Google appreciates the opportunity to share our concerns about Schedule 1 of the Telecommunications and Other Legislation Amendment (Assistance and Access) Act 2018 (“the Act”) in conjunction with the Independent National Security Legislation Monitor’s (INSLM) mandate under the Act. Over the past year, Google has participated in various coalitions of civil society groups, trade associations, and other companies (including the Open Technology Institute, Access Now, the Communications Alliance, DIGI, the AIIA and others) to raise concerns about Schedule 1 of the Act. Our comments in this submission echo and expand upon the issues broached in those submissions.

Google was involved in extensive and forthright engagement with the Australian government as this legislation moved through the legislative process. Still, we have fundamental concerns with the Act, which could significantly weaken the cybersecurity posture of companies and users alike at a time when we should be collectively reaffirming the importance of robust cybersecurity tools, technologies, and protocols. Furthermore, we maintain that a persuasive case has not been made by the government as to why this legislation is needed, particularly when companies are already working proactively with Australian agencies to address their concerns.

The changes we propose below could mitigate some provisions in the Act that threaten to weaken cybersecurity and undermine users’ trust in products and services.   Our comments focus on three discrete concerns we have with Schedule 1 of the Act: (i) certain definitions of key terms; (ii) the absence of judicial review mechanisms for technical capability notices (TCNs) and technical assistance notices (TANs); and (iii) non-disclosure provisions.

Definitions of Key Terms

The definitions of “systemic vulnerability” and “systemic weakness” create ambiguity about the types of TCNs and TANs that government officials may authorise. Below are the definitions for these key terms under the Act:

> “***systemic vulnerability*** means a vulnerability that affects a whole class of technology, but does not include a vulnerability that is selectively introduced to one or more target technologies that are connected with a particular person. For this purpose, it is immaterial whether the person can be identified.”
>
> “***systemic weakness*** means a weakness that affects a whole class of technology, but does not include a weakness that is selectively introduced to one or more target technologies that are connected with a particular person. For this purpose, it is immaterial whether the person can be identified.”

By their own terms, a “systemic vulnerability” and “systemic weakness” both encompass vulnerabilities and weaknesses that affect a whole class of technologies and yet exempt any such vulnerability or weakness “that is selectively introduced to one or more target technologies that are connected with a particular person.” In effect, the exception swallows the rule, and therefore effectively negates the prohibition in Section 317ZG on requiring providers to implement or build systemic vulnerabilities or weaknesses. For example, a TCN that requires a company to engineer a weakness or vulnerability into its software update infrastructure would necessarily create a systemic vulnerability and weakness because once a weakness or vulnerability is engineered, it can subsequently be used or exploited in connection with any number of other individuals. Under the Act, however, that same TCN would purportedly not be systemic in nature so long as it is “selectively introduced to one or more target technologies that are connected with a particular person.”

Even where the introduction of such vulnerabilities or weaknesses could be easily and subsequently exploited or repurposed to target similarly situated users of a product or service, the very fact that such vulnerabilities and weaknesses are initially introduced in connection with a particular person would render such vulnerabilities or weaknesses purportedly non-“systemic”. The acute risks to the broader cybersecurity posture of service providers and users alike militate against this approach, precisely because mechanisms such as software updates are designed in many instances to be systemic for the benefit of a broad (if not entire) universe of users. A recent report entitled “Moving the Encryption Policy Conversation Forward”, authored by a diverse panel of experts assembled by the Carnegie Endowment for International Peace and Princeton University, articulates the view of some experts that selectively introducing weaknesses or vulnerabilities will still create systemic risk in this context:

> “...code updates, typically delivered by service providers over the internet, patch known flaws in software and hardware and are considered a foundation of basic cybersecurity hygiene. Companies and cybersecurity specialists worry that consumers will be less likely to accept updates - thus exposing themselves to exploitation by hackers and governments - if they are suspicious of potential government interventions through such means. Such disincentives, even if they only were to affect a percentage of users, *would have a systematically negative impact on cybersecurity that could outweigh the benefits of lawful access*.” (emphasis added).

The requirement for a weakness or vulnerability to be systemic only when it ‘*affects a whole class of technology*’ is also deeply problematic. The concept of ‘class of technology’ is undefined and unclear, but the term seems to encompass something much broader than a single service, device, platform or system. This element of the definition provides no protection against TCNs and TANs being issued to require that a product or service-wide weakness or vulnerability be built into a provider’s individual products and services. For example, a TCN or TAN which required a weakness or vulnerability to be introduced into a carriage service provided by a single provider would not introduce a ‘systemic’ weakness or vulnerability if carriage services are viewed as the relevant class of technology. The definition of systemic weakness and vulnerability should refer to those which are systemic within an individual service, device, platform or system.

While Google appreciates the limitations on TCNs and TANs that are described in Section 317ZG, those limitations only apply to the extent a “designated communications provider” (DCP) is required to create a “systemic weakness” or “systemic vulnerability”. As outlined above, if the weakness or vulnerability is selectively introduced to one or more target technologies that are connected to a particular person, or fail to affect a “whole class of technology”, it falls outside the scope of a “systemic weakness” or “systemic vulnerability”. Moreover, while the limitations described in Section 317ZG clarify that DCPs may not be required to “build a new decryption capability”, there is nothing that would prohibit the Australian government from issuing a TCN that requires a DCP to circumvent a technology’s encryption protocols, so long as its application in a particular case was sufficiently targeted. In short, the current definitions are overly narrow and allow much room for interpretation and difference of opinion in what constitutes a systemic weakness or systemic vulnerability. This renders them both inadequate as a means to protect TCNs and TANs from being used to create product or service-wide weaknesses and vulnerabilities, and difficult to work with in practice.

As articulated in previous joint submissions, the terms “systemic weakness” or “systemic vulnerability” should encompass weaknesses or vulnerabilities that may be exploited or repurposed to target other users of the same or different service, device, platform, or operating system. In many instances, the fruits of a TCN could be the subject of a future TAN - i.e., it will be relatively trivial in at least some cases to ask a DCP to provide assistance once that DCP has created a capability to exploit a weakness or vulnerability pursuant to a TCN. For example, a TCN that requires a DCP to create a vulnerability in its software update infrastructure may be initially targeted at a particular person, but it strains credulity to contend that this vulnerability would not be exploited or repurposed in future cases. The definition(s) of “systemic weakness” and “systemic vulnerability”, therefore, should reflect the dangers inherent in creating such weaknesses and vulnerabilities without recognising that they will likely be exploited in future cases.

Finally, the broad definition of “designated communications provider” (DCP) captures business-to-business (i.e., enterprise focused) infrastructure and services. This could have a negative impact on cloud service providers’ ability to deliver innovative solutions to customers in Australia, including important security and productivity products. More broadly, the legislation works against the Government’s longstanding commitment to cloud first policies by undercutting trust in technology providers, increasing costs, slowing down cloud adoption, and weakening security.

We also query the public policy objective behind including these services within the scope of this law when it is highly unlikely that offenders of serious crimes will be using enterprise platforms to communicate with other offenders. We urge the INSLM to recommend the appropriate scoping of this law in light of what the agencies know of how these offenders are communicating with each other.

Judicial Oversight and Review

There are no judicial review or appellate mechanisms in the Act, and any recourse to judicial review by a provider based on other law(s) is both unclear and potentially limited to procedural -- rather than substantive -- review. Given the broad and invasive powers conferred to the government, the Act itself should have clearly defined judicial review and appellate mechanisms, with precise standards of *de novo* procedural and substantive review.

The administrative review procedures in Schedule 1 of the Act are insufficient in light of the powers that are granted to the Australian government. Under Section 317WA, a DCP may request that the Attorney General undertake an assessment of whether a TCN should be executed. Subsection (4) of Section 317WA requires that one of the two assessors assigned to review a TCN be a former judge who served in such a capacity for at least five years.

The standard of review by which the assessors may review a TCN is notably vague. Assessors decide whether a TCN "should be given", by "consider[ing]" five enumerated factors and "consult[ing]" with three parties. However, it is unclear how much deference the assessors should give the original ministerial action or what standard the assessors must reach in order to overturn that action. Also, as noted below, there is no mechanism in the Act for appellate review of the assessors' decision.

The assessment procedures require that assessors "consider" a variety of factors that may bear upon the appropriateness of a particular TCN. While we appreciate that those factors encompass privacy and cybersecurity concerns, the Act should make those factors more concrete and practical for assessors to evaluate. The principles covering mobile phone encryption access in "Moving the Encryption Policy Conversation Forward" (i.e. law enforcement utility, equity, specificity, focus, authorisation, limitation, auditability, and transparency, evaluation, and oversight) can be useful lodestars both in the assessment phase and as part of robust judicial review mechanisms that should be incorporated into the law. Under the Act, the Attorney General remains the ultimate and sole arbiter of whether to proceed with the TCN. This is true even if serious concerns are raised in the course of conducting the assessment.

The Attorney General must "have regard" for the assessment, but "regard" is not a not a standard of review. Moreover, "regard" for the assessment falls well short of meaningful, independent judicial review; the participation of a former judge in the assessment process is not tantamount to judicial review or even judicial oversight. Absent the power to reject or approve the TCN under robust standards, the participation of a former judge is a procedural nicety, but it should not be confused for judicial review.

The absence of judicial review or oversight is stark in other places as well.

First, there is no provision for judicial review (or meaningful oversight) of TANs in the Act. The issuance of TANs can raise similar concerns to TCNs. Both TANs and TCNs may require DCPs to do specified “acts or things” that are identical or substantially similar in nature. As defined in Section 317E, these “acts or things” may include “removing one of more forms of electronic protection”, “installing, maintaining, testing, or using software or equipment”, or even “modifying...the characteristics of a service”. These “acts or things” described in Section 317 E are not trivial and are non-exhaustive. Legal procedures that conscript private companies into government service to (i) engineer vulnerabilities into their products of services or (ii) defeat security protocols that are intended to protect users should be invoked rarely, if at all, and not without separate, robust procedures for judicial review. At a minimum, there should be clear and robust standards for judicial review and oversight of TANs in the Act, in addition to TCNs.

Second, neither the issuance of TANs, nor TCNs, are subject to appellate review under the Act. DCPs, therefore, have no clear judicial remedy in circumstances where the Australian government may authorise extraordinary measures that may weaken trust, undermine sound cybersecurity hygiene, or otherwise contravene the reasonable and proportionate factors under 317JC and/or the limitations under Section 317ZG. The powers conferred to a court under the Act are limited: “a court may make such orders as the court considers appropriate in relation to the disclosure, protection, storage, handling, or destruction” of TCN and TAN information. Courts have no explicit authority under the Act to review the execution of TCNs or TANs, nor is there a meaningful and strong standard of review for such orders. Administrative appeals remain the only option for judicial review. Any such review, however, would be confined to the procedural elements of a decision, which could not reach the merits of any decision made by the Attorney General under this legislation.

Non-Disclosure Provisions

Given the broad powers under the Act and the absence of meaningful judicial review, it is particularly important to shed light on the volume, nature, and scope of TANs and TCNs that are issued under Schedule 1 of the Act. Although subsection 13 of Section 317ZF authorises DCPs to publish transparency reports that disclose the aggregate number of TANs and TCNs received during a six-month period, DCPs are prohibited from disclosing additional information about the nature of the capabilities they are being asked to develop or assistance they are being compelled to render. While there may be legitimate reasons to delay disclosure of details related to the issuance of TCNs and TANs, transparency is critical in order for the broader public to understand how the Australian government is using the authorities granted under Schedule 1 of the Act.

Relatedly, there should be clear rules that govern the issuance of such non-disclosure orders vis a vis the targets of criminal investigations. For example, as a general rule, DCPs should be permitted to notify the targets of TCNs and TANs unless such orders are accompanied by non-disclosure orders of limited duration. The Act should further describe the limited circumstances (e.g. endangering the life of an individual) under which a non-disclosure order may be sought, require judicial approval for the issuance of non-disclosure orders, and, as noted above, limit the duration of non-disclosure orders.

Finally, the Act should clarify that security researchers who discover weaknesses and vulnerabilities in products and services are not prohibited from disclosing technical information about such weaknesses and vulnerabilities. There may be circumstances where security researchers discover weaknesses and vulnerabilities that are created and/or exploited as a result of the issuance of a TCN/TAN. Particularly in light of the broad non-disclosure obligations placed on DCPs, security researchers should not be constrained from disclosing information that they independently discover.

****

Thank you for the opportunity to provide this submission in conjunction with the INSLM's mandate under the Act. We look forward to engaging with the INSLM in advance of its March 2020 report.