Responses to Written Questions Submitted by Chairman John Thune to Keith Enright

Question 1. Google is subject to an FTC order issued in 2011 that, among other things, required Google to establish and implement, and thereafter maintain, a privacy program designed to: "(1) address privacy risks related to the development and management of new and existing products and services for consumers, and (2) protect the privacy and confidentiality of covered information." In 2012, Google paid a $22.5 million civil penalty to settle allegations that it violated the FTC order. Please describe the steps Google has taken to comply with the FTC's order.

Response. We are committed to ensuring compliance with our FTC consent decree and have dedicated significant attention and resources to ensuring the commitments we made to the FTC are met. Our comprehensive compliance program — one of the most sophisticated in the world — includes:

- Yearly risk assessments, strong internal policies that guide our employees, and training and advice by our privacy and security experts to ensure compliance with our FTC Consent decree and the protection of our users.

- Compliance checks and controls. Product launches and changes to existing products are gated by a unified privacy and security review process ensuring the product's data lifecycle, covering data collection, use, notice and control, sharing, storage and access, and deletion and retention.

- Strong incident response procedures. Potential privacy issues are addressed through our robust incident response program for privacy- and

security-related events. Employees are required to report any suspected security or privacy incidents to our dedicated 24x7x365 worldwide incident response teams, so that we can respond, including by securing and protecting users' data and handling user notifications.

- Regular external assessments — three separate, globally recognized external assessors assess our program on a bi-annual basis. They each have found that our program was operating effectively.

But we will not stop there. We are always looking for ways to refine and improve our privacy program and continue to focus that effort.

Question 2. What, if any, additional steps did Google take to comply with the FTC order after agreeing to settle allegations that it violated the order in 2012?

Response. In this instance, Google moved swiftly to correct an inaccurate statement on our Help Center and to fix an issue with cookies being set where we did not intend as a result of changes made in third-party software. We also took steps to ensure this type of issue could not arise in the future. This work and our agreement with the FTC resolved the matter and allowed us to move on to launching great, privacy-protective products and features.

Question 3. On July 22, 2018, Google responded to questions posed by me, Telecommunications Subcommittee Chairman Wicker, and Consumer Protection Subcommittee Chairman Moran in a letter dated July 10, 2018, regarding the use of Gmail users' personal data by third party email app developers. I have some follow-up questions based on Google's response, which did not fully answer all of the questions we posed in July.

Google said that it pre-installs "Google Play Protect" on all Google-licensed Android devices to continuously monitor users' phones, along with apps in Play and across the Android ecosystem, for potentially malicious apps. Google also asserted that it scans more than 50 billion apps every day and warns users to remove apps Google identifies as malicious.

a. Does Google take any steps to protect consumers from the potentially malicious apps it has identified, other than warning consumers about them? If so, what are these steps?

Response. Yes, Google Play Protect includes on-device capabilities that protect users from potentially harmful apps in real-time, as well as services that analyze device and app data to identify possible security concerns. In 2016, we started scanning all devices for potentially harmful applications once a day. Daily scanning allows Google Play Protect to respond quickly to a detected threat, reducing how long users could be exposed to the threat and how many devices may be affected. Google Play Protect leverages cloud-based app-verification services to determine if apps are potentially harmful. If it finds a potentially harmful application, Google Play Protect warns the user. In cases of severe malware, Google Play Protect can remove the potentially harmful application from affected devices and block future installs. This furthers user safety and mitigates the potential harm the app can cause while providing minimal inconvenience to the user and providing them with control over their device.

Question 4. How many potentially malicious apps has Google identified over the last year? How many consumers heeded Google's warning and removed these apps? Did Google follow up with consumers who did not remove these apps? If so, how?

Response. In 2017, daily Google Play Protect scans led to faster identification and removal of approximately 39 million potentially harmful applications. We automatically disabled such applications from roughly 1 million devices. In November 2017, we updated Google Play Protect to disable potentially harmful applications without uninstalling them. When we can, we try to leave as much choice in users' hands as possible. To walk the line between user choice and safety, when Google Play Protect detects certain kinds of potentially harmful applications, it automatically disables the app. Users are asked to uninstall the app or re-enable it without losing their data. This mitigates the potential harm the app could cause while providing minimal inconvenience to the user while they decide what to do. If the user decides to not respond to a Google Play Protect warning, detected harmful apps that remain on the device will always be presented to the user on the main Play Protect settings screen, so they can get back to this later if they want.

Question 5. In its July 22, 2018, response Google said it acts promptly on user reports about privacy and security issues, rewards researchers and developers who flag privacy and security

issues, and engages in research and community outreach on privacy and security issues to make the internet safer.

a. What specific steps does Google take in response to user reports about privacy and security issues?

Response. Whenever a user reports a potential privacy or security issue, the appropriate team promptly investigates the circumstances surrounding the report to determine whether further action is warranted. The action taken depends on the result of that investigation. For example, if a user reports an app's potential violation of the Play Store's policies, and our review confirms the violation, the relevant app may be removed from the Play Store. As another example, a user report about a security bug in a third-party app may lead us to contact the relevant developer and suspend the app until the bug is fixed.

Question 6. How many such reports has Google received during the last year? How many of these reports warranted action by Google, and what action did Google take? How long did it take, on average, for Google to act on these reports?

Response. Google incentivizes researchers to identify and report vulnerabilities in Google's products and apps as well as in apps developed by partners under various vulnerability reward programs. See https://www.google.com/about/appsecurity/ for more information. We unfortunately do not have information we can share about the average response time, but we can share that we receive thousands of such leads every year and pay out millions of dollars to researchers for their submissions.

Question 7. How does Google reward researchers and developers who flag privacy and security issues?

Google has a close relationship with the security research community, and has maintained a Vulnerability Reward Program (VRP) for Google-owned web properties since 2010. Through its VRP, Google encourages researchers to promptly submit reports about any design or implementation issue that substantially affects the confidentiality or integrity of user data. After the bug is confirmed and remediated, a panel of Google's security team reviews and considers the impact of the reported issue and chooses a reward accordingly. Rewards for qualifying bugs can exceed $30,000. In our latest VRP annual report from February 2018, we noted that the program had paid out nearly $12M in rewards to date.

Question 8. How does Google engage in research and community outreach on these issues?

Response. In addition to the VRP described above, Google maintains a permanent privacy and security research team that is dedicated full time to researching privacy and security issues. This research serves both to inform the teams building products about important privacy and security issues, as well as to engage and contribute to the vibrant research community, and is frequently published and presented in external journals and conferences. These teams also engage directly with users through user experience studies, to ensure that our products and policies are built with users in mind and based on their feedback.

Question 9. In July, we asked Google to provide a list of all instances in which Google has suspended an app for failing to comply with Google's policies, with an explanation of the circumstances for each. Google's written response to our letter did not provide this information.

a. How many apps did Google find to have violated its policies over the last two years? Please provide a list of these apps and describe the nature of the violation found by Google.

b. How many apps has Google suspended over the last two years? How many of these apps were suspended for violating Google's policies?

c. How many apps have been denied access by Google over the last two years? How many of these apps were denied access for violating Google's policies?

d. How many apps did Google provide warnings about over the last two years? How many of these apps violated Google's policies?

e. Did Google subsequently "un-suspend" or restore access to any of these apps? If so, why? Please identify any such apps and explain the circumstances.

We answer questions 5.a through 5.e together.

Response. In 2017, Google took down over 700,000 apps that violated Google Play policies. This was an approximately 70% increase over the number of apps taken down in 2016. This increase was possible due to significant improvements in our ability to detect potential abuse — such as impersonation, inappropriate content, or malware — through new machine learning techniques.

Our agreements with developers outline a range of actions we may take in the event of policy violations, depending on the severity of the violation. For example, we may disable the relevant app or remove it from user devices where it could cause serious harm to the user's device or data; reject, remove, suspend, or demote applications on the Play Store; or ban developers from the Play Store completely. When a developer engages in repeated or serious violations of our policies, such as developing malware or other apps that may cause user or device harm, we terminate their accounts. For example, we employ signals like app similarity and other details about developers to detect such repeat abuse. We've also developed detection models and techniques that can identify repeat offenders and abusive developers at scale. This resulted in suspending and removing the apps of 100,000 bad developers in 2017, and made it more difficult for malicious actors to create new accounts and attempt to publish more bad apps.

Google also works with the developers of apps that are not malicious but nevertheless do not meet our standards, to ensure that they improve and clarify their practices for our users. If those developers accept our recommendations, the app may ultimately be approved.

While the number of apps that are suspended for violations of our policies makes it impossible to describe the individual circumstances underlying each, we provide illustrative examples below to demonstrate how these processes work as safeguards against apps that violate our policies and to protect our users:

1. An app that helped users track another person, in violation of our Malicious Behavior policy

(https://play.google.com/about/privacy-security-deception/malicious-behavior/).

2. A game that collected personal information but lacked a privacy policy, in violation of our Personal and Sensitive Information policy (https://play.google.com/about/privacy-security-deception/personal-sensitive/)

3. An app that collected location data from users in an unsanctioned manner, in violation of our Device and Network Abuse policy (https://play.google.com/about/privacy-security-deception/device-network-abuse/)

In addition to the protection provided by Google Play, Google Play Protect provides another layer of security against malware that finds its way on the device, whether it was downloaded from Play or side loaded. As outlined in our responses to question 3(b), Google Play Protect identified and removed approximately 39 million potentially abusive apps in 2017.

Question 10. Google asserted that it tightly restricts its own employees' access to the content of users' Gmail accounts, and that Google conducts routine auditing of employee access to user email message contents.

a. What specific steps does Google take to audit employee access to user email message contents?

Response. Access to email data is restricted by technical measures to a limited number of Google employees performing legitimate business processes, such as performing debugging, responding to law enforcement requests, or reviewing spam reported by users. Additionally, we have a team of security engineers dedicated to detecting any abuse involving user data access. By design, access to email content by a Google employee is logged, analyzed, and subject to automated detection for potentially malicious behavior, such as accessing an account without an appropriate justification. On top of the automated detection, security engineers regularly analyze logs to detect new anomalous access patterns to investigate.

Question 11. Has Google detected any improper employee access to user email message contents? If so, please describe the circumstances and explain any steps taken by Google to address such improper employee access?

Response. As we described above, our first goal is always to prevent such misuse via technical and policy means, but a small number of employees must have access to fix issues with our systems.

We have strict policies to ensure that our employees sufficiently protect our users' data. While we are not aware of any misuse by an employee in the last five years, to the extent we do identify misuse by an employee, we will take strong action, including termination.

Question 12. Our July letter asked whether Google was aware of any instance of an app developer sharing Gmail user data with a third party for any purpose. Google responded

that it allows developers to share data with third parties "so long as they are transparent with the users about how they are using the data."

a. What specific steps does Google take to ensure that app developers do not improperly transfer Gmail user data to third parties?

Response. We support our policies on third party access to Gmail user data with verification, monitoring, and enforcement. In addition to the measures described in our previous response, Google's proactive review of apps seeking access to user data also include the use of machine learning tools to detect signals indicative of malicious apps. Depending on the results, a developer's history, and user feedback, we identify apps that need additional manual review. This review can include testing their app directly and reviewing their website materials, among other investigative steps.

In addition, we recently announced even stronger privacy controls. These controls include an improved user permission flow that provides a finer-grained ability to choose what data they share, limiting the types of apps that can request access from Gmail users, and imposing new requirements on how developers must treat Gmail data. These policy changes are going into effect on January 9, 2019.

More specifically, beginning in January 2019, we will only allow specific types of applications — such as email clients and productivity tools (the new policy is available at https://developers.google.com/terms/api-services-user-data-policy#additional-requirements-for-specific-api-scopes) — to access certain Gmail APIs. When users grant Gmail access to applications that do not require regular direct user interaction (for example, services that provide background reporting or monitoring to users) users will be provided with additional warnings and be required to re-grant access at regular intervals.

We are also continuing work to ensure compliance with our policy that developers should only request access to information they need. During application review, we will be tightening our review for compliance with this existing policy. For example, if an app does not need full or read access and only requires send capability, we require the developer to request narrower permission scopes so the app can only access data needed for its features.

Finally, our new policies include additional, strict limitations on how data may be used. Apps accessing these APIs can only use Gmail data to provide prominent, user-facing features and may not transfer or sell the data for other purposes, such as targeting ads, market research, email campaign tracking, and other purposes unrelated to these features. As an example, with a user's permission, consolidating data from a user's email for their direct benefit, such as expense tracking, is a permitted use case. However, consolidating the expense data for market research that benefits a third party is not permitted. We have also clarified that human review of email data must be strictly limited.

Question 13. Has Google ever suspended an app for improperly transferring Gmail user data to third parties? If so, please provide a list of these apps and describe the nature of the violation found by Google, including any action Google has taken to recover data.

Response. As one example of how Google enforces its policies, in June 2018, we identified an app in our verification and review process that appeared to imitate another legitimate company. The deceptive app claimed to make sending and receiving emails easier. Our review identified that the app had no such apparent functionality.

In addition, the app exhibited numerous suspicious signals. For example, the app's login page simply redirected users to their Gmail page and was otherwise non-functional. The app also claimed to have a demonstration page for users that did not actually exist. The app's request for verification was rejected.

Question 13. In the event of a security lapse involving user data, how does Google determine what constitutes a "significant risk of harm" for its users?

Question 14. As you may know, there is ongoing discussion as to whether there should be some disclosure requirement in the event of a security lapse or vulnerability involving user data, even if it does not constitute a breach or create a "significant risk of harm." Does Google believe some form of disclosure, such as public notice or notice to a relevant regulator, would be appropriate, even if there is not a significant risk of consumer harm?

We answer Questions 13 and 14 together.

Response. Google operates a robust incident response program for privacy and security-related events. Under this program, employees are required to report any suspected security or privacy incidents to our dedicated 24x7x365 worldwide incident response teams, so that we can respond, including by securing and protecting users' data and handling user notifications. Risk of harm to users is an important criterion for determining whether notification is appropriate. We do, for example, consider whether the type of data at issue is particularly sensitive or whether there is evidence of misuse. But we go beyond that and any legal obligations by applying several considerations focused on our users. These include whether we could accurately identify the users to inform, whether there was evidence of misuse, and whether there were any actions a developer or user could take in response.

While notification is often the right response, it is also important to avoid over-notification, which could impair users' ability to recognize and take action upon the most important notifications. We are acutely aware of the importance of the trust our users have in us. That is why we have — and will — continuously examine our approach to user notifications, always with a focus on the user.

Question 15. Google's Framework for Responsible Data Protection Legislation advocates for legislation that would, among other things: 1) require organizations to take responsibility for the use of consumers' personal data (i.e., data that can be linked to a person or personal device), 2) mandate transparency and help consumers be informed, and 3) require organizations to secure consumers' personal data and expeditiously notify individuals of security breaches that create a significant risk of harm.

a. How does Google's decision not to disclose the vulnerability that potentially exposed the personal data (including name, email address, profile photos) of nearly 500,000 Google+ users comply with the principles reflected in the Framework?

Response. When we become aware of a potential incident, we always review our legal notification obligations and determined whether a notification is required. But we always go further at Google — looking beyond our legal obligations and applying several considerations focused on our users in determining whether to voluntarily provide notice beyond what the law may require. These include whether we could accurately identify the users to inform, whether there was evidence of misuse, and whether there were any actions a developer or user could take in response. Here, with respect to the G+ bug, the answers to each of those considerations was no, and we decided against notification.

- We did not find evidence of misuse or user harm. As discussed above, we undertook a number of steps in an attempt to determine whether the developers who may have accessed non-public profile data because of this bug abused that access in any way.

- We could not accurately identify affected users. As discussed above, we have not been able to identify the set of specifically affected users and therefore do not know how many users were actually affected or who they are.

- There was nothing users or developers could do in response to notification. Finally, we considered whether this was a situation where we would recommend steps to users or developers in response to the notification. Once we patched the bug, there was nothing more that we could identify that could be done to mitigate the consequences of the bug. Indeed, once we fixed the bug, developers would have had no way after the fact to identify which of the data they accessed may have been non-public at the time they accessed it nor would we be able to confirm any users' claim that their data was always set to private.

Finally, giving notification in these types of situations, without being able to even tell a user they were affected—frustrates users and contributes to breach notification fatigue, where users begin ignoring important warnings because they are overwhelmed but the number of notifications they receive. We balance that risk with our desire for transparency and our desire to ensure the continued long-established trust of our users.

All of these factors weighed against voluntary notification, and so we made a considered decision not to provide one. Given that there was no notification to users we also did not discern a reason to make a separate notification to Congress. With that said, we are always looking to improve and will continue to look at our approach to user notifications, always with a focus on the user.

Question 16. How did Google determine that the potential exposure of nearly 500,000 Google+ users' personal data, including name, email address, date of birth, and profile photos, to outside app developers without their consent did not warrant notification of the affected consumers?

Response. First, it is important to understand the limited scope of data that was at issue. The bug allowed apps to potentially access nonpublic profile information that had been shared with a user by another G+ user.

This data was limited to profile fields including name, profile photo, occupation, and gender. (The full list of G+ Profile fields that could be fetched by the relevant API is documented on our developer site at

https://developers.google.com/+/web/api/rest/latest/people.) It is important to note, however, that many of these profile fields are or may be set to public, and therefore may not have been implicated by the G+ bug at all. Additionally, the fields of data potentially at issue did not include data like Google+ posts, messages, Google account data or G Suite content, nor did it include information about a person's home address, phone numbers, or other types of data typically used for identity theft.

As described above, whenever we become aware of a potential incident, we review our legal obligations. In the course of this review we are reminded of the standards that legislators and regulators around the world have deliberately chosen to implement after having carefully weighed the potential benefits and harms to users of notice following incidents with certain features. If we are legally obliged to give notice, we do so. But we always go further than that — looking beyond our legal obligations and applying several considerations focused on our users in determining whether to voluntarily provide notice even under circumstances where doing so is not legally required.

Those include whether we could accurately identify the users to inform, whether there was evidence of misuse, and whether there were any actions a developer or user could take in response. Here, the answer to each of those considerations was no, and we decided against notification.

Question 17. In its October 8, 2018 blog post about this issue, Google stated that it found no evidence that any of the potentially exposed consumer data was misused. What level of confidence does Google have in this assessment? Does Google have a way to verify that none of the affected consumers' data was impermissibly accessed?

Response. While we had limited data to analyze, based on what we did review, we are confident that the misuse of this data was highly unlikely.

We undertook a number of steps in an attempt to determine whether the developers who may have accessed non-public profile data because of this bug abused that access in any way.

- First, we surveyed our access logs to identify whether the developers appeared to be making API calls that appeared unusual, for example large volumes of calls. We saw no

evidence of an unusually large volume of calls from an app, which, had it occurred, may have suggested that the developer knew it was accessing nonpublic data and was taking advantage of that.

- Second, we took a harder look at the top developers accessing the API to determine whether there was any reason to believe they were not calling the API for a legitimate purpose, or likely misusing the data in some way. The vast majority were apps created by well-known and reputable companies or apps that had already undergone vetting by our teams to be allowed to participate in other developer programs.

- Third, we reviewed individual apps to determine whether their use of the API seemed appropriate given their apps' functionality. The apps were generally the type of apps that would have a legitimate purpose to use the API, e.g., apps that connect users with their friends, manage their social media profiles, or connect individuals through enterprise apps designed to connect employees within others in their organisation.

- Finally, we have also undertaken a review to determine whether there was any reason to believe the developers were engaged in misuse or deception (for example, whether those apps were known to us to be subject to regulatory inquiries or were otherwise publicly identified as misusing user data). We found no reason to be concerned.

Question 18. Has Google received any complaints from affected consumers surrounding its decision not to notify them of this potential exposure of their personal information?

Response. We are committed to addressing user concerns on all privacy and security topics, including this one. We are both proactively and reactively reaching out to users to help them understand the G+ bug, and have yet to identify a specific case of user harm.

Responses to Written Questions Submitted by Honorable Jerry Moran to Keith Enright

Question 1. The GDPR included a "data portability" requirement that allows consumers to request and receive their personal information from companies in a structured, commonly-used and machine-readable format that can be imported by competing companies and services. Your testimony indicated that Google supports the idea of "data portability," and even enables consumer data export from its variety of products. Could you please explain what compliance and enforcement with this GDPR provision looks like? Please describe the consumer benefit of this requirement.

Response. Google strongly supports the notion that users should be able to export the personal information they have provided to an organization, in a format that allows them to understand the information, store a local copy, and/or to import it into another provider's systems. This has two important consumer benefits. First, it empowers individuals to understand and control their personal information. It also keeps the market innovative, competitive, and open to new entrants by allowing users to easily move to new services without losing the benefit of their accumulated data.

The GDPR is only now entering the earliest stages of enforcement, and we cannot tell precisely how this provision will be interpreted and enforced, but we believe our program meets or exceeds the law's requirements.

More generally, we have worked on portability for over a decade and were the first to offer a portability tool in 2011. We updated and broadened this tool, Download Your Data, last spring so that it now covers more products and data types. The tool allows users to take personal information about them stored in more than 50 Google products, including search queries, Gmail messages and contacts, YouTube videos, and many others. The output is provided in formats designed to be importable into software on the user's own devices or other services.

The ability for users to transfer data directly from one provider to another, without downloading and re-uploading it, is a significant advancement in making portability practical for users all over the world. We are working with partner companies on the Data Transfer Project (https://datatransferproject.dev/), an open-source initiative to expand this capability and make it even easier for users to try a new service or otherwise control their data. The current partners (Google, Microsoft, Twitter, and Facebook) are working on building a user interface as well as bringing new and more diverse partners into the project.

Question 2. Would you expect issues of interoperability to arise for companies aiming to comply with this requirement, especially for smaller businesses that have less resources to change their data practices and equipment?

Response. Our proposed privacy framework suggests applying general principles in ways that reflect the different resources of different organizations. The overall touchpoints should be accountability and preventing harm, rather than inflexible one-size-fits all rules. Accordingly, we urge the Committee to explore ways to develop the data portability principle to work for businesses of all types and sizes.

One way to further this goal is for industry organizations and government entities like the Federal Trade Commission to explore best practices and methodologies that can be adopted by smaller players — perhaps via open-source projects or other low-cost options. We are working on this already: we recently launched the Data Transfer Project with several industry partners. As we described above, it's an open-source project that provides tools for any company, big or small, to build direct service-to-service data portability.

Question 3. Efforts to draft meaningful federal legislation on consumer data privacy will heavily rely upon determinations of what types of personally identifiable data are classified as "sensitive" and what are not. While some have suggested that expanded FTC rulemaking authority is necessary to flexibly account for new types of data sets coming from innovative technologies, I have concerns that excessive rulemaking authority could lead to frequent reclassifications of the types of data with ensuing liability adjustments. Do you have suggestions on how to best identify "sensitive" personally identifiable information?

This concern is valid, and should be considered when drafting regulatory proposals. Our regulatory framework suggests that "sensitivity" of personal information should be tied in law to risk of harm to individuals and communities, rather than a specific list of data types that might quickly become out of date. We think this is the right approach, but does require thought to avoid unnecessarily shifting regulatory standards.

One possibility is to ensure that regulatory authority over this issue is closely bound to an articulation of risk of harm. While regulators may have the ability via rulemaking or other process to define certain data types that meet this criteria, they must tie such rules to findings that those data types present such a risk of harm.

Question 4. NTIA issued a request for comment on ways to advance consumer privacy without harming prosperity and innovation. I commend the administration for their attention to this important issue. The "High Level Goals for Federal Action" that NTIA is seeking comments for includes interoperability and the development of a regulatory landscape that is consistent with the international norms and frameworks in which the U.S. participates. How do you foresee federal legislation affecting cross-border data flows?

Response. A comprehensive federal data protection law would help promote and sustain U.S. global leadership around the free and open Internet, including promoting cross-border data flows. Digital trade has become an engine of economic growth for large and small businesses around the world, and the flow of data now contributes more to GDP growth than the flow of goods.

Some countries have taken steps to limit cross-border data flows through forced data localization requirements. Such requirements fail to recognize the way that modern distributed networks function and could have the unintended consequence of weakening privacy and security protections (https://www.blog.google/products/google-cloud/freedom-data-movement-cloud-era/). A comprehensive federal data protection law that eschews data localization would serve as a bulwark against data localization requirements and lend credence to the idea that countries can protect privacy on a cross-border basis without compromising key digital trade principles. A federal law could also build on recent steps taken by the US, Mexico, and Canada in the

USMCA to require protection of the personal information of users of digital trade and to promote compatibility between different privacy frameworks. As NTIA recognized in its request for comments, it is important to promote a regulatory landscape that is consistent with international frameworks for protecting privacy, including the APEC Cross-Border Privacy Rules System.

Question 5. Also included in NTIA's request for comments, how should the U.S. government encourage more research and development of products and services that improve privacy protection?

We believe the federal government has an important role to play in enabling the development of privacy and security enhancing technologies.

We encourage the federal government to continue providing funding for the research and development of products, services, and techniques that improve privacy and security protection. Basic research remains cost intensive and educational institutions and research organizations need sustained funding to make the critical long-term investments that lead to new and improved ways to protect privacy and security. However, in its support, the government should not only focus only on the products and services that consumers see as an end-result, but also on expanding the types of tools and training available to practitioners. For example, techniques for internal data management and expanded availability of ethics training in schools can promote better outcomes for consumers.

The government should also consider establishing local centers of excellence for privacy and security research and applications, perform privacy and security research at government labs and agencies, create frameworks and mechanisms to facilitate public-private sector collaboration, and explore incentives for researchers who receive public funding to explore priority areas of research. Google has long supported open-source research, and we encourage open access to publicly funded research.

In addition the U.S. government should leverage its convening power to disseminate best practices to ensure that every organization that processes personal data, including the government itself, can keep abreast of and implement the state of the art. Publications, public events, technical workshops, digital literacy programs, and advisory committees, are potential ways the government could achieve this goal.

Lastly, the U.S. government should leverage its convening power to disseminate best practices to ensure that every organization that processes personal data, including the government itself, can keep abreast of and implement the state of the art. Publications, public events, technical workshops, digital literacy programs, and advisory committees, are potential ways the government could achieve this goal.

Question 6. As GDPR includes requirements like the "right to portability" and the "right to be forgotten," it is clear that these provisions aim to promote the consumer's ownership of their data by requiring companies to abide by their requests to permanently delete or transport their personal data to another company. However, how are these concepts enforced when the

consumer's data is submitted as an input to one or multiple proprietary algorithms employed by the company?

Response. The GDPR requirements of 'right to portability' and 'right to be forgotten' are separate concepts and we would encourage policy makers to consider them apart from each other.

With regard to user control, we believe that individuals should retain control over personal information, including when used as an input into proprietary machine learning or other algorithms.

However, we take into account privacy design principles in our development and deployment of machine learning or other algorithms, including to first see if the algorithms can be effective using anonymous data rather than personal information. For example, Google can input an aggregate of users' search queries into algorithms to learn about which search results are most relevant to which queries, without the need to include specific user information or store outputs in a way connectable to a specific user.

If the function of the algorithm does require the use of personal data, we aim to provide the user with transparency and control. So, for example, Google's algorithms use a specific user's search history (if their settings permit it) to predict the best search results for that user. Such uses can be managed, in Google's case, through easy-to-use tools for individuals to delete data stored in their account. This will cause future predictions to exclude the deleted data.

Methods like these can enable systems to continue to work and innovate while still keeping individuals in control.

Question 7. Are the outputs of the company's algorithm decidedly the consumer's personal information and required to be deleted or transported at the request of the consumer? If so, do these requirements remain the same if the data outputs are anonymized?

Response. At Google, we view all information tied to an identified individual as "personal information", whether it is information they provided to us, information our systems associated with them, or outputs of our algorithms. For example, our advertising systems sometimes attempt to determine topics of interest for a signed-in user based on his or her activity. These results are available to see, change, and delete in Ad Settings, and we consider them personal information.

Question 8. Since companies often use aggregated data outputs to study and improve their existing algorithms, services, and products, what impacts do you expect these vague GDPR requirements to have on companies' abilities to innovate?

Response. The GDPR helpfully excludes data that is no longer capable of being associated with an individual. It also creates specific exemptions for "pseudonymized" data for research purposes. While we will learn more about how these provisions will be interpreted, we generally think this principle is the right one: the law should encourage organizations to store and use data in the least identifiable manner that is compatible with the purposes for which it collected it.

Question 9. In July 2018, I joined my colleagues Senators Thune and Wicker in a letter to Google requesting more information on the data privacy practices of their Gmail service, and more specifically, third party app developers' access to email contents. In response to questions posed in our letter, Ms. Susan Molinari, Vice President of Public Policy and Government Affairs of Americas Google Inc., indicated that the verification process of third-party web apps that request access to sensitive data, such as the contents of Gmail message, undergo manual reviews of the app's privacy policy and the "suitability of the permissions the app is requesting." Will you please further explain the specific considerations of the manual review process as it determines the "suitability of the permission the app if requesting?"

Response. As we describe in our previous response, developers that request access to sensitive data, like Gmail data, must complete a verification process, described at https://developers.google.com/apps-script/guides/client-verification. This process is designed to prevent apps from misrepresenting themselves to users or accessing data that they do not need in order to perform their function. That process involves a manual review of the app's privacy policy to ensure that it adequately describes the types of data it wants to access and a manual review of the suitability of permissions the app is requesting compared to its functionality.

Google's proactive review also includes the use of machine learning tools to detect metadata signals that could indicate an app is malicious. Depending on the results and a developer's history and user feedback, we identify apps that need additional manual review for verification. This review can include testing their app directly, reviewing their website materials, among other investigative steps.

In addition, we are launching stricter limits through our appropriate access policy. Starting in January 2019, we will only allow specific types of applications — such as email clients and productivity tools (the new policy is available at https://developers.google.com/terms/api-services-user-data-policy#additional-requirements-for-specific-api-scopes) — to access certain Gmail APIs. In addition, when users grant Gmail access to applications that do not require regular direct user interaction (for example, services that provide background reporting or monitoring to users) users will be provided with additional warnings, and we will require them to re-grant access at regular intervals.

We are also continuing work to ensure compliance with our policy that developers should only request access to information they need. During application review, we will be tightening our review for compliance with this existing policy. For example, if an app does not need full or read access and only requires send capability, we require the developer to request narrower permission scopes so the app can only access data needed for its features.

Finally, our new policies add strict limitations on how data may be used. Apps accessing these APIs can only use Gmail data to provide prominent, user-facing features and may not transfer or sell the data for other purposes such as targeting ads, market research, email campaign tracking, and other unrelated purposes. (And Gmail users' email content is not used by Google for ads personalization.) As an example, with a user's permission, consolidating data from a user's email for their direct benefit, such as expense tracking, is a permitted use case. However,

consolidating the expense data for market research that benefits a third party is not permitted. We have also clarified that human review of email data must be strictly limited.

Question 10. Ms. Molinari's response also indicated that Google's Security Checkup Tool described in the letter would flag unverified apps for users. From the descriptions and graphics provided in the response, it remains unclear exactly the granularity of information that is relayed to the consumer in checking on the status of a third-party app and its access to their data. For instance, does third-party access to sensitive information to Gmail, Google Calendar, Google Contacts, and Google Hangout allow the third-party to retain the data for a certain period of time? If the consumer opts-out of sharing sensitive information with the third-party shown on the Security Checkup Tool, is that information deleted immediately, and if not, how long is it retained before deletion.

As described in our previous responses and in our Help Center (available at https://support.google.com/accounts/answer/3466521?hl=en), to help users safely share their data, Google lets them give third-party sites and apps access to different parts of their account. By visiting their account permissions page (available at https://myaccount.google.com/permissions) or using Security Checkup, users can review and control all apps that have access to their account, including viewing exactly which permissions each app currently has. If a user gives account access to a site or app they no longer trust or otherwise want to remove, they can remove its access to their Google Account at any time. That site or app won't be able to access any more information from the user's Google Account, but the user may need to request that the third party delete the data they already have.

Question 11. In the response, Ms. Molinari explained the most common reasons for Google suspending or removing third-party apps' access to Google customers' sensitive information should they fall out of compliance with Google's policies. Our original inquiry in July requested a list of all instances in which Google has suspended an app in this way, with an explanation of the circumstances for each. Will you please provide this list in your written response for the committee record?

Response. We support our policies on third party access to Gmail user data with verification, monitoring, and enforcement. In addition to the measures described in our previous response, Google's proactive review of apps seeking access to user data also include the use of machine learning tools to detect signals indicative of malicious apps. Depending on the results, a developer's history, and user feedback, we identify apps that need additional manual review. This review can include testing their app directly and reviewing their website materials, among other investigative steps.

In addition, we recently announced even stronger privacy controls. These controls include an improved user permission flow that provides a finer-grained ability to choose what data they share, limiting the types of apps that can request access from Gmail users, and imposing new requirements on how developers must treat Gmail data. These policy changes are going into effect on January 9, 2019.

More specifically, beginning in January 2019, we will only allow specific types of applications — such as email clients and productivity tools (the new policy is available at https://developers.google.com/terms/api-services-user-data-policy#additional-requi rements-for-specific-api-scopes) — to access certain Gmail APIs. When users grant Gmail access to applications that do not require regular direct user interaction (for example, services that provide background reporting or monitoring to users) users will be provided with additional warnings and be required to re-grant access at regular intervals.

We are also continuing work to ensure compliance with our policy that developers should only request access to information they need. During application review, we will be tightening our review for compliance with this existing policy. For example, if an app does not need full or read access and only requires send capability, we require the developer to request narrower permission scopes so the app can only access data needed for its features.

Finally, our new policies include additional, strict limitations on how data may be used. Apps accessing these APIs can only use Gmail data to provide prominent, user-facing features and may not transfer or sell the data for other purposes, such as targeting ads, market research, email campaign tracking, and other purposes unrelated to these features. As an example, with a user's permission, consolidating data from a user's email for their direct benefit, such as expense tracking, is a permitted use case. However, consolidating the expense data for market research that benefits a third party is not permitted. We have also clarified that human review of email data must be strictly limited.

As one example of how Google enforces its policies, in June 2018, we identified an app in our verification and review process that appeared to imitate another legitimate company. The deceptive app claimed to make sending and receiving emails easier. Our review identified that the app had no apparent functionality. In addition, the app exhibited numerous suspicious signals. For example, the app's login page simply redirected users to their Gmail page and was otherwise nonfunctional. The app also claimed to have a demonstration page for users that did not actually exist. The app's request for verification was rejected and the app was suspended from requesting user data.

Responses to Written Questions Submitted by Honorable Shelley Moore Capito to Keith Enright

Question 1. According to a study by Pew Research, only 38% of consumers know how to limit what information they give online. Consider me among those consumers who do not know what is being collected and how to keep my information to myself. Even with privacy settings and assurances that my data is not being collected and used without my consent, I still have concerns.

I believe the root of this issue is transparency and consumer confidence. What are your companies doing to increase the transparency when it comes to the type of data you collect?

Response. Transparency is a core principle, and we provide users with clear, simple explanations of what we collect and our use of it. We realize privacy policies aren't user's first choice of reading material, but we worked to make ours best-in-class, with illustrations, videos and other interactive content designed to convey key concepts and choices.

But we go beyond transparency to try to really help users understand, in real time and in context as they use our services. We gave some examples in my testimony, and just last week added new transparency in our flagship product, Search, that shows users exactly how their data is being used to improve their search results, along with direct access to controls. (https://www.blog.google/technology/safety-security/making-it-easier-control-your-data-directly-google-products/)

Question 2. What difficulties have your companies faced when developing more transparent privacy policies?

Response. We value being transparent with our users about how and why we use data to operate our business. Making this information available is critical to building and maintaining user trust, and specifically helps our users make important decisions about their privacy. One challenge we have encountered is how to ensure users have the information they need, without overwhelming them with extraneous details. We are constantly refining this balance based on feedback from our users.

We recently updated our privacy policy to incorporate some of the insights we have gained.

While we don't solely rely on the privacy policy for that purpose, we still want to make our policies understandable and accessible to users who spend the time to review them as well as being full and complete statements of our data practices for experts and regulators to hold us accountable. We try to meet both needs, via clear headings, easy navigation, overlays and examples, and explanatory videos.

We regularly conduct surveys and interviews with users to inform our approach and ensure we strike this balance effectively.

Question 3. West Virginia has a high elderly population that is rapidly increasing as baby boomers retire. I am positive that a lot of my elderly constituents are among those individuals who do not know how to limit their online information.

What are some of the measures your companies are doing to teach consumers – and specifically older consumers – about what data they share on your platforms?

Response. Google invests directly and through partnerships with expert organizations to inform individuals about the data they share on our platforms, and ways to protect their privacy and security. In addition to our efforts on transparency mentioned above, we offer Privacy Checkup (https://myaccount.google.com/privacycheckup) and Security Checkup (h ttps://myaccount.google.com/intro/security-checkup). These tools are helpful for all our users, but we believe seniors who feel less comfortable with online services could particularly benefit from the guided review of their privacy and security settings.

In terms of partnerships, we support organizations that provide general audience and senior-specific digital literacy information, through financial assistance, take-home kits that include security keys for two-factor authentication, and general privacy and security advice. For example, we helped support ConnectSafely's development of the Senior's Guide to Online Safety (https://www.connectsafely.org/seniors/).

Project GOAL (Getting Older Adults Online) is an organization dedicated to helping older adults access broadband and address barriers like digital literacy to making that possible, safely. We have provided financial support to Project GOAL to support their work to educate older consumers and to generally raise the profile of this important work.

Question 4. I know advertising through data collection has a monetary value, and appreciate the business model, however, I find it hard to know what is being collected and how I can keep my information to myself. Even with privacy settings and assurances my data is not being used without my consent, I still have concerns.

Please explain how your business model allows both data to be used to make suggested recommended purchases on your site? As well as how you use that data to target ads to consumers? And how do you do that while protecting personal data?

Response. The online advertising model enables advertisers to reach audiences who are more likely to be interested in purchasing their products or services, and therefore to waste less of their marketing budgets on audiences who are not. This model has lowered barriers to entry for scores of small businesses, enabling them to compete for access to global markets.

At Google, supporting this economy while promoting the privacy and security of user data is essential. Crucially, users' personal information stays within Google and is not shared with or sold to advertisers. Advertisers instead have access to our services through dashboards and other interfaces that enable them to decide how to show ads to aggregated audiences with certain characteristics.

While providing free, ad-supported services, we remain committed to putting users in control of their privacy, so we are constantly improving our privacy disclosures, settings and controls. Users can opt out of personalized ads via Ad Settings and the AdChoices industry program, via a notice served in every ad Google shows.

Our Ad Settings page not only allows you to turn off targeted ads; it also shows you what data we use to personalize ads, as well as the topics and advertisers we think you are interested in (and why we think you are interested).

Our industry-leading policies prohibit the use of sensitive categories for ad targeting. As we explain to users in our privacy policy, we do not show personalized ads to people based on sensitive categories, such as race, religion, sexual orientation, or health conditions. And we don't allow advertisers to use these sensitive interest categories to select audiences for their ads.

Building systems that users trust is essential to the continued success of Google and the Internet at large. So, we keep the information we collect confidential and under users' control, and work every day to maintain that trust.

Question 5. How can Congress ensure that data collected is used responsibly without shutting down the collection of data completely?

Response. While companies and regulators like the FTC are doing a lot today to protect privacy and security, we think it is has long been appropriate to adopt a comprehensive data protection law in the United States. Our model framework includes a set of principles, based on established regimes like the Fair Information Practices Principles (FIPPs), OECD Privacy Principles, Asia-Pacific Economic Cooperation (APEC) Privacy Framework, and aspects of the European General Data Protection Regulation (GDPR). We also bring our 20 years of experience offering services that depend on information, privacy protections, and user trust.

We specifically urge a law that requires organizations to be responsible to individuals whose data they collect:

Organizations must operate with respect for individuals' interests when they process personal information. They must also take responsibility for using data in a way that provides value to individuals and society and minimizes the risk of harm based on the use of personal information.

We think an approach like this can bolster trust and use of online services in a way that still encourages new and innovative uses of data.

Question 6. In April, the European Union (EU) passed the General Data Protection Regulation (GDPR) in order to protect personal data and uphold individual privacy rights. These new regulations have created uncertainty for U.S. firms, despite several already coming into compliance.

Response. Innovation is important to small businesses, especially in rural America. The new European standards have created massive hurdles for these businesses to be in compliance. Many small companies in Europe are already expressing an inability to afford the legal consequences. For example, if a rural grocery store advertises online and provides a link to coupons. Under the GDPR compliance rules, this simple practice can result in expensive legal consequences.

Question 7. For those who do business in Europe, do you think GDPR has the potential to have negative impacts on rural small businesses in Europe?

Response. We do worry about the impact of data protection regulation on small businesses and new entrants, who lack the resources to build complex compliance programs like Google has built. This would be a bad outcome for everyone.

This difficulty can be managed, in our view, with flexible regulation that focuses on the principles of accountability and protecting users from harm. There should be many ways to demonstrate accountability, which scale with the size and scope of the organization. For example, small businesses should have access to industry best practices, or open-source projects for obligations like data portability, that can allow them to meet the requirements of the law without significant cost or expertise.

Question 8. California has already passed a sweeping consumer protection law that threatens established business models throughout the digital sector. I appreciate the industry taking the initiative in creating a framework, in addition to the privacy principles released by the US Chamber of Commerce.

Response. As we begin discussing the appropriate position of the federal government, can you describe what actions we should investigate more closely for any potential national framework?

As you consider creating a federal framework, we hope you will also take into account our model framework, as we referenced above, which sets out substantive requirements and enforcement and scoping principles. An important component of our framework is the principle that new privacy regulations should apply across the board, to all sectors of the economy.

Question 9. Who, in your opinion, is the appropriate regulator to oversee any framework and why?

Response. Many different regulators have been involved with data protection in recent years, and they all have an important role to play in the future. In particular, the Federal Trade Commission has been the primary federal privacy regulator, and have built a track record of strong enforcement and deep expertise on this issue.

Question 10. According to recent research by Magid, a media research firm, 35% of millennials share their password to access streaming services. I certainly understand that the terms and conditions of these services already note that access is for personal use and not to be shared with others. And that the account holder remains responsible for the actions of that third party. However, as the number younger generations sharing their password grows so has the potential for abuse. This "overly sharing of passwords" and the younger generation operate differently than many my age.

Are your policies flexible to cover a third party that may use a friend's or spouse's password? Is this something we should consider as we create federal guidelines?

Response. We strongly discourage password sharing which can create serious security and other problems for those involved. We have mechanisms for families to share content via their

Google Family Group, which allows users to share apps, movies and books through the Play Family Library or subscriptions like YouTube Premium and Google Play Music.

For shared devices like Google Home, others in your home can request music from the device if you have logged into the device with your personal music account. We ask users for additional permissions before surfacing sensitive information like calendar entries (https://support.google.com/googlehome/answer/7684543?hl=en&ref_topic=7549809&visit_id=636764503506028117-3686497603&rd=1) or payments (https://support.google.com/googlehome/answer/7276665).

We created Be Internet Awesome (https://beinternetawesome.withgoogle.com/en), our flagship digital literacy program, to give younger users the knowledge and tools to navigate the Internet safely. One of the five lessons, Be Internet Strong, specifically teaches kids to not share passwords with anyone other than parents or guardians. We worked with the Family Online Safety Institute and ConnectSafely as well to build a program that aims to encourage parents, educators and kids alike to exhibit all the traits that comprise "Awesome" online: to be Smart, Alert, Strong, Kind and Brave.

Responses to Written Questions Submitted by Honorable Todd Young to Keith Enright

Question 1. GDPR establishes a right of data portability, which some believe is key to driving new innovation and competition within the emerging data ecosystem. Others are concerned that data portability rights, depending on how crafted, could further entrench incumbent companies.

What questions should policymakers be asking in developing data portability rights?

Response. Google strongly supports the notion that users should be able to export the personal information they have provided to an organization, in a format that allows them to understand the information, store a local copy, and/or to import it into another provider's systems. This has two important consumer benefits. First, it empowers individuals to understand and control their personal information. It also keeps the market innovative, competitive, and open to new entrants by allowing users to easily move to new services without losing the benefit of their accumulated data.

More generally, we have worked on portability for over a decade and were the first to offer a portability tool in 2011. We updated and broadened this tool, Download Your Data, last spring so that it now covers more products and data types. The tool allows users to take personal information about them stored in more than 50 Google products, including search queries, Gmail messages and contacts, YouTube videos, and many others. The output is provided in formats designed to be importable into software on the user's own devices or other services.

The ability for users to transfer data directly from one provider to another, without downloading and re-uploading it, is a significant advancement in making portability practical for users all over the world. We are working with partner companies on the Data Transfer Project (https://datatransferproject.dev/), an open-source initiative to expand this capability and make it even easier for users to try a new service or otherwise control their data. The current partners (Google, Microsoft, Twitter, and Facebook) are working on building a user interface as well as bringing new and more diverse partners into the project.

Question 2. What improvements would you make, if any, to Art. 20 of GDPR, which addresses the right to data portability?

Response. This is the very early stages of GDPR enforcement, and so we cannot tell precisely how this provision will be interpreted and enforced. We agree it will be important to continue to observe the experience of data portability under the GDPR to best learn what is working well and what can be improved.

The GDPR's portability provision applies to personal data "provided to a controller", avoiding data types like inferred data or observed data that often cannot be practically made available for download. While some observed or inferred data can be made available (and we do so in Download Your Data), it may not be appropriate to mandate such a requirement. Our model principles follow a similar line.

Question 3. How best can data portability rights be crafted to create new competition, but not further entrench incumbent companies?

Response. Our proposed privacy framework suggests applying general principles in ways that reflect the different resources of different organizations. The overall touchpoints should be accountability and preventing harm, rather than inflexible one-size-fits all rules. Accordingly, we urge the Committee to explore ways to develop the data portability principle to work for businesses of all types and sizes.

One way to further this goal is for industry organizations and government entities like the Federal Trade Commission to explore best practices and methodologies that can be adopted by smaller players — perhaps via open-source projects or other low-cost options.