2013), https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2013/wp203_en.pdf.
[28] GDPR, *supra* note 1, art. 5(c).
[29] GDPR, *supra* note 1, art. 89.
[30] Case C-131/12, Google Spain SL v. Agencia Española de Protección de Datos, 2014 E.C.R 317 at para. 93-94.

appropriate safeguards. For example, in working with multinational financial customers, our data minimization and purpose justification tooling has unlocked analytics workflows for users with a legitimate business purpose to access sensitive data, while ensuring privacy protection and security meets a strict regulatory bar. While this may create some short-term cost, the long-term gains such as improved data governance and reduced data storage cost can be significant.

Certifications

As a final area of contribution, our response to Question 51 offers explicit suggestions on Data Use and Data Retention Certifications frameworks, how they can be tailored and implemented most effectively, and existing models that such standards may be able to draw upon.

Response to Question 51

Question 51: To what extent, if at all, should the Commission require firms to certify that their commercial surveillance practices meet clear standards concerning collection, use, retention, transfer, or monetization of consumer data? If promulgated, who should set those standards: the FTC, a third-party organization, or some other entity?

We believe that a general expectation should be established through rules or other means for firms to certify how their practices meet standards of collection, use, retention, transfer, or monetization of consumer data. However, as a company that provides business-to-business (B2B) enterprise software and does not engage in any forms of consumer-facing data collection, transfer, or monetization, we will limit our remarks to data **use** and data **retention**, as they relate to functions that our products are built to support or enable for our clients.

With respect to these two specific areas of standards development, we observe — based on nearly 20 years of experience — that strong standards are both technologically and operationally feasible. The tools and tradecraft that allow for responsible data use and sensible data retention handling (up to and including deletion) are not abstract ideas, but rather fully within the realm of possibility and are actively and effectively used already by many institutions. Certifying the use of these capabilities should come as a fairly natural extension of their adoption by companies, especially those in higher-risk workflows.

Data Use Certifications should focus not just on identifying intended purposes for using data, but also on explaining how they will process data to achieve desired outcomes, and how the privacy and security risks relate to those intended uses will be mitigated through, at minimum, practical constraints on data access and processing. As one example of how this might be done, Palantir builds and deploys in our commercial product, Foundry, a system referred to as Purpose-Based Access Control that enables a tight integration of data governance process into the underlying access control of the platform.[31] This system demonstrates how the data use infrastructure, itself, can promote structure and clarity for data access decisions, can be used to capture missing context and make it available to the people who need it, and can provide intuitive tooling for non-technical data governance teams to enforce requisite rules for designated uses.

[31] *See supra* note 11 and accompanying text.

The system is further configurable to generate usage reports for external certification purposes.

Data Retention Certifications could provide a valuable framework for firms to leverage in solidifying their deletion protocols and articulating the assumptions and tradeoffs that factor into their adopted approaches. Inter alia, such certifying statements would encourage organizations to specify not only the timeframes for retention and ultimate deletion, but also the method of deletion chosen (e.g., soft deletion, hard deletion), the threat vectors motivating the chosen deletion method (as compared with other methods), standards for documenting metadata or other records as proof of the deletion even, as well as measures and controls for ensuring the propagation of deletion events to downstream consumers of the data.[32]

Data Use and Data Retention Certifications outlined above offer just two examples of potential certification requirements for firms to follow. Such standards certification regimes, it's worth noting, can be modeled after strong precedents. Standards setting has been elsewhere successfully implemented, for example with ISO standards[33], as well as with Record of Processing Activities (RoPA) requirements laid out in Article 30 of the GDPR.[34] Similar certification frameworks could be extended to the data security principles identified in this question, especially Data Use and Data Retention.

As for the promulgation of these standards, we advocate that standards are best established in context, with knowledge of the industry, applications, and other localized considerations. For example, standards for the specific method and time frame of data retention and deletion regimes will be specific to the nature and risk profile of a given data asset in context. Third-party organizations with industry or use case specialized knowledge (e.g., industry standards groups, competent third-party assessors, or other entities with localized proficiencies) will be best positioned to determine the most relevant or necessary specifications. The general classes of techniques and tools used for carrying out various types of regulated uses and retention protocols, however, may be articulable at a higher level by the Commission.

[32] *See generally* Paula Cipierre & Annabelle Larose, *Designing for Deletion (Palantir Explained, #6)*, PALANTIR BLOG (2022), https://blog.palantir.com/designing-for-deletion-palantir-explained-6-adfe25fda810 (outlining these and related considerations).
[33] *See, e.g.,* INTERNATIONAL ORGANIZATION FOR STANDARDIZATION, ISO/IEC 27701:2019, Security techniques — Extension to ISO/IEC 27001 and ISO/IEC 27002 for privacy information management — Requirements and guidelines (Aug. 2019), https://www.iso.org/standard/71670.html.
[34] GDPR, *supra* note 1, art. 30.