Proposed Principles for consumer IoT Security Transparency

1. A printed label must not imply real time assurance

Unlike food labels, digital security labels must be "live" labels, where security information is conveyed on a central maintained source, which could be the same website hosting the evaluation scheme and could be leveraged at scale by relying parties. A physical label, either printed on a box or visible in an app, should be used to the extent it encourages users to visit the further source or website (e.g. scan a QR code or click a link) to obtain the real-time status. Importantly, labels should be attestable.

Security is dynamic. Organizational security processes can change, the threat landscape is always evolving, and new vulnerabilities can be discovered at any time. If there are significant changes to a digital product's compliance or it loses its certification, the live label should reflect that. For connected device security, printed labels lacking this element of live information should be avoided. If they convey trust implicitly such as, "certified to NNN standard" or, "3 stars", they run the danger of influencing consumers to make decisions based on potentially inaccurate information. A consumer may purchase a webcam with a "3-star" security label only to find when they return home the product has non-mitigatable vulnerabilities that make it unsafe. Or, a product may sit on a shelf long enough to exceed its expected life span and thus call the initial certification into question. Security labeling programs should help consumers make better security decisions.

2. Labels must reference strong international evaluation schemes

The challenge of utilizing a security labeling scheme is not limited to the physical manifestation of the label, but also extends to ensuring the label references a meaningful security standard. To be meaningful the security label should accurately represent or describe the security/privacy capabilities, limitations, and status or posture of the device in a manner consistent with a trustworthy security/privacy evaluation scheme, such as the ones being developed by the Connectivity Standards Alliance (CSA) and GSMA. And these descriptions should be rendered in a machine readable format capable of being ingested and analyzed at scale. Both of these organizations are actively developing IoT security/privacy evaluation schemes that reference well-regarded standards, including recent IoT baseline security guidance from NIST, ETSI, ISO, and OWASP. Some important principles for industry evaluation schemes leveraged by a national labeling program include:

***Strong governance*:** The organization administering the evaluation scheme, particularly if it is a non-governmental organization (NGO), must have strong governance. For example, NGOs that house both a scheme and their own in-house evaluation lab introduce potential conflicts of interest that should be avoided.

***Strong track record for managing evaluation schemes at scale*:** Adequate planning, capacity, and security expertise are needed to conduct consistent and timely evaluations based on clearly delineated inputs using procedures that scale globally. The scheme requirements and evaluation methodology should be tech-neutral and provide consistent results across their authorized testing laboratories. Managing a high quality, global scheme is hard. National authorities have struggled at this for many years, especially in the consumer realm. An NGO that has no prior track record of managing a scheme with significant global adoption may face challenges. In the security and connected device space, CSA and GSMA have had success in managing global security schemes that have stood the test of time.

***Choice with a risk-based high quality bar*:** To maximize effectiveness, a small set of high quality evaluation schemes could act as the hub within a hub and spoke model for enabling national security

labeling schemes across the globe. Evaluation schemes should authorize a range of labs for lab-tested results, providing price and quality competition for lab engagements. Multiple schemes would encourage competition among evaluation schemes, as they too will levy fees for membership, certification, and monitoring vendor claims and should be incentivized to maintain quality over time. However, balance is key, as too many schemes could be challenging for governments and the private sector to monitor and trust. Setting a high bar for governance and track record, as described above, will help curate global security evaluation scheme choices.

***International participation*:** National labeling schemes must recognize that many manufacturers sell products across the world. A national label that does not reference security evaluation schemes that serve the global community may result in multiple inconsistent national labeling schemes that confuse consumers and could be prohibitively expensive for small and medium size product developers. Misaligned or non-harmonized national efforts may become a significant barrier to entry for smaller vendors and run counter to the intended goals of competition-enhancing policies in their respective markets. Schemes should endeavor to maximize the reuse of conformance artifacts to ease the cost to manufacturers that participate in multiple evaluation schemes.

***Assurance maintenance*:** The security evaluation schemes should provide a mechanism for independent researchers to pressure-test conformance claims made by manufacturers. Moment-in-time certifications have historically plagued security evaluation schemes, and for cost and other reasons, forced annual re-certifications are not a suitable answer. For connected consumer products, crowdsourced research can help identify weaknesses that may question a certification result. This approach has succeeded in helping maintain the security of numerous global consumer products and platforms and is especially helpful in monitoring the results of self-attestation certifications that will be needed in any national scale security labeling program. Also there should be incentives for manufacturers to offer security bug bounty programs that will leverage the skills of the security community to pressure test evaluation scheme results. These reward programs are also a great way to recruit more people into the cybersecurity field.

3. A minimum security baseline must be coupled with flexibility above it

Any standard used in association with a security labeling scheme should require a robust minimum baseline for security before the label can be applied. A minimum security baseline must further be coupled with flexibility to define additional requirements and/or levels to accelerate ecosystem improvements. Most security labeling schemes to date are focused on common sense baseline requirement standards. These standards set an important minimum bar for digital security, reducing the likelihood that consumers will be exposed to poor security practices. However, security is not a binary state. Applying a minimum set of best practices will not magically make a product free of vulnerabilities, but it is likely to discourage the most common security foibles—and as noted above, the minimum requirements for a baseline should shift over time as experience informs our understanding of risk; as technology evolves; and as threat actors adapt. Furthermore, while providing broad protection against common opportunistic attackers, even carefully developed baseline security standards may not protect against all advanced persistent threat actors. The Mirai botnet attack was successful because many digital products lack the most rudimentary security functionality: the ability to apply a security update in the field.

Over time IoT security needs to do better. IoT security evaluation schemes should be sufficiently flexible to allow for additional security functional requirements to be measured and rated across products, based on risk. For example, the current baseline security requirements do not cover things like the strength of a biometric authenticator (important for a growing range of consumer digital products) nor do they provide a standardized method for comparing the relative strength of security update policies

(e.g. a product that receives regular updates for five years should be valued more highly by consumers than one that receives updates for two years). Communities that focus on specific vertical markets of product families are motivated to create security functional requirement profiles (and labels) that go above and beyond the baseline and are more tailored for that product category. Labeling schemes must allow for this flexibility, as long as profile compliance is managed by high quality evaluation schemes.

Similarly, in addition to functionality such as biometrics and update frequency, labels need to allow for assurance levels, which inform the question of how much confidence a consumer should have in a product's security functionality claims. For example, emerging consumer evaluation schemes may permit a self-attestation of conformance or a lab test that validates basic security functionality. These kinds of attestations yield relatively low assurance, but still better than none. Today's schemes do not allow for an assessment that emulates a high potential attacker trying to break the system's security functionality. To date, due to cost and complexity, high potential attacker vulnerability assessments have been limited to a vanishingly small number of products, including secure elements and small hypervisors. Yet for a nation's most critical systems, such as connected medical devices, a higher level of assurance will be needed, and any security labeling scheme must not preclude future extensions that offer higher levels of assurance.

4. Broad-based transparency is just as important as the minimum bar

While it is desirable that security labeling schemes provide consumers with simple guidance on safety, the desire for such a simple bar forces it to be the lowest common denominator for security capability so as not to preclude large portions of the market. It is equally important that labeling schemes increase transparency in security. So much of the discussion around security labeling schemes has focused on selecting the best possible minimum bar rather than promoting transparency of security capability, regardless of what minimum bar a product may meet. This is short-sighted and fails to learn from many other consumer rating schemes (e.g. Consumer Reports) that have successfully provided transparency around a much wider range of product capabilities over time.

While a common baseline is a good place to start, it's also important to encourage the use of more comprehensive security requirement specifications developed by high-quality NGO standards bodies and/or schemes against which products can be assessed. The goal is not to mandate every requirement above the baseline, but rather to mandate transparency of compliance against those requirements. Similar to many other consumer rating schemes, the transparency across a wide range of important capabilities (e.g. the biometrics example above) will enable easy side-by-side comparison during purchasing decisions, which will act as the tide to raise all boats, driving product developers to compete with each other in security. This already happens with speeds and feeds, battery life, energy consumption, and many other device features that people care about. By developing more comprehensive transparency in the labeling scheme, consumers can learn about a wider range of security capabilities, with that awareness driving demand for device developers to do better.

Greater IoT transparency fosters conversations—not only between vendors, retailers, and consumers, but also between devices and networks. Devices that transparently describe their capabilities, limitations, and status or posture will enable networks to better understand what behaviors are expected or anomalous. This conversation helps enable networks to address serious device security flaws by potentially blocking malicious actions or through isolation and segmentation.

5. Security labeling schemes need adoption incentive

Transparency is a core concept that can raise demand and improve supply of better security across the IoT. However, what will cause products to be evaluated so that security capability data will be published

and made easily consumable?

“Voluntary” regimes at their start will attract the developers that are already doing good security work and depend on doing so for their customers and brands. The key is how to keep momentum across the IoT ecosystem to encourage even further device developers to participate. Many avenues can lead to increased economic incentives for improved security. That means a mix of carrots and sticks will be necessary to incentivize developers to increase the security of their IoT products.

National security labeling schemes should focus on a few of the biggest market movers, in order of decreasing impact:

***National mandate*:** Some national governments are moving towards legislation or executive orders that will require common baseline IoT security requirements to be met, with corresponding labeling to differentiate compliant products from those not covered by the mandate. If an IoT security national mandate is created, it would be most appropriate to include the following conditions. National mandates and labeling schemes must reference broadly applicable, high quality, NGO standards and schemes (as described above) so that they can be reused across multiple national labeling schemes. Global normalization and cross-recognition is not a nice-to-have, national schemes will fail if they do not solve for this important economic reality upfront. Ideally, government officials who care about a successful IoT national labeling scheme should encourage the NGO schemes to try to solve this problem globally.

Mandating a poor IoT security labeling scheme can do more harm than good. Nations should not create bespoke evaluation schemes, as it causes small and medium size developers to be priced out of the market due to the need to recertify and label their products across all these schemes. Non-harmonized approaches will also inhibit innovation as developers create less inclusive products to avoid nations with bespoke and non-harmonized labeling regimes.

***Retailers*:** Retailers of IoT products could have a huge impact by incentivizing and promoting baseline security standards compliance. Retailers could mandate compliance for all products listed for sale, provide visual labeling and/or search and discovery incentives for IoT products that meet the requirements specified in high quality security evaluation schemes.

***Platform developers*:** Many IoT devices exist as part of technical platforms, such as devices built on various software and app distribution platforms (i.e. Android Open Source Project (AOSP)) or interoperability standards such as Matter and Bluetooth which act as a platform certifying products that meet those interoperability standards. These platform developers may use security compliance within larger certification, compliance, and business incentive programs that can drive adoption at scale. The impact depends on the size and scale of the platform and whether the carrots provided by platform providers are sufficiently attractive.