



A First Look at Governments’ Enterprise Security Guidance*

Kimberly Ruth
Stanford University
kcruth@cs.stanford.edu

Raymond Buernor Obu
Stanford University
oburay@stanford.edu

Ifeoluwa Shode[†]
Fisk University
josephshode@gmail.com

Gavin Li
Stanford University
hawklin@cs.stanford.edu

Carrie Gates
FS-ISAC
cgates@fsisac.com

Grant Ho
University of Chicago
grantho@uchicago.edu

Zakir Durumeric
Stanford University
zakir@cs.stanford.edu

Abstract

To combat the deluge of enterprise breaches, government agencies have developed and published a wealth of cybersecurity guidance for organizations. However, little research has studied this advice. In this paper, we conduct the first systematic analysis of government guidance for enterprise security. We curate a corpus of prominent guidance documents from 41 countries and analyze the availability of advice, the coverage provided by the advice, and the consistency of advice across countries. To facilitate detailed analysis and comparisons, we develop a tree-based taxonomy and quantitative comparison metric, and then apply these tools to analyze “essential” enterprise best practice documents from ten countries. Our results highlight a lack of consensus among the governments’ frameworks we analyzed—even among close allies—about what security measures to recommend and how to present guidance.

1 Introduction

Governments worldwide are working to educate and support constituent organizations in improving their cybersecurity. Many countries even have dedicated agencies, such as the U.S.’s Cybersecurity and Infrastructure Security Agency (CISA) and Singapore’s Cyber Security Agency (CSA), whose missions include helping companies to defend against and recover from attacks. Operating in a complex ecosystem composed of both continually evolving adversary behavior and embellished industry marketing claims, CISOs and other senior security leaders carefully track recommendations from these government agencies, which they view as an important and impartial source of cybersecurity guidance [60].

Yet, despite the critical role that government guidance plays in shaping security programs, little prior analysis has been done of governments’ cybersecurity guidance. Understanding the guidance companies receive is critical; without it, we lack

the foundation to determine if these government efforts are effective. Our work aims to bridge this gap by answering the following research questions:

- RQ1:** Availability: Which countries publish recommendations to enterprises? Is this guidance generic, or is specific guidance available (e.g., targeting specific industry sectors, business sizes, or technologies)?
- RQ2:** Coverage: Do these recommendations span all the security themes observed in existing security frameworks? Is there comprehensive coverage of specific security controls? What level of detail is provided?
- RQ3:** Consistency: Are security themes consistently covered? Are recommended controls consistent? Do countries provide diverging or contradicting advice?

To answer these questions, we introduce a taxonomy for building a structured representation of enterprise security guidance that allows us to compare guidance across countries. We survey prominent security guidance from 41 countries, and systematically analyze a sample of 10 documents from different countries using this taxonomy. Additionally, we develop and use a metric that allows us to compare the advice between countries’ frameworks to determine the consistency of recommended controls and their implementation guidance.

Rather than observing consensus, we find that the government documents we analyzed vary substantially in what they recommend and how they construct guidance. Most of the 41 countries we examined publish voluminous guidance material via a bevy of differently scoped documents. For their general-purpose documents intended for any corporate entity, these governments differ in whether the intent of the document is completeness or prioritization, how compact or extensive it is, whether it is further decomposed into maturity levels, and what kind of incentive is provided (if any) for following it.

The essential controls that our sampled government sources identify have little overlap. For a set of guidance documents from 10 countries, only 2 of 166 observed controls were universally agreed upon. Even the most widely promoted controls are explained in varying degrees of detail, ranging from

*This version corrects minor errors in Table 1 (§3.2). See Appendix G.

[†]Work done while visiting Stanford University.

a single generic sentence to multiple pages of implementation instructions. Across the board, the countries we analyzed differ from each other pairwise by about 53%—and even close allies (US, UK, and Australia) differ nearly as substantially. Moreover, we also find direct contradictions between the advice from different countries about the same security controls.

In summary, we contribute:

- The first systematic review of government advice for enterprise security, spanning 41 countries.
- A taxonomy of enterprise security advice, including a systematization of 10 government guidance documents.
- An analysis of prominent enterprise guidance provided by government agencies in ten countries showing significant lack of consensus on advice and its presentation.
- Public release of our data and code to allow future studies to extend our work.¹

Together, our results call into question the expert consensus backing enterprise cybersecurity guidance. We argue that the differences that we observe are largely a byproduct of subjective opinions based on limited evidence. Our work serves as a call to the community to build stronger empirical evidence on what security controls are most effective.

2 Related Work

We draw inspiration from the body of work that studies security advice for individuals. Ion et al. [18] and Busse et al. [4] compared what security experts vs. non-experts view as the most valuable security practices, and Redmiles et al. [38] studied at scale how users perceive the quality of security and privacy advice. These studies have concluded that advice for individuals is too voluminous to be usable [38, 39] and proposed interventions for making advice more actionable [42]. Other work has looked at individuals’ advice seeking behaviors [30, 36, 37] and adherence to advice [9, 65]. Finally, Acar et al. [1] studied security advice for software developers.

The literature on enterprise security guidance, however, is sparse. Wolf et al. [60] conducted interviews with security managers at small businesses and found that organizations perceive government security advice as more trustworthy, albeit less usable, than industry sources. However, their work does not analyze advice itself. Other work has looked at security compliance requirements and identified gaps and inaccuracies in three sector-specific compliance standards [50] as well as what motivates organizations to implement security measures beyond such requirements [51]. There is also criticism of how enterprise security primitives are commonly framed, coupled with counterproposals such as formal analysis [34] and explicit trust relationships [47].

Additionally, a growing body of work has studied how individual players operate in the enterprise security landscape, including CISOs [29, 43, 60], cyberinsurance providers [6,

31], security operations centers (SOCs) [21, 23], and threat hunters [3, 27]. Other work has documented enterprises’ security governance processes [45, 46, 59], described and critiqued the degree of governments’ security oversight [10, 56], and developed political and sociological theory of governments’ cybersecurity actions [7, 48]. However, these prior works do not systematically analyze enterprise security guidance itself.

3 Published Guidance Documents

In this section, we describe how we identified and selected guidance documents, the availability of guidance, and how guidance is presented.

3.1 Identifying Guidance

Below we describe how we built our dataset of the most prominent government security guidance documents.

3.1.1 Selecting Prominent Countries

We start our search for guidance documents by identifying countries whose governments are prominent in the security landscape. We combine six rankings of countries with complementary objectives and methodologies: four cybersecurity, one technological development, and one population:

1. **ITU Cybersecurity Commitment.** The International Telecommunication Union (ITU) ranks countries based on their commitment to cybersecurity, as demonstrated through legal, technical, organizational, educational, and collaborative efforts [53].
2. **National Cyber Security Index.** The National Cyber Security Index (NCSI) ranks countries’ security maturity based on policy, diplomacy, education, research, and incident response capacity indicators [8].
3. **Cybersecurity Exposure.** The Global Cybersecurity Exposure Index uses data from Microsoft to rank countries based on targeted attacks (e.g., ransomware) [33].
4. **Cyber Power Index.** The National Cyber Power Index, from the Harvard Kennedy School, ranks countries on “cyber power” using capability and intent indicators such as international engagement, education/awareness campaigns, offensive actions, and security technology industry development [57]. We used the Defense ranking.
5. **Global Finance Tech Advancement.** Global Finance magazine ranks countries by technological advancement, based on Internet and cell network penetration, digital competitiveness, and R&D spend [15].
6. **Population Rank.** Finally, we considered the largest countries by population [16].

Taking the union of the top 10 countries from each list yielded 41 total countries across 6 continents (Appendix A, Table 3).

¹<https://doi.org/10.5281/zenodo.15612457>

3.1.2 Finding Guidance Documents

For each of the 41 countries, we conducted an Internet search for government-published enterprise security guidance. We conducted searches manually, because prior work found that over 90% of security advice for individuals was missed by automated methods [38]. During our search, we likewise observed a long tail of website idiosyncrasies that were best navigable through human judgment (e.g., long tail of descriptive terms and open-ended or generic-sounding index pages). Five researchers participated in the search, conferring with each other as needed during the process, each researching a subset of the 41 countries.

First, we collected a list of potentially relevant government agencies. We used two sources: (1) a list of government agencies from the United Nations (UN) Cyber Policy Portal [12], which releases cybersecurity strategy, policy, or legal documents;² and (2) independent Google searches for government agencies with public-facing pages about cybersecurity ("`<country> cybersecurity <ministry/agency/department>`"). Second, we systematically investigated each agency's website to find guidance. This systematic search consisted of manually iterating through menus, sitemaps, and cross-references to find resources. We did not search for state/province resources.

For agencies that do not focus on cybersecurity (e.g., a defense agency) but had a site search function, we searched within the site for "cybersecurity" or a local-language translation (as found through Google translate). If we found no guidance on an agency's site (e.g., a national police force that informed the public about their cybercrime-fighting divisions, but did not publish security advice), we dropped the agency from consideration. We used Google Translate to navigate and understand non-English web pages and PDFs. For each country, we also identified 1–3 primary documents: those most prominently featured on agency websites and/or linked to most frequently from other pages on the websites.

For comprehensiveness, to supplement this agency-based search, we performed direct Google searches of the form "`<country> government cybersecurity recommendations`", iterating through substitute terms such as "business security" and "guidelines." We pursued search results beyond the first page until we reached saturation on relevant results. This direct search turned up no materially new guidance resources.³

3.1.3 Limitations

While we attempted to find guidance from each country through multiple avenues, both traversing agency websites identified by the UN and using public search engines, we cannot guarantee that we were able to find all documents. The

²Except for Taiwan, which was not included in the UN database, but which has unique cybersecurity interests.

³Direct searching found two outdated versions of materials already cataloged: one Luxembourg resource that had been relocated to a different domain, and one English translation of a superseded Estonian standard.

manual nature of searching by multiple researchers may have missed documents and the dynamic nature of search engines might lead to other results at different points in time. We conducted our searches from within the U.S. using Google in English, which could affect the documents we found compared to local search engines (e.g., Yandex). Further, our Google translations for search terms may lack needed nuance.

It is impossible to prove that we found all documents, but the fact that our validation searches for guidance on Google did not surface additional documents provides some reassurance that we found the most important documents publicly promoted by each country. We were able to find documents for every country except Russia, which has repeatedly indicated its intent to block access to government websites from Western search engines [24]. It may be that we are simply unable to access Russian guidance without bypassing country security mechanisms. In other cases, we may miss niche documents and underestimate the guidance provided.

3.2 Availability and Types of Guidance

With the exception of Russia, we found published cybersecurity guidance from all of the countries we investigated (Table 1). Oftentimes, we found dozens of advice resources. For most countries, these resources emanated from 2–5 agencies or ministries, not all of which had a primary security mission. Except for Indonesia, Nigeria, and Mexico, nearly all countries (37/41) have at least one general-purpose guidance document available (i.e., addressed to all organizations about broadly securing their infrastructure). Most countries had multiple such documents, and about a third had five or more documents with this same general-purpose scope. The U.S. far outflanks other countries in sheer volume of guidance, with 18 general-purpose resources.

Beyond general-purpose guidance, the vast majority of countries also publish a suite of guidance documents that are more *targeted* in scope. We identified six classes of targeted guidance: documents intended for a particular industry *sector* or for companies of a specified *size*, documents about securing a particular *technology*, documents addressing a specified *threat model*, documents focused on a particular *mitigation*, and documents contextualized in a moment in *time*. Below, we outline key features of these classes of guidance.

3.2.1 Audience-Focused: Sector, Size, Technology

These classes of guidance documents provide advice on topics specific to a particular audience or setting. The vast majority of countries (39/41) published at least some guidance about securing particular types of technologies, but specific coverage varied substantially. The top two technologies were cloud systems (27 countries) and network appliances (23), followed by cryptographic systems, web, and mobile systems (each 20). A long tail of technologies also surfaced in our search, includ-

	Country	Agencies	General Purpose	Timely	Size-specific	Critical Infra	Sector Specific						Tech Specific								Threat Model Specific						Mitigation Specific																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																				
							Any	Govt./Defense	Telecom/IAP	Finance	Transportation	Health	Education	Any	Cloud	Networks	Web	Mobile	Crypt.	IoT	OT	AI	Any	(D)DoS	Ransomware	Phishing/spam	Malware	Remote/BYOD	Data breach/loss	Any	IR	IAM/Auth	Risk Mgmt.	Governance	Logs/alerts	Network defenses	For individuals	Pro. development																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																									
	Australia	5	6	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●

Table 1: **Guidance Availability.** Nearly every country releases general-purpose guidance as well as more targeted documents.

ing robotics (Bangladesh), printers/copiers (U.S.), dashcams (Malaysia), and hearing aids (Norway).

Beyond specific technologies, most countries also issued guidance targeted to particular industry sectors (36/41), but again exhibited high variance in coverage. The most common sectors were government and defense (25 countries) and finance (13). Other common sectors correspond to industries often considered as critical infrastructure, such as health, transportation (esp. rail and ships), education, water, and energy; 29/41 countries published guidance aimed at critical infras-

tructure broadly. However, we also observed a long tail of sectors, including political campaigns (France), lawyers (UK), e-sports (Luxembourg), and computer-controlled chicken coops (Israel, Appendix Figure 7). The long-tail of sector-specific guidance also came in part from countries with no guidance for the six most common sectors (e.g., Lithuania and Czechia released guidance for hotels and satellites, respectively, which most countries omit).

Guidance tailored to companies of a particular size was more uniform in coverage. About half of countries published

such guidance, and when present, these resources almost exclusively focused on small to medium sized businesses (SMBs). Only four countries had resources addressed to larger organizations: Canada, U.K., Australia, and the U.S. While not our focus, we also observed many countries that published guidance addressed to individual employees or citizens (37 countries) and/or training program materials for security professionals (19 countries), indicative of a rich ecosystem of government efforts to help different audiences.

3.2.2 Threat-Focused: Threat, Mitigation, Time

These classes of guidance center on the adversarial attack/defense landscape in which companies operate. Most countries published some form of guidance addressing a specific threat model of concern; the most common issues were (D)DoS (24 countries), ransomware (22), and phishing/spam (19). Unlike audience-focused guidance, threat model specific guidance coverage was more bimodal: countries either had targeted guidance for most of the top threat models in our dataset, or had guidance for nearly none at all. However, we still observed a long tail of coverage including deepfakes (US), drone attacks (Spain), scattered electromagnetic radiation (Finland), and malicious e-books (Lithuania).

In accordance with a shifting threat landscape, nearly every country also released some form of time-sensitive advice, typically (but not exclusively) from a national CERT. Postings covered topics such as defending against new ransomware strains, patching recent CVEs, and navigating contemporary events like COVID-19 or national elections.

Finally, some guidance focused on specific methods for addressing threats (e.g., MFA or network segmentation). Grouping mitigations using the top-level themes from our analysis framework (Section 4.1.1), we identify the top themes of targeted guidance as incident response (24 countries), IAM and authentication (24), and risk management (16).

3.2.3 Who Publishes Guidance?

Governments' guidance authorship is frequently decentralized. Countries had a median of three agencies publishing enterprise security guidance, but with high variance stemming largely from sector-specific guidance released by agencies whose primary mission centers on that sector. Even for security-focused agencies, nested sub-agencies sometimes have separate websites from their parent agency and host guidance that might not be co-authored by their parent agency (e.g., for India, CSK $\subset$ CERT-IN $\subset$ MeitY, and NCIIPC $\subset$ NTRO). The U.S. overwhelmingly had the largest number of agencies issuing advice (36 agencies), followed by India (10) and New Zealand (9), and subsequently Ukraine and Austria (each with 8 agencies). Even looking only at general-purpose guidance, the U.S.'s 18 documents originate from four separate agencies (CISA, NSA, FBI, and NIST). Given potential

limitations of our search, these results may still underestimate guidance, particularly in non-English speaking countries.

3.3 Presentation

Previously, we identified general-purpose "primary" documents that serve as governments' main modes of addressing organizations broadly. In this section, we examine how these documents are presented. We find that the format, intent, size, and complexity of primary guidance documents varies widely from a simple page on an agency website to a PDF dozens of pages long.

3.3.1 Document Intent

Many documents contained some form of preface text that explicitly indicates the document's purpose, audience, and/or approach. Using this preface text, we categorized documents into four groups based on their *intent*:

- **Essential Control Lists:** Guidance that only present a pared-down list of controls deemed most worthwhile. Distinguished by words like "essentials," "top N," "basics," or, by inference, short length. Example: Belgium's CyberFundamentals Framework contains a "*Basic*" level document that describes security measures that "*provide an effective security value*" for all enterprises.
- **Control requirements:** Security criteria with mandatory or auditable compliance for designated organizations (e.g., defense contractors). Example: Pakistan's primary guidance describes itself as required for government and critical sector suppliers but also "*encouraged for private and commercial sector organizations*".
- **Catalog:** Non-opinionated lists of controls which rely on the readers' judgment to select and prioritize. Distinguished by descriptive text to this effect (e.g., that choice of controls may depend on the organization and its risk tolerance). By inference, extensively long documents that do not frame themselves as either basics or requirements are also categorized as such; we apply this inference using consensus of three researchers. Example: The US's NIST Cybersecurity Framework states that "*the way organizations implement the CSF will vary*" and that it "*assists its users in learning about and selecting*" specific controls/mitigations.
- **Not Controls:** Guidance that lacks an itemized list of security measures to implement, and instead presents an abstract framework or approach. Example: New Zealand's Cyber Security Framework "*sets out how the NCSC thinks, talks about, and organises cyber security efforts*", without listing any associated controls.

We catalog the full list of primary documents and their intent in Appendix E, Table 5. For our investigation, we were primarily interested in the guidance documents that captured

what governments consider essential security advice for enterprises. For example, our analysis focuses on the U.S.’s essentials-focused Cybersecurity Performance Goals rather than the catalog-type NIST Cybersecurity Framework, since the former encodes its authors’ value judgments on worthwhile controls in a way that the latter does not. Thus, we focused on only the first of these four categories (33 / 57 primary documents).

3.3.2 Length, Complexity, and Tone

Our set of “essential” guidance documents varies in length, complexity, and tone (Appendix F, Table 6). In terms of length, documents varied from India’s 217-word document that spanned less than one page, to Japan’s two-part guideline of a 15K+ word “Cybersecurity Management Guidelines” and accompanying implementation guide of over 35K words. Although some variation in word-length comes from more controls (as we address in Section 4), some documents simply contained far more words per control than others, ranging from more implementation details (Section 4.5), to providing explanations that justified the purpose of controls, to simply having more verbose language. For example, New Zealand spends 3,391 words on patching alone, while Bangladesh spends 3 words: “*Implement patch management.*”

Beyond length, documents varied in the complexity of language used and tone. We use the Flesch-Kincaid Grade Level [22], based on Flesch reading ease [11], to estimate the reading grade level of English guidance documents. This method is the most widely used reading-ease metric and has been shown to be a meaningful, albeit imperfect, metric for security domain-specific text [35]. As it is only designed for English, we do not compute it for guidance documents with no official English version available. Using this metric, we see a wide range of reading levels: New Zealand’s documents can be read with a high school education, while the US, France, and Japan are written at a graduate degree level.

Some countries, such as Luxembourg, used colloquial language that appeared to mimic casual banter (e.g., “*Passwords must be looong. To make it even more secure, use a combination of numbers, capital and lower-case letters, as well as punctuation marks.*”). Others, such as advice from Singapore, presented using more formal and structured imperatives: “*The organisation shall change all default passwords and replace them with a strong passphrase, e.g., it should be at least twelve (12) characters long and include upper case, lower case, and/or special characters.*”

3.3.3 Limitations

Of the 33 essential guidance documents we analyzed, 22 are in English or had an official English version (Appendix E, Table 5); we used English versions when present. For the other 11 documents, we relied on Google Translate, which could

miss nuances in tone. We also acknowledge that our tone interpretation may be biased or miss cultural norms, and our reading complexity metric does not measure global readers’ perceived complexity.

4 Guidance Analysis

In this section, we analyze the guidance provided in ten “essential” guidance documents from ten countries spanning five continents: Australia, Egypt, India, Israel, New Zealand, Norway, Singapore, Ukraine, U.K., and U.S. (Appendix F, Table 6). This includes the U.S.’s Cybersecurity Performance Goals (CPGs), Australia’s Essential Eight, Singapore’s Cyber Essentials, and Norway’s Basic Principles for ICT Security. We chose these ten countries as a subset of sources representing: (1) a diversity of global regions, (2) examples of allied countries that coordinate on cybersecurity (U.S., U.K., and Australia), and (3) a balance of document lengths. While we chose documents to compare across these dimensions, our results may not generalize beyond these ten countries.

4.1 Enterprise Guidance Taxonomy

Guidance in the documents we analyzed varies widely in terms of suggested controls, implementation details, and presentation format. To compare these documents, we introduce a hierarchical taxonomy for systematically mapping and comparing guidance. Our taxonomy has two high-level goals:

1. *Comprehensive*: Our taxonomy must provide broad coverage to compare and recommendations across countries.
2. *Hierarchical*: Our taxonomy must facilitate fine-grained comparisons that identify conceptual and implementation level similarities and differences.

Since no existing taxonomy is sufficiently hierarchical, we developed our own by combining two existing control frameworks to meet Goal 1 (§4.1.1), and then using a structured set of attributes to add further hierarchical depth to meet Goal 2 (§4.1.2). We specifically map the recommended controls in guidance documents into a five-level semantic tree:

Level 1: Themes: Broad security topics that group together sets of related controls, such as *Asset Management* or *Network Security*.

Level 2: Subthemes: Focused groups of controls united around a common goal within a given theme. For instance, subthemes within *Network Security* include *Architecture and Segmentation*, *DNS Security*, and *Traffic Inspection*.

Level 3: Controls: Distinct security components, countermeasures, or practices to advance security within a given subtheme. For instance, controls for *Traffic Inspection* include *Visibility of Encrypted Communications* and *Content Disarm and Reconstruction*.

Level 4: Attribute Identifiers: Boolean indicators for whether a guideline contains specific types of information about a control’s implementation. For instance, does the *Patching and Remediation* control specify *what* assets scope to patch, the target timeline for *when* to patch, and/or *how* patching should be conducted?

Level 5: Attribute Values: The concrete attribute values for a control. For example, for *when* assets should be patched, the attribute might be *2 days*, while another might specify *2 weeks*.

This hierarchical structure allows us to model guidance documents as trees for subsequent analysis (Section 4.2).

4.1.1 Defining Themes and Controls

The top three levels of our analysis framework—themes, subthemes, and controls—dictate how we can distinguish and sort advice content from different countries. To ensure broad coverage when creating these top levels, and to avoid overrelying on a single industry source, we merged two popular *common controls frameworks* designed to help compliance teams in enterprises harmonize their many security compliance requirements: first, Secure Controls Framework [5] (SCF), developed by industry expert volunteers; and second, Adobe Common Controls Framework [2] (CCF). Each contains several hundred controls loosely binned into top-level themes such as *Governance* and *Vulnerability Management*.

We used a qualitative methods approach to build themes, subthemes, and controls based on these two frameworks. We treat themes, subthemes, and controls as a set of semantically-grouped deductive codes, forming our codebook’s top-level themes by merging the top-level categories between the two frameworks, which had high consistency. Within each theme, we merged and grouped similar controls (which may differ in wording, depth, or coverage) to form our own list of subthemes and controls. We then iteratively refined the codebook based on group discussions while applying the codebook to government guidance data. The finalized codebook—containing 29 themes, 93 subthemes, and 228 controls (Appendix D)—became the first three levels of our taxonomy.

4.1.2 Normalizing Control Implementation Details

As described, our merged control framework enables us to compare *which* controls and themes are included, but it lacks the granularity to compare the implementation guidelines for each control. To support more granular systematic analysis, we extend our taxonomy by adding a third and fourth layer that describe the *attributes* of each control using a set of structured questions drawn from classic journalism practice [26]: Who, What, When, Where, and How. This framing provides us with a set of deductive codes that are simple to apply to each

control. For instance, attributes for a control about patching can include attributes around *what* issues to patch, *where* (what systems) to patch, and *when* (how fast) to patch.

Specifically, the third layer nodes indicate whether or not the guidance contains any Who, What, When, Where, or How details about a control. The final layer contains the control’s specific details for each attribute (*attribute values*). For instance, the U.S.’s CPGs specify to patch “within a risk-informed span of time,” while Australia’s Essential Eight say to patch critical vulnerabilities “within 48 hours of release”. Both address the same attribute of patching timeline (“When?”), but differ in their specifics. We developed this layer inductively from the data during analysis.

4.2 Analyzing Guidance Documents

Our taxonomy allows us to transform unstructured guidance documents into a normalized and structured form for comparison (e.g., Appendix Figure 9).

4.2.1 Systematizing Guidance Documents

We used our taxonomy to construct semantic trees for each of our ten guidance documents as described below.

Mapping Themes and Controls. First, we mapped each document’s controls into the top three layers of our taxonomy (themes and controls). Five researchers conducted this mapping using descriptive first-cycle coding [44]. For each country, two researchers independently coded controls, then met to resolve disagreements. Although we did not aim to (nor did we) hit saturation with guidance document content, we saw reasonably little long-tail guidance content that defied categorization, and no clear patterns therein (Section 4.4.2), which lends confidence that our taxonomy captures key concepts with reasonable comprehensiveness.

Mapping Attributes and Specifics. Next, we built out attributes for the controls in the three most frequently occurring subthemes across the ten guidance documents: Backups and Redundancy, Patching and Remediation, and Authentication. Focusing on the subthemes with (near-)unanimous support from guidance documents provided us with the clearest comparison of how these documents discuss control implementation. We assigned each phrase of each control to one of our taxonomy’s attributes; Table 2 shows an example. One researcher constructed the attribute values layer inductively using straightforward in-vivo open coding [44] of a control’s phrases which had been labeled with the relevant attribute.

Limitations. Mapping guidance content into our taxonomy hinges on a clear understanding of what the guidance communicated. For the three non-English documents among our ten selected (Israel, Ukraine, Egypt), there is a risk that meaning was lost in translation. Our focus was on concrete imperatives

Attribute	Description	Control Phrase
What?	Patch what issues?	“All known exploited vulnerabilities (listed in CISA’s Known Exploited Vulnerabilities Catalog)”
Where?	Patch which systems?	“internet-facing systems”
When?	Patch when?	“within a risk-informed span of time”
How?	Patch what first?	“prioritizing more critical assets first”

Table 2: **Example Control Decomposition.** Text from original CISA CPG Control: “All known exploited vulnerabilities (listed in CISA’s Known Exploited Vulnerabilities Catalog) in internet-facing systems are patched or otherwise mitigated within a risk-informed span of time, prioritizing more critical assets first.”

(e.g., enforcing password rotation every 90 days) rather than nuances or tone. A spot-check by a native Hebrew speaker between the Israeli guidance document and our translation, as well our own spot-check between the English version of Norway’s document and our Google translation of the Norwegian version, did not surface substantive differences. Additionally, we checked for off-topic or hard-to-parse text that might indicate mistranslation. We identified one such case: Egypt’s “Activate ‘Register’ in order to better investigate any security issues” would have been better translated as “Activate logging.” Our three translated guidance documents could possibly contain other errors that went unnoticed by our team. However, the vast majority of translated text made understandable use of domain-specific language, with no obvious oddities.

4.2.2 Comparing Guidance Documents

To concretely capture the overall differences between advice represented in our framework’s tree structure, we developed a quantitative tree-comparison metric that represents the (dis)similarity between two guidance documents. For our particular taxonomy, this metric has three goals:

1. The comparison metric is designed for rooted trees where the order of child nodes under a parent does not matter but their labels do.
2. The metric allows comparing documents that contain partially non-overlapping controls or mismatched depths.
3. The metric should weight differences near the root more heavily than near the leaves (i.e., two documents have a higher dissimilarity if they cover different themes than documents with identical themes but different controls).

Unfortunately, we find that existing tree comparison metrics do not serve our purposes. For example, existing metrics use domain-specific tree properties that do not match our use case, such as for physical-world simulations [49] or user purchase histories [63]; compare taxonomic rearrangements of a fixed set of leaf nodes (e.g., candidate evolutionary trees), rather than different node sets [28, 62, 64]; struggle with scalability [62]; weight changes near the root the same as changes near the leaves [20, 25]; or operate on non-rooted trees [52].

Thus, we introduce a new tree edit-distance metric for our analysis. Our comparison metric considers only leaf-node insertion and deletion operations. The weight of inserting or




Figure 1: **Visualization of Tree Distance Metric.** Transforming the orange tree (left) into the blue tree (right), both mapped into the same taxonomy (gray), involves removing the nodes labeled ‘-’ and adding the nodes labeled ‘+’, each weighted by its distance from the root. The total weight of the trees’ union is $2 + \frac{3}{2} + \frac{5}{4} + \frac{2}{8} + \frac{3}{16} = 5.1875$. The total distance (dissimilarity) between trees is thus $(1 + \frac{2}{2} + \frac{4}{4} + \frac{1}{8} + \frac{2}{16}) / 5.1875 = 0.63$.

deleting a node depends on its distance from a dummy root node at the top of a document’s tree (the parent node to all Level 1 “Theme” nodes), and drops off exponentially with distance. For example, inserting a top-level “theme” node of Vulnerability Management has a weight of 1; a second-level “subtheme” node of Vulnerability Identification, $1/2$; a third-level “control” of Vulnerability Scanning, $1/4$; a fourth-level “attribute identifier” about what systems to scan, $1/8$; and a fifth-level “attribute value” about web servers, $1/16$.

To compute our metric for two documents’ trees, we identify the minimal set of nodes that must be inserted or deleted to transform one tree into the other, sum these nodes’ weights, and normalize by the total weight of the union of nodes in the two trees.⁴ Thus, our metric provides a score that quantifies the (dis)similarity between two documents’ trees: lower scores mean higher similarity (i.e., a score of zero means identical advice) and higher scores mean greater differences. Figure 1 shows a simple example of computing this metric. This metric satisfies all of our goals: it operates on rooted unordered labeled trees, handles trees that contain some disjoint or mutually exclusive nodes, and employs an intuitive weighting scheme that equally weights nodes at the same conceptual level in our framework.

⁴For those familiar with the Jaccard index, this distance can alternatively be conceptualized as a weighted Jaccard dissimilarity.

4.3 Content Volume and Depth

We first investigate the differences between the structure of countries' primary cybersecurity guidance documents. Across the board, we find that guidance differs substantially, in terms of both the number of controls highlighted and level of detail. This variance suggests disagreement both in what terms of what governments believe is essential for organizations' protection as well as the level of detail that organizations need to correctly implement those controls.

To quantitatively compare the content in each guidance document, we consider a "universal tree" reference point, which consists of the union of all document trees with our original analysis framework based on common controls. We then use our tree distance metric to compute dissimilarity between each country's guidance tree from our universal tree (quantifying *lack* of coverage), then take one minus distance to get coverage. New Zealand provides the most total content, covering 71% of the universal tree, followed by Norway (66%) and the US (58%). At the other extreme, India covers only 21% of the tree, with Ukraine and Egypt following at 23%.

Higher content volume could result from greater depth (e.g., more details about controls) or greater breadth (e.g., more recommended controls). To better understand differences, we visualize each country's guidance document as a tree. (Note that we only see depth in a few subthemes because we only build out attributes for the controls in the most popular subthemes: see Section 4.2.1.) As shown in Figure 2, control breadth and the density of attribute values correlate, but there exist exceptions. For example, New Zealand is the most specific in its recommendations and is second only to Norway in breadth. In contrast, Singapore's content volume comes more from its breadth, while Australia's comes more from its depth of detail about each control. Similarly, India and Ukraine have comparable content volume (21% vs. 23% of the overall taxonomy), but India's tree shows greater breadth while Ukraine has more depth.

4.4 Content Agreement: Coverage

We next examine the consensus (and lack thereof) over recommended essential security controls.

4.4.1 Full and Mixed Consensus

Themes. As can be seen in Figure 4, there is mixed agreement about what security themes and controls are essential. Indeed, only three themes are covered by all ten countries' documents: business continuity and disaster recovery, identity and access management, and vulnerability and patch management. Beyond these three, we see a significant drop in agreement. Of the 27 themes that appear in at least one country's guidance, only 8 themes are covered by more than 75% of guidance documents. Over half of the content—15 of 29 themes (52%)—sits in the middle, neither consistently included nor consis-

tently excluded. This includes security training (5 countries), web security (4), and threat intelligence (3), and subthemes such as asset inventory (6 countries), spam/phishing email protection (5), and third-party risk management (4).

Controls. The lack of strong consensus between countries becomes even starker when looking at subthemes and individual controls: 75 subthemes appear in at least one guidance document, but only 9 / 75 (12%) are recommended by over 75% of countries. For controls, this drops to 5 / 166 (3%) recommended as essential by over 75% of countries. Most controls appear in only one or a few countries' guidance documents, forming a long tail of recommendations (Figure 5). Over a third of included controls remain in the middle, neither consistently included nor consistently excluded: 59 / 167 (35%). For instance, while all 10 countries' documents discussed patching vulnerabilities, far fewer provide recommendations on how to scan for (3 countries) or prioritize vulnerabilities (6 countries), and/or how organizations should manage their vulnerability program (2 countries). Strikingly, only 2 out of the 166 controls unanimously appear in all countries' guidance: backups and patching.

4.4.2 Long Tail of Coverage

A long tail of advice is recommended as essential by only a small minority of countries. Four themes (15% of the 27 covered themes) appear in less than 25% of countries' guidance: maintenance (2 countries), compliance and auditing (1 country), enabling customers' security (1 country), and human resources security (i.e. considering employees as potential insider risk, 1 country). Among subthemes and controls, even more sparsity exists: 33 / 75 (44%) subthemes and 102 / 166 (61%) controls appear in less than 25% of the guidance documents.

Part of this long tail results from countries with significantly longer essential-guidance documents. For example, Norway has the largest coverage and presents 9 / 16 subthemes that appear in only one country's document. However, this long tail of advice also resulted from unique or unusually specific security controls that defied typical categorization. Examples include Australia's advice to annually validate "*Microsoft Office's list of trusted publishers*", Norway's advice to "*build back better*" when restoring services after an incident, and Israel's advice to avoid unnecessarily disclosing "*details that are not essential to the functioning of the organization's system*". Perhaps the most unexpected long-tail advice came from the U.S.: "*Organizations sponsor at least one "pizza party" or equivalent social gathering per year that is focused on strengthening working relationships between IT and OT security personnel and is not a working event (such as providing meals during an incident response).*" This long tail of rarely proffered advice could either result from other countries explicitly deciding the content is less important, or because




Figure 2: Tree Shapes. Overall shape of tree coverage shown for selected 10 documents. Colored subtrees show the popular controls for which we built out control attributes (blue = backups, green = authentication, orange = patching); we explore these subtrees further in Section 4.5. New Zealand specifies attributes with the most detail, while Norway has the broadest overall coverage of controls. Despite some differences in these overall shapes, control breadth and detail coverage largely correlate.




Figure 3: Tree Dissimilarity. Distances between countries' primary guidance documents (country codes defined in Appendix A). Most documents differ from each other by over half their combined coverage; India's is the most different.

countries simply overlooked and omitted it. In either case, the rarity of these controls implicitly indicates that most countries do not consider them a top priority.

4.4.3 Uniformly Omitted Recommendations

The ten documents we analyze represent a set of "essential" security advice (i.e., curated and branded as a set of core, important security measures by government agencies). Thus, we also examined what themes and controls countries frequently *omitted* from their essential recommendations (relative to the comprehensive industry security controls frameworks that we




Figure 4: Top-Level Theme By Country. Countries lack consensus on whether to include over half of all themes.

used to generate our taxonomy).

All ten documents omit controls around two top-level themes from the common controls frameworks: AI and security operations. At the subtheme level, 18 / 93 (19%) subthemes in the taxonomy were never used. All countries omitted several HR mitigations for insider threats such as separation of duties and vetting new hires. Also notably absent was (D)DoS protection, addressed from neither a capacity planning nor network configuration perspective. Remaining omitted controls range from data minimization, to endpoint defense testing, to system/product documentation.

While we may reasonably expect omitted topics to be in scope for consideration due to their presence in common con-




Figure 5: **Distribution of Subtheme and Control Coverage.** Very few subthemes or controls are common across countries, and most form a long and heavy tail of guidance content.

controls frameworks, there may be multiple reasons that some of these themes, subthemes, or controls do not appear at all. We observe that a topic’s lack of coverage in an “essentials” document does not preclude it from appearing in the ten countries’ supplemental, targeted guidance documents. Although countries may treat some of these topics separately because they are new and rapidly evolving (e.g., AI, which six countries have standalone documents for), others are well-established: three countries issue guidance specifically about security operations, six countries have DDoS specific documents, and two countries for insider threats. Excluding such content from essentials-focused guidance may indicate that countries explicitly view this content as less essential than other heavily discussed security topics, or it may have simply not come to mind when assembling essential controls.

4.4.4 Quantifying Country Similarity

To unify our per-layer observations and quantify the overlap between countries’ recommendations, we use our tree similarity metric (§4.2.2) to compute pairwise comparisons between documents (Figure 3). We find that while countries do have overlap, there are large differences in the content and structure of their advice: the average pairwise dissimilarity is 53%.

New Zealand and Norway have the highest similarity: both countries’ documents have blanket coverage of most themes and controls we observed across our data (indeed, their intersection alone covers 22/27 themes and 42/75 subthemes). Similar to both of these countries, the U.S. and Singapore also aim for fairly broad coverage of themes and controls (overall coverage 58% and 46%, respectively, just behind Norway’s 66% and New Zealand’s 71%). By contrast, India’s guidance has the highest differences with other countries, with a dissimilarity score ranging from 46% to 73%. When compared to other guidance of similar size (Ukraine, Israel, Egypt), India’s recommendations have large differences as a result of uncommon choices for themes and controls (e.g., omitting network security but including web security).

Surprisingly, even countries with strong geopolitical alignment show substantial difference. For example, the “Five




Figure 6: **Ally Agreement.** Overlap between controls from three high-profile security allies: U.S. CISA, U.K. NCSC, and Australia ASD. The three agencies agree unanimously on only 13% of the union of the security controls they recommend.

Eyes” countries, which include the U.S., U.K., and Australia, are a set of high-profile allies who regularly put out joint security statements [54, 55]. Nonetheless, documents from U.S. CISA, U.K. NCSC, and Australia ASD have comparable dissimilarity with each other as they do with other countries: they differ from each other by an average of 48% of their combined content, compared to the overall average pairwise dissimilarity of 53%. These three countries unanimously agree on only 13% of the total controls they cover as essential, despite regularly publishing targeted security guidance together (Figure 6). Indeed, the U.K. has higher similarity with Israel and Singapore than it does with Australia or the U.S. For instance, the U.K.’s document is the only one of the three to discuss network architecture and segmentation (indeed, Australia’s does not cover networking at all), and only Australia’s covers asset discovery. And while the U.S.’s document is the only one of the three to cover incident response (IR) procedure drills, IDS, or DNS-based email protocols (e.g., SPF, DKIM, DMARC), it uniquely omits malware detection.

4.5 Comparing Consensus Content

Although our results show that most controls are not consistently included in a majority of countries’ essential guidance, three subthemes had broad consensus. Two subthemes unanimously appeared across all ten of the guidance documents we analyzed: “Backups and Redundancy” and “Patching and Remediation.” One additional subtheme, “Authentication,” appeared in all countries’ guidance except India’s. However, when we analyzed the attributes and details provided for the controls in these subthemes, we found high variance in the specific information that countries provided.

Backups and Redundancy. All countries’ documents included a recommendation about creating backups for assets, but they differed both in how much advice they provided and the specific actions for implementing a secure back-up process. In terms of agreement, nearly all documents explicitly advised organizations to store back-ups offline or on segmented parts of the network (only Australia and Ukraine omitted this detail). Most countries (except Israel, India, and Egypt) also provide some advice to create back-ups periodically. However, this timing advice was nearly universal in its vague and ambiguous nature, where most advice simply said to “regularly” make back-ups without a more precise

definition about what constitutes “regular” and/or how often a back-up should be kept. Similarly, although all documents explicitly mention concrete assets to back-up, they differ in what assets they explicitly mention: data (6 countries), business critical/high risk assets (4), configurations or settings (3), software/applications (3), and a long tail of other assets appearing in only 1-2 countries’ guidance, ranging from log data (India and New Zealand) to engineering drawings (U.S.).

Furthermore, we found a long-tail of backup advice that only a few countries’ guidance documents mention. In particular, details about securing backups are sparse. Only three documents provide explicit recommendations about access control over backups; Australia’s is the most detailed, enumerating a long list of access control policies such as “*Backup administrator accounts are prevented from modifying and deleting backups during their retention period*”.

Beyond differences in the amount and types of detail, we observed several cases where advice explicitly contradicted. First, while some countries’ guidance states that organizations can store back-ups in either online cloud storage platforms or via external offline media (the U.K. and Singapore), other countries’ guidance explicitly states that organizations should only store back-ups offline (Israel and Egypt). Second, we observed isolated instances where a country offered differing levels of strictness within the same guidance document. In particular, with respect to testing back-ups, New Zealand’s recommendations specifies in one section that full system restoration tests should be conducted at least “*at least once a year*”, but elsewhere specifies “*every couple of months*”.

Patching and Remediation. Patching also appeared in all of the ten countries’ guidance documents, but similar to back-ups, we observed wide variance in the specific details and implementation advice provided. Seven countries’ guidance presented criteria for what vulnerabilities organizations should patch. However, the most common criteria appeared in at most three countries’ documents, with the other 7–8 countries completely omitting it: organizations should patch vulnerabilities exploited in the wild (2 countries), exploitable vulnerabilities (not necessarily exploited) (2 countries), and/or high severity vulnerabilities, either according to the software’s authors or by CVSS scores (3 countries). In terms of what systems or software to patch, guidance documents also lacked any agreement, with a long tail of 18 items such as servers (4 countries), network devices (3), and cloud applications (2).

As another example, most guidance documents either lacked details or provided differing advice on the timing of patching. Australia’s document provides the most specific guidance on how quickly to patch, with detailed target timelines of 48 hours, two weeks, or one month depending on the vulnerability’s severity and the type of affected system. In contrast, New Zealand’s guidance document omits any specific timeframe recommendations and instead discusses the balance between applying a patch during convenient business hours or during periods of minimal business disruption. Sur-

prisingly, the majority of guidance documents omit any mention of automatic updates, with only Norway, New Zealand, Singapore, and the UK discussing this option at all.

More concerning, we also observed contradictions in patching advice from different countries. We saw three disagreeing stances about what version to use when patching: Australia recommends updating to “*The latest release, or the previous release*”, whereas Norway, New Zealand, and Singapore specify only the latest release. Potentially diverging from both of these two recommendations, India states that organizations patch using the “*latest stable (non-vulnerable) version*”. In terms of testing patches, New Zealand specifies a carve-out for looser testing (such as “*limited testing*” for emergency patches); but Singapore does not, saying to “*carry out compatibility tests on updates for operating system and applications before installing them*”, without any listed exceptions.

Authentication. Similar to the prior two subthemes, we see a lack of consensus in the specific details surrounding authentication. All countries’ essentials-focused documents except Ukraine’s and India’s discuss deploying MFA as part of implementing secure authentication, but only six mention specific forms of MFA with a wide range in their recommendations: four countries mention hardware MFA and five recommendations include software authenticator apps. Among the four countries that discuss SMS-based MFA (US, UK, NZ, and Israel), all recommend avoiding it; but three (excluding Israel) caveat that using it is better than nothing at all. Other authentication attributes with wide variance include what systems/applications should require authentication (ranging from remote access tools, to sensitive/important systems, to cloud services) and what accounts require authentication (five countries discuss specific user classes including privileged administrator accounts, while the rest have no details).

We observe several instances of overt incompatibilities across different countries’ guidance more severe than in other subthemes. In particular, password requirements often had contradictory advice. Ukraine and Singapore encourage organizations to employ password character diversity requirements, with Ukraine describing a strong password as “*contain[ing] letters, numbers and special characters*”; however, the U.K. and U.S. explicitly advise against it, with the U.K. advising to “*Support users to choose unique passwords for their work accounts by [...] not enforcing password complexity requirements*” and instead advocating for a “*three random words*” approach. We also observe direct contradictions in length requirements: Ukraine and the U.K. offer secure options for 8 character passwords, but the U.S. explicitly rejects this length as insecure and recommends 15+ characters, noting: “*Modern attacker tools can crack eight-character passwords quickly*”. Beyond passwords, Norway advocates for using biometrics “*on clients frequently used in public areas*”, and the U.K. describes it as equivalent to a password, but New Zealand questions biometrics’ accuracy and uniqueness and deems it “*not a secure authentication method on its own*”.

5 Discussion and Conclusion

Our investigation into government security guidance uncovers not only differences between the countries we study, but methodological challenges for rigorously studying security guidance and future research directions.

5.1 Methodological Lessons Learned

Below, we describe the methodological lessons that we learned during the course of our study.

Finding Guidance Resources. Systematically finding even the most prominently and publicly promoted documents proved difficult given decentralized sources, potential translation issues and geofencing, and time-intensive manual website traversal. It is also possible that we missed some publications. For instance, appropriate search terms, especially non-English ones, may differ by country. Government guidance is a form of “grey literature,” which is, by definition, published outside traditional channels and difficult to retrieve [40]. While grey literature search in computer science is still nascent [61], there has been decades-long development of search methods and resource repositories in fields such as health [17, 32] and social sciences [41, 58] (e.g., to reduce English-language bias), which may provide inspiration for future work.

We see two avenues for expanding upon our initial foray. First, while our goal was to find prominent guidance rather than all guidance, there is room for future work to more comprehensively capture the global landscape of governments’ advice, understand how governments are publishing guidance, and gather perspectives from government representatives. Second, future work may investigate what guidance is seen by security teams in practice and how this advice is found. Unlike prior work on security advice for individuals, which approximated a user’s advice-finding experience using crowdsourced search terms and subsequent Google search results [38], enterprises receive advice not only through public postings, but also through an opaque network of peer information-sharing, threat intelligence sources, and closed forums [43].

Improved Comparison Frameworks. Mapping sources’ enumerated controls into a common taxonomy was surprisingly nuanced and painstakingly time consuming. Even the two industry frameworks we used, written specifically to unify enterprise security actions, did not always align on control delineation. Controls in government advice were even more disparate: in some cases, two enumerated “controls” in a document differed only by an attribute (implementation detail), while in others, a single “control” consisted of conjoined clauses spanning multiple themes. Advice sources use the term “control” non-uniformly, and qualitative analysis is vital for accurately understanding and comparing sources (rather than taking enumerated control lists at face value).

Beyond differing levels of abstraction, we faced two forms of ambiguity during our analysis. First, to serve our purpose of *reliably mapping similar guidance content to the same taxonomic location*, we needed to build a taxonomy with minimally-overlapping control definitions. Common controls frameworks focus on helping organizations demonstrably achieve *coverage* of multiple standards’ requirements, which differs from minimizing overlap. Combined with aforementioned misalignments, reconciling the two frameworks with each other sometimes triggered lengthy discussions about taxonomic structure and mapping rules. Despite this, in some cases we ultimately decided to assign the same guidance content to multiple taxonomic branches. While our taxonomy served our purposes, we also see room for further refinement.

Second, the guidance itself often contained ambiguous wording. For example, several documents described “hardening” systems without providing a concrete definition of hardening. Rigorously resolving these ambiguities involved robust discussion between qualitative coders, closest-match judgments using surrounding document context and domain knowledge, and aiming for consistency with prior decisions. Continuous refinement of our codebook—such as clarifying that “monitoring” should be sorted into the same sub-theme as alerts—also supported consistency. Fundamentally, leading industry and government guidance documents lack a common lexicon to communicate controls, progressive capability levels and associated milestones, and consistent nuanced instructions for their correct implementation; building out such a lexicon is a foundational area for future community efforts.

5.2 Future Directions in Enterprise Guidance

Among the governments we studied, our results highlight a lack of consensus on how enterprises should protect themselves, despite facing many of the same threats. We discuss two key areas for future research informed by our work.

First, our Section 4 analysis shows substantial disagreement between what security themes and controls the government documents we study emphasize—even among close allies. Out of 228 distinct security controls, only two unanimously appear in all 10 countries’ guidance, only 3% appear in over 75% of countries’ recommendations, and the majority appear in less than one-quarter of documents. Even among the controls that appeared in all countries’ recommendations, we observed high variance in the specific details and implementation advice for how organizations should effectively deploy the security control(s). In the most extreme cases, advice between countries was directly contradictory (§4.5).

Despite many globally shared threats (e.g., ransomware and business email compromise) [13], there may be country-specific cultural or geopolitical context that shapes advice. Some industry surveys have found mild-to-moderate global differences in how worried practitioners are about threats [14] and how they perceive security training and culture [19].

Some of the differences we observe might be influenced by these factors. Nevertheless, we are unaware of work showing globally-differing efficacy of technical measures against global threats. Indeed, we observe allied countries spanning continents and cultures releasing joint advisories, and many documents cite foreign sources for inspiration or supplementary reading. We posit that the striking content differences we observe are likely not due to cultural factors alone.

Instead, we hypothesize that disagreement stems, at least in part, from a lack of a rigorous understanding of how well security controls work in practice, leading to advice driven by anecdotal evidence and experts' folk wisdom. Beyond security efficacy, we also have little insight into the costs and burdens imposed by security measures. As a result, government agencies and the broader security community lack any scientific model or empirical basis for determining how many and which security controls should be recommended. Our work concretely illustrates one likely consequence of this problem—a widespread lack of consensus across ten countries' essential security recommendations—and serves as a call for empirical research that rigorously quantifies the efficacy, costs, and relative trade-offs of security controls.

Second, looking beyond which security controls qualify as “essential,” our work demonstrates that government agencies present and structure advice differently. As shown in §3.3 and §4.4, guidance varies in format, linguistic complexity, and detail. Whereas some countries mention a security control, outline its general importance, and provide only high level advice (e.g., patch vulnerabilities on critical systems periodically), other countries elaborate extensively on conditions for when the advice applies and how to implement it. Furthermore, almost every country we studied publishes a range of security advice, from general “essentials” to more specific “targeted” guidance that covers a particular audience or threat. However, countries vary widely in what topics their agencies chose to publish targeted guidance for.

These differences in presentation and scoping of guidance highlight the need for the security community to better understand how best to present security advice and construct security advice frameworks. For example, how are organizations interacting with and using this advice? What level of detail and presentation would be most effective for practitioners: do security teams find more detail helpful, or treat it as an overly restrictive checklist? And is it feasible to provide the desired level of detail while ensuring advice remains accurate over time? Moreover, the bevy of “targeted” advice guidance, aimed at specific threats or audiences, raises the question of whether it is even possible to design a single list of “essential” security controls for organizations, as opposed to a body of more custom and specific advice; or whether instead these “targeted” guidance documents are largely redundant. Finally, we also acknowledge that one size may not fit all: cultural and social context may shape how guidance choices such as presentation and tone are perceived. Future work should explore

how these factors should influence guidance design.

Ultimately, our research demonstrates that the governments we studied lack a clear consensus for what security measures enterprises should deploy. Further, even when there exist agreed-upon controls, sources differ on how to structure and present advice. Our work serves as a call-to-action for the community to develop strong empirical evidence for what advice is most effective for enterprises and how to best structure this advice for meaningful impact. Absent this understanding, we anticipate that the landscape of enterprise security advice will remain fractured, discordant, and even contradictory, leaving organizations without clear and scientifically-sound direction for how to best secure themselves.

Acknowledgments

We thank Olga Livingston, Tara Dixit, Parker Ruth, Natalia Iwach, our anonymous shepherd, and members of the Stanford Empirical Security Research Group for valuable discussions, feedback, and support. This work was supported in part by a Sloan Research Fellowship, the National Science Foundation under Grant Number #2319080. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF or other funding organizations.

Open Science

For open science, we have released the following research artifacts at <https://doi.org/10.5281/zenodo.15612457>, enabling full reproducibility of our work:

- Our dataset of links to countries' guidance resources, sorted by type, with associated metadata (§ 3.1 and §3).
- Our hierarchical taxonomy and our mappings of 10 guidance documents (§4.1 and § 4).
- Our code for our tree edit distance metric (§4.2.2).

Ethics Considerations

Our work does not raise ethical concerns. Our data collection methodology was rate-limited by manual effort and did not generate high web traffic volumes that could overwhelm a studied website. This paper involved only publicly available data, without sensitive or identifying information about individuals or groups. We acknowledge our position as US-based researchers studying guidance from regions of the world where we do not reside; we refrain from passing judgment on whose advice is “best” and focus our investigation on consensus between sources. Our work presents no clear harms, and we anticipate only positive societal impacts: we hope that gathering and analyzing our dataset of global enterprise security guidance can prompt governments to rethink and improve their guidance, and therefore enterprise security as a whole.

References

- [1] Yasemin Acar, Christian Stransky, Dominik Wermke, Charles Weir, Michelle L Mazurek, and Sascha Fahl. Developers need support, too: A survey of security advice for software developers. In *IEEE Cybersecurity Development (SecDev)*, 2017.
- [2] Adobe. The Adobe common controls framework (CCF). <https://www.adobe.com/trust/compliance/adobe-ccf.html>, 2024.
- [3] Priyanka Badva, Kopo M Ramokapane, Eleonora Pantano, and Awais Rashid. Unveiling the hunter-gatherers: Exploring threat hunting practices and challenges in cyber defense. In *USENIX Security Symposium*, 2024.
- [4] Karoline Busse, Julia Schäfer, and Matthew Smith. Replication: No one can hack my mind revisiting a study on expert and non-expert security practices and advice. In *SOUPS*, 2019.
- [5] SCF Council. Secure controls framework. <https://securecontrolsframework.com>, 2024.
- [6] Savino Dambra, Leyla Bilge, and Davide Balzarotti. SoK: Cyber insurance–technical challenges and a system security roadmap. In *IEEE Symposium on Security and Privacy*, 2020.
- [7] Myriam Dunn Cavelty and Andreas Wenger. Cyber security meets security politics: Complex technology, fragmented politics, and networked science. *Contemporary Security Policy*, 41(1):5–32, 2020.
- [8] e-Governance Academy. National cyber security index (NCSI). <https://ncsi.ega.ee/ncsi-index/?order=rank&type=c>, 2024.
- [9] Michael Fagan and Mohammad Maifi Hasan Khan. Why do they do what they do?: A study of what motivates users to (not) follow computer security advice. In *SOUPS*, 2016.
- [10] Tobias Fiebig. How to stop crashing more than twice: A clean-slate governance approach to IT security. In *IEEE European Symposium on Security and Privacy Workshops*, 2020.
- [11] Rudolph Flesch. A new readability yardstick. *Journal of applied psychology*, 32(3):221, 1948.
- [12] United Nations Institute for Disarmament Research. Cyber policy portal. <https://cyberpolicyportal.org/>.
- [13] World Economic Forum. Global cybersecurity outlook 2025. https://reports.weforum.org/docs/WEF_Global_Cybersecurity_Outlook_2025.pdf, 2025.
- [14] Kevvie Fowler, Kevin Urbanowicz, William Burns, Diana Kearns-Manolatos, and Iram Parveen. Cybersecurity threats and incidents differ by region. <https://www2.deloitte.com/us/en/insights/topics/cyber-risk/global-cybersecurity-threat-trends.html>, 2023.
- [15] Marc Getzoff. Most technologically advanced countries in the world 2023. <https://gfmag.com/data/non-economic-data/most-advanced-countries-in-the-world/>, 2023.
- [16] World Bank Group. Population, total. https://data.worldbank.org/indicator/SP.POP.TOTL?most_recent_value_desc=true, 2023.
- [17] Julian PT Higgins and Sally Green. Cochrane handbook for systematic reviews of interventions, 2008.
- [18] Iulia Ion, Rob Reeder, and Sunny Consolvo. “...No one can hack my mind”: Comparing expert and Non-Expert security practices. In *SOUPS*, 2015.
- [19] Ivanti. International inconsistencies: How cybersecurity preparedness varies across countries. <https://www.ivanti.com/blog/international-inconsistencies-how-cybersecurity-preparedness-varies-across-countries>, 2023.
- [20] Zheng Jia, Xudong Lu, Huilong Duan, and Haomin Li. Using the distance between sets of hierarchical taxonomic clinical concepts to measure patient similarity. *BMC medical informatics and decision making*, 2019.
- [21] Kailani R Jones, Dalton A Brucker-Hahn, Bradley Fidler, and Alexandru G Bardas. Work-from-home and COVID-19: trajectories of endpoint security management in a security operations center. In *USENIX Security Symposium*, 2023.
- [22] JP Kincaid. Derivation of new readability formulas (Automated Readability Index, Fog Count and Flesch Reading Ease Formula) for Navy enlisted personnel. *Chief of Naval Technical Training*, 1975.
- [23] Faris Bugra Kokulu, Ananta Soneji, Tiffany Bao, Yan Shoshitaishvili, Ziming Zhao, Adam Doupe, and Gail-Joon Ahn. Matched and mismatched SOCs: A qualitative study on security operations center issues. In *ACM CCS*, 2019.
- [24] Kommersant. The bot is not a friend of the ministry. <https://www.kommersant.ru/doc/6932790>.
- [25] Pramodh Krishna D. and Venu Gopal Rao K. Generalized weighted tree similarity algorithms for taxonomy trees. *EURASIP Journal on Information Security*, 2016.
- [26] John Kroll. Digging deeper into the 5 W’s of journalism. <https://ijnnet.org/en/story/digging-deeper-5-ws-journalism>, 2018.
- [27] William P Maxam III and James C Davis. An interview study on third-party cyber threat hunting processes in the US Department of Homeland Security. *arXiv:2402.12252*, 2024.
- [28] Matt McVicar, Benjamin Sach, Cédric Mesnage, Jeffrey Lijffijt, Eirini Spyropoulou, and Tjil De Bie. SuMoTED: An intuitive edit distance between rooted unordered uniquely-labelled trees. *Pattern Recognition Letters*, 2016.
- [29] Tyler Moore, Scott Dynes, and Frederick R. Chang. Identifying how firms manage cybersecurity investment. In *WEIS*, 2016.
- [30] James Nicholson, Lynne Coventry, and Pamela Briggs. “If it’s important it will be a headline”: Cybersecurity information seeking in older adults. In *CHI*, 2019.
- [31] Jason RC Nurse, Louise Axon, Arnau Erola, Ioannis Agrafiotis, Michael Goldsmith, and Sadie Creese. The data that drives cyber insurance: A study into the underwriting and claims processes. In *International conference on cyber situational awareness, data analytics and assessment*, 2020.
- [32] Arsenio Paez. Gray literature: An important resource in systematic reviews. *Journal of Evidence-Based Medicine*, 2017.

- [33] PasswordManagers.co. Cybersecurity exposure index (CEI). <https://passwordmanagers.co/cybersecurity-exposure-index/>, 2020.
- [34] Sean Peisert, Ed Talbot, and Matt Bishop. Turtles all the way down: a clean-slate, ground-up, first-principles approach to secure systems. In *Proceedings of the 2012 New Security Paradigms Workshop*, pages 15–26, 2012.
- [35] Elissa Redmiles, Lisa Maszkiewicz, Emily Hwang, Dhruv Kuchhal, Everest Liu, Miraida Morales, Denis Peskov, Sudha Rao, Rock Stevens, Kristina Gligorić, et al. Comparing and developing tools to measure the readability of domain-specific texts. In *EMNLP-IJCNLP*, 2019.
- [36] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. How I learned to be secure: a census-representative survey of security advice sources and behavior. In *CCS*, 2016.
- [37] Elissa M. Redmiles, Amelia R. Malone, and Michelle L. Mazurek. I think they’re trying to tell me something: Advice sources and selection for digital security. In *IEEE Symposium on Security and Privacy*, 2016.
- [38] Elissa M. Redmiles, Noel Warford, Amritha Jayanti, Aravind Koneru, Sean Kross, Miraida Morales, Rock Stevens, and Michelle L. Mazurek. A comprehensive quality evaluation of security and privacy advice on the web. In *USENIX Security Symposium*, 2020.
- [39] Robert W Reeder, Iulia Ion, and Sunny Consolvo. 152 simple steps to stay safe online: Security advice for non-tech-savvy users. *IEEE Security & Privacy*, 2017.
- [40] Hannah R Rothstein and Sally Hopewell. Grey literature. *The handbook of research synthesis and meta-analysis*, 2009.
- [41] Hannah R. Rothstein, Herbert M. Turner, and Julia G. Lavenberg. Information retrieval policy brief, 2004.
- [42] Anna Lena Rotthaler, Harshini Sri Ramulu, Lucy Simko, Sascha Fahl, and Yasemin Acar. “It’s time. Time for digital security.”: An end user study on actionable security and privacy advice. In *IEEE Symposium on Security and Privacy*, 2025.
- [43] Kimberly Ruth, Veronica Rivera, Gautam Akiwate, Aurore Fass, Patrick Gage Kelley, Kurt Thomas, and Zakir Durumeric. “Perfect is the enemy of good”: The CISO’s role in enterprise security as a business enabler. In *CHI*, 2025.
- [44] Johnny Saldaña. *The coding manual for qualitative researchers*. SAGE publications Ltd, 2021.
- [45] Stef Schinagl and Abbas Shahim. What do we know about information security governance? “From the basement to the boardroom”: towards digital security governance. *Information & Computer Security*, 2020.
- [46] Stef Schinagl, Abbas Shahim, and Svetlana Khapova. Paradoxical tensions in the implementation of digital security governance: Toward an ambidextrous approach to governing digital security. *Computers & Security*, 122:102903, 2022.
- [47] Abe Singer and Matt Bishop. Trust-based security; or, trust considered harmful. In *New Sec. Paradigms Workshop*, 2020.
- [48] Rebecca Slayton. Governing uncertainty or uncertain governance? Information security and the challenge of cutting ties. *Science, Technology, & Human Values*, 2021.
- [49] Raghavendra Sridharamurthy, Talha Bin Masood, Adhitya Kamakshidasan, and Vijay Natarajan. Edit distance between merge trees. *IEEE Trans. on Vis. and Comp. Graphics*, 2018.
- [50] Rock Stevens, Josiah Dykstra, Wendy Knox Everette, James Chapman, Garrett Bladow, Alexander Farmer, Kevin Halliday, and Michelle L Mazurek. Compliance cautions: Investigating security issues associated with US digital-security standards. In *NDSS*, 2020.
- [51] Rock Stevens, Faris Bugra Kokulu, Adam Doupé, and Michelle L Mazurek. Above and beyond: Organizational efforts to complement US digital security compliance mandates. In *NDSS*, 2022.
- [52] Andrea Torsello, Antonio Robles-Kelly, and Edwin R Hancock. Discovering shape classes using tree edit-distance and pairwise clustering. *International Journal of Computer Vision*, 2007.
- [53] International Telecommunications Union. Global cybersecurity index. <https://www.itu.int/pub/D-STR-GCI.01>, 2020.
- [54] U.S. CISA. ASD’s ACSC, CISA, and Partners Release Secure by Design Guidance on Choosing Secure and Verifiable Technologies. <https://www.cisa.gov/news-events/alerts/2024/05/09/asds-acsc-cisa-and-partners-release-secure-design-guidance-choosing-secure-and-verifiable>.
- [55] U.S. CISA. CISA, NCSC-UK, and Partners Release Advisory on Russian SVR Actors Targeting Cloud Infrastructure. <https://www.cisa.gov/news-events/alerts/2024/02/26/cisa-ncsc-uk-and-partners-release-advisory-russian-svr-actors-targeting-cloud-infrastructure>.
- [56] Michel Van Eeten. Patching security governance: an empirical view of emergent governance mechanisms for cybersecurity. *Digital Policy, Regulation and Governance*, 19(6), 2017.
- [57] Julia Voo, Irfan Hemani, and Daniel Cassidy. National cyber power index (NCPI). <https://www.belfercenter.org/publication/national-cyber-power-index-2022>, 2022.
- [58] C Anne Wade, Herbert M Turner, Hannah R Rothstein, and Julia G Lavenberg. Information retrieval and the role of the information specialist in producing high-quality systematic reviews in the social, behavioural and education sciences. *Evidence & Policy*, 2006.
- [59] Susan P Williams, Catherine A Hardy, and Janine A Holgate. Information security governance practices in critical infrastructure organizations: A socio-technical and institutional logic perspective. *Electronic Markets*, 23:341–354, 2013.
- [60] Flynn Wolf, Adam J Aviv, and Ravi Kuber. Security obstacles and motivations for small businesses from a CISO’s perspective. In *USENIX Security Symposium*, 2021.
- [61] Affan Yasin, Rubia Fatima, Lijie Wen, Wasif Afzal, Muhammad Azhar, and Richard Torkar. On using grey literature and Google Scholar in systematic literature reviews in software engineering. *IEEE Access*, 2020.
- [62] Kaizhong Zhang and Dennis Shasha. Simple fast algorithms for the editing distance between trees and related problems. *SIAM journal on computing*, 1989.

[63] Yinjia Zhang, Qinpei Zhao, Yang Shi, Jiangfeng Li, and Weixiong Rao. Category tree distance: a taxonomy-based transaction distance for web user analysis. *Data Mining and Knowledge Discovery*, 37(1):39–66, 2023.

[64] Yang Zhong, Christopher A Meacham, and Sakti Pramanik. A general method for tree-comparison based on subtree similarity and its use in a taxonomic database. *Biosystems*, 1997.

[65] Yixin Zou, Kevin Roundy, Acar Tamersoy, Saurabh Shintre, Johann Roturier, and Florian Schaub. Examining the adoption and abandonment of security, privacy, and identity theft protection practices. In *CHI*, 2020.

A List of Countries In Scope

Table 3 lists all countries selected for our search.

Continent	Country	Code	Cont.	Country	Code
NA	Canada	CA	NA	United States	US
	Mexico	MX			
SA	Brazil	BR			
EU	Austria	AT	EU	Netherlands	NL
	Belgium	BE		Norway	NO
	Czechia	CZ		Poland	PL
	Denmark	DK		Russia	RU
	Estonia	EE		Slovakia	SK
	Finland	FI		Spain	ES
	France	FR		Sweden	SE
	Germany	DE		Switzerland	CH
	Lithuania	LT		Ukraine	UA
	Luxembourg	LU		United Kingdom	UK
AS	Bangladesh	BD	AS	Pakistan	PK
	China	CN		Saudi Arabia	SA
	India	IN		Singapore	SG
	Indonesia	ID		South Korea	KR
	Israel	IL		Taiwan	TW
	Japan	JP		UAE	AE
	Malaysia	MY			
AF	Egypt	EG	AF	Nigeria	NG
OC	Australia	AU	OC	New Zealand	NZ

Table 3: List of 41 Countries Searched.

B Example Images of Enterprise Guidance

Figures 7 and 8 show several examples of enterprise guidance.

C Example Guidance Mapping

Figure 9 shows a small example of mapping a guidance document into a tree structure using our framework. For clarity, we use a selection of controls from Singapore and show only a small portion of our framework.

D First Codebook Level

Table 4 shows the first level of the codebook that we built deductively from preexisting common controls frameworks.



Figure 7: Example of Long-Tail Sector-Specific Guidance. Cover page and representative interior page for Israel’s controlled chicken coop guidance.

This forms the theme level of our five-level analysis framework. Our full codebook, forming the first three layers of the taxonomy, is available as part of our Open Science artifacts.

Governance	Human Resources Security
AI Risk	Identity and Access Management
Asset Management	Incident Response
Business Continuity and Disaster Recovery	Maintenance
Change Management	Network Security
Cloud Security	Physical Security
Compliance and Auditing	Risk Management
Configuration Management	Security Operations
Logging, Monitoring, and Alerting	Security Engineering and Secure Software Development
B2B Customer Security	Security Training and Awareness
Enablement	Third-Party Management
Cryptography	Threat Hunting and Threat Intelligence
Data Management	Vulnerability and Patch Management
Embedded Systems Security	Web and Browser Security
Endpoint and Mobile Device Management	Other

Table 4: Top Level of Codebook.

E List of Primary Guidance Documents

Table 5 shows primary document(s) from each of the 41 countries we investigated, along with their authoring agency, language, and intent (see Section 3.3).

F Essential Guidance Document Metadata

Table 6 shows where each of the essential-guidance documents fall on a range of design dimensions. We use these results to choose 10 documents to analyze further in Section 4.





Figure 8: **Guidance Presentation.** Examples of the varied ways that countries format/present their guidance to the public.

ID	Control Text
A.2.4(a)	An up-to-date asset inventory of all the hardware and software assets shall be maintained in the organisation. Organisations may meet this requirement in different ways, e.g., use of spreadsheet or IT asset management software to maintain the IT asset inventory.
A.2.4(l)	Before disposing of any hardware asset, the organisation shall ensure that all confidential information have been deleted, e.g., encrypting hard disk before reformatting and overwriting it.
A.8.4(a)	The organisation shall identify business-critical systems and those containing essential business information and perform backup. What needs to be backed up is guided by identifying what is needed for business recovery in the event of a cybersecurity incident.
A.8.4(d)	The backup process should be automated where feasible.
A.8.4(h)	Backups shall be stored separately (i.e., offline) from the operating environment. Where feasible, backups should be stored offsite, e.g., separate physical location.



(b) Mini-taxonomy mapping

Figure 9: **Mapping Example.** Example of mapping a guidance document (here, a subset of controls from Singapore) into our taxonomy (small portion shown for clarity). Taxonomy nodes shaded gray are part of mini-Singapore’s tree.

Co.	Agency	Primary Document	Language	Intent
Australia	ASD	Essential Eight Maturity Model	English	Essential
Austria	A-SIT	Austrian Information Security Handbook	German	Catalog
Bangladesh	BGD e-GOV CIRT	Security Best Practices	English	Essential
	ICTD	Government of Bangladesh Information Security Manual	English	Requirement
Belgium	CCB	Safeonweb@work CyberFundamentals Framework	English	Essential
Brazil	GSI/PR	Information security management manual	Portuguese	Not controls
Canada	CCCS	Top 10 IT security actions to protect Internet connected networks and information	English	Essential
China	SAC/TC260	fied criteria for security protection of computer information system	Chinese	Catalog
	SAC/TC260	Common security techniques requirement for information system	Chinese	Requirement
	SAC/TC260	Information security technology — Baseline for classified protection of cybersecu- rity	Chinese	Requirement
Czechia	NÚKIB	Recommendations for administrators, version 4.0	Czech	Essential
Denmark	CFCS	Effective Cyber Defense	Danish*	Essential
Egypt	EG-CERT	Potential cyber threats	Arabic	Essential
Estonia	RIA	Estonian Information Security Standard (E-ITS) - Catalog of measures	Estonian	Catalog
Finland	Traficom	Strengthening cyber security in Finnish organizations	Finnish*	Essential
	Traficom	Guide to cyber security and company board responsibility	Finnish	Not controls
France	ANSSI	Guideline for a healthy information system in 42 measures	French*	Essential
Germany	BSI	Basic Measures of Cybersecurity	German	Essential
India	MeitY	General Guidelines for Secure Application and Infrastructure	English	Essential
	CERT-IN	System Security Guidelines	English	Catalog
Indonesia	BSSN	Information security self-assessment (Paman Kami) for SMEs	Indonesian*	Essential
Israel	INCD	The Ten Recommendations that Organizations Should Adopt	Hebrew	Essential
	INCD	Cyber Defense Doctrine 2.0	Hebrew*	Not controls
Japan	IPA	Cybersecurity management guidelines v3.0	English	Essential
	IPA	Cybersecurity Management Guidelines v3.0 Practices for Implementation 4th Ed.	Japanese	Essential
Lithuania	NCSC	Security control measures (Critical Controls, CC)	Lithuanian	Essential
Luxembourg	LHC	Cybersecurity Essentials	English	Essential
	LHC	Be prepared at all times - Test and Improve	English	Essential
	LHC	Common incidents - Detect & React	English	Essential
Malaysia	NACSA	Guidelines and Best Practices for Business	English	Requirement
	NACSA	National Cyber Security Baseline (NCSB)	English	Essential
Mexico	SPF	Cybersecurity Guide for Public Facilities	Spanish	Essential
Netherlands	NCSC	5 basic principles of digital resilience	Dutch	Essential
New Zealand	CERTNZ	Critical Controls	English	Essential
	NCSC	Cyber Security Framework	English	Not controls
	NCSC	Own Your Online - Top Online Security Tips for Your Business	English	Essential
Norway	NSM	Basic Principles for ICT Security	Norwegian*	Essential
Pakistan	NTISB	IT Security Guide Book	English	Requirement
Poland	NSC	Security and Privacy Protection for Information Systems and Organizations (800- 53)	Polish	Catalog
	NSC	Baseline Security for Information Systems and Organizations	Polish	Catalog
Saudi Arabia	NCA	Essential Cybersecurity Controls (ECC)	English	Requirement
Singapore	CSA	Cyber Essentials	English	Essential
	CSA	Cyber Trust Mark	English	Requirement
Slovakia	CSIRT	Methodology for systematic security of public administration organizations in the field of information security	Slovak	Requirement
	CSIRT	Critical safety measures	Slovak	Essential
South Korea	KISA	CISO Guide Basics	Korean	Not controls
Spain	INCIBE	IMC: Indicators for Improving Cyber Resilience	Spanish	Catalog
Sweden	MSB	Security measures for information systems	English	Requirement
	NCSC	Cyber security in Sweden 2022 Part 2: Recommended security measures	Swedish	Essential
Switzerland	BACS	Minimum standard for improving ICT resilience	English	Essential
Taiwan	NICS	Reference Guidelines on Security Control Measures	Chinese	Catalog
UAE	TRA	Information Assurance Regulations	English	Requirement
Ukraine	CERT-UA	Basic Rules of Cyber Hygiene	Ukrainian	Essential
United Kingdom	NCSC	Cyber Essentials	English	Essential
	NCSC	10 Steps to Cybersecurity	English	Essential
United States	CISA	Cross-sector Cybersecurity Performance Goals (CPGs)	English	Essential
	NIST	Cybersecurity Framework (CSF) 2.0	English	Catalog

Table 5: **Primary Document(s) From Each Country.** Nigeria and Russia are omitted since we found nothing matching our definition of a primary document. *Official English translation available also.

Cont.	Co.	Document	Year	International Refs	Words	Grade	Controls	Levels	“Why”
OC	Australia	Essential Eight Maturity Model	2023	-	11876	16	152	3	No
AS	Bangladesh	Security Best Practices	2023	-	370	13	43	1	No
EU	Belgium	CyberFundamentals Framework	2023	NIST, ISO, IEC, CIS	6168	14	38	2	No
NA	Canada	Top 10 IT security actions to protect Internet connected networks and information	2021	NIST	4604	13	10	1	No
EU	Czechia	Recommendations for administrators	2020	-	1395*	-	46	1	No
EU	Denmark	Effective Cyber Defense	2023	ITIL, CIS, NSA, AU, OECD, CA, UK, NO	5446	14	15	1	Yes
AF	Egypt	Potential cyber threats	2022	CISA, CCCS	842*	-	13	1	No
EU	Finland	Strengthening cyber security in Finnish organizations	2022	MITRE, CISA, NIST, NSA, MICROSOFT, AWS, GCP, UK, IE	1858	16	18	1	Yes
EU	France	Guideline for a healthy information system	2017	-	10061	18	42	2	Yes
EU	Germany	Basic Measures of Cybersecurity	2018	-	4727*	-	29	1	Yes
AS	India	General Guidelines for Secure Application and Infrastructure	2017	-	237	14	10	1	No
AS	Indonesia	Information security self-assessment (Paman Kami) for small and medium enterprises	2020	ISO, NIST, AU, JP	7009	12	79	1	Yes
AS	Israel	The Ten Recommendations that Organizations Should Adopt	2020	-	419*	-	10	1	No
AS	Japan	Cybersecurity management guidelines	2023	ISO, NIST, CIS	15206	18	13	1	Yes
		Cybersecurity Management Guidelines Practices for Implementation	2023	OWASP, CIS, FIRST	35735*	-	31	1	Yes
EU	Lithuania	Security control measures (Critical Controls, CC)	2015	SANS	1318*	-	20	2	No
		Cybersecurity Essentials	2021	-	1647	15	68	1	Yes
EU	Luxembourg	Be prepared at all times - Test and Improve	2021	ISO	1328	16	8	1	No
		Common incidents - Detect & React	2021	-	1387	13	13	1	No
AS	Malaysia	National Cyber Security Baseline (NCSB)	2024	NIST, ISO	2806	13	33	1	No
NA	Mexico	Cybersecurity Guide for Public Facilities	2018	NIST, ISO, IETF, ENISA, ITU, AXIS, ES, KASPERSKY	11368*	-	14	1	Yes
EU	Netherlands	5 basic principles of digital resilience	2024	ISO	4580*	-	27	1	Yes
OC	New Zealand	Critical Controls	2023	-	23025	12	192	1	Yes
		Top Online Security Tips for Your Business	2024	-	2096	11	11	1	Yes
EU	Norway	Basic Principles for ICT Security	2024	ISO, ITIL, CIS, UK NCSC, NIST, ASD	19662	14	118	1	Yes
AS	Singapore	Cyber Essentials	2022	ISO, CIS, CISA, NIST, HiTrust, PCI, SOC, CA, UK, US, AU	6657	14	78	1	Yes
EU	Slovakia	Critical safety measures	2012	SANS	13749*	-	66	4	Yes
EU	Sweden	Cyber security in Sweden 2022 Part 2: Recommended security measures	2022	ISO, ITIL, CIS, UK NCSC, ASCS/ASD	9515*	-	11	1	Yes
EU	Switzerland	Minimum standard for improving ICT resilience	2023	NIST, ISO, COBIT, ENISA, BSI	14196	13	106	1	Yes
EU	Ukraine	Basic Rules of Cyber Hygiene	2018	-	808*	-	15	1	No
EU	United Kingdom	Cyber Essentials	2023	-	4773	15	41	1	Yes
		10 Steps to Cybersecurity	2021	-	11132	14	43	1	Yes
NA	United States	Cybersecurity Performance Goals	2023	-	4328	18	38	1	Yes

Table 6: **Document Metadata.** Metadata for “essential” guidance documents. Countries vary widely along every dimension we examined, including length (by word count and enumerated control count), explanatory goals (i.e., describing purpose of controls), complexity and structure (reading ease, maturity levels), age, and influences from international or overseas guidance sources. From these documents, we selected 10 for further analysis (in bold). *Word counts are approximate for non-English documents since they are computed from English translations.

G Table 1 Correction

This paper version corrects minor errors in Table 1. The original published table is shown in Table 7. Consequent textual updates are also documented below.

- nearly a third had five or more documents with this same general-purpose scope
- + about a third had five or more documents with this same general-purpose scope

- The vast majority of countries (37/41) published at least some guidance about securing particular types of technologies, but specific coverage varied substantially. The top two technologies were cloud systems (25 countries) and network appliances (22), followed by cryptographic systems (19), web, and mobile systems (each 18).
- + The vast majority of countries (39/41) published at least some guidance about securing particular types of technologies, but specific coverage varied substantially. The top two technologies were cloud systems (27 countries) and network appliances (23), followed by cryptographic systems, web, and mobile systems (each 20).

- guidance targeted to particular industry sectors (34/41)
- + guidance targeted to particular industry sectors (36/41)

- The most common sectors were government and defense (22 countries) and finance (13).
- + The most common sectors were government and defense (25 countries) and finance (13).

- (e.g., Japan, Lithuania, and Czechia released guidance for factories, hotels, and satellites, respectively, which most countries omit).
- + (e.g., Lithuania and Czechia released guidance for hotels and satellites, respectively, which most countries omit).

- guidance addressed to individual employees or citizens (35 countries)
- + guidance addressed to individual employees or citizens (37 countries)

- top themes of targeted guidance as incident response (24 countries), IAM and authentication (22), and risk management (15)
- + top themes of targeted guidance as incident response (24 countries), IAM and authentication (24), and risk management (16)

H Artifact Appendix

H.1 Abstract

Governments globally publish a wealth of cybersecurity guidance for enterprises; however, little prior work has studied this guidance ecosystem. This paper presents the first systematic investigation of governments’ enterprise security guidance. We build a corpus of prominent guidance documents from 41 countries, develop a tree-based content taxonomy and quantitative comparison metric, and compare a selection of ten guidance documents’ content. Our analysis shows that governments differ widely in what and how they advise companies.

H.2 Description & Requirements

Our artifact materials contain the following main components:

- `Guidance document links.xlsx`: Our dataset of links to countries’ guidance resources, sorted by type. Functionally, this spreadsheet is a qualitative codebook constructed during our document search process as described in Section 3.1.2. For each resource we found during our search, we determined its category according to the high-level document typology we outline in Section 3.2, and added its URL directly to the corresponding spreadsheet column. We then added columns to provide a more granular breakdown of sector-, tech-, threat-, and mitigation-specific documents (using our qualitative observations of common elements within each of those four categories).
- `Roots of shrubs.xlsx` and `v2 Shrubs to trees.xlsx`: Our hierarchical taxonomy and our mappings of ten guidance documents. The former spreadsheet contains the first three levels of our taxonomy; the latter spreadsheet contains the final two levels.
- `tree_similarity.ipynb`: Our code for our custom tree edit distance metric. Usage and supplementary files are described in `README.md`.

In addition, our `v2 Shrubs to trees.xlsx` document contains supplementary data in the tabs `themes_count`, `subthemes_count`, and `controls_count`. These are simple raw counts over the data structure that `tree_similarity.ipynb` constructs (with the addition of “Other” subthemes and controls). The data structure records, for each country document, which nodes in our taxonomic tree the document contains content for. The supplementary data counts how many country documents contain content for each node in the taxonomic tree.

H.2.1 Security, privacy, and ethical concerns

Our artifact materials involve no security, privacy, or ethical concerns for evaluators. This work only involves analysis of publicly available data, without sensitive or identifying

	Country	Agencies	General Purpose	Timely	Size-specific	Critical Infra	Sector Specific						Tech Specific								Threat Model Specific						Mitigation Specific															
							Any	Govt./Defense	Telecom/IAP	Finance	Transportation	Health	Education	Any	Cloud	Networks	Web	Mobile	Crypt.	IoT	OT	AI	Any	(D)DoS	Ransomware	Phishing/spam	Malware	Remote/BYOD	Data breach/loss	Any	IR	IAM/Auth	Risk Mgmt.	Governance	Logs/alerts	Network defenses	For individuals	Pro. development				
	Australia	4	6	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
	Austria	8	2	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Bangladesh	5	2	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Belgium	2	4	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Brazil	2	1	●	○	○	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Canada	2	6	●	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	China	2	5	●	○	○	●	●	●	●	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Czechia	4	1	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Denmark	4	7	●	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Egypt	2	1	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Estonia	1	3	●	○	○	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Finland	2	3	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	France	4	9	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Germany	2	4	●	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	India	10	4	●	○	○	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Indonesia	3	0	●	●	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Israel	3	5	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Japan	5	3	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Lithuania	1	2	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Luxembourg	6	3	●	●	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Malaysia	1	3	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Mexico	0	6	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Netherlands	5	2	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	New Zealand	8	6	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Nigeria	5	0	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Norway	1	3	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Pakistan	2	1	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Poland	1	6	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Russia	0	0	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Saudi Arabia	3	3	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Singapore	1	5	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Slovakia	3	3	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	S. Korea	3	1	●	●	○	●	●	●	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Spain	5	10	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Sweden	4	4	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Switzerland	5	3	●	●	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Taiwan	2	5	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	UAE	3	3	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	Ukraine	8	1	●	○	○	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	
	UK	3	7	●	●	●	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
	US	36	18	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
	# Countries	39	37	40	23	29	34	22	12	13	10	12	8	37	25	22	18	18	19	16	17	17	33	24	22	19	18	15	14	35	24	22	15	14	14	14	14	37	19			

Table 7: **Guidance Availability.** Nearly every country releases general-purpose guidance as well as more targeted documents.

information about individuals or groups. Our code is non-destructive and can be run safely on evaluators’ machines.

H.2.2 How to access

All artifact materials are accessible through Zenodo at the following URL: <https://doi.org/10.5281/zenodo.15612457>

H.2.3 Hardware dependencies

None.

H.2.4 Software dependencies

Our code requires `pygraphviz` as well as a suite of supplementary Python packages viewable in the `requirements.txt` file. Full installation instructions, using `uv` for dependency management, are provided in the repository’s README.

H.2.5 Benchmarks

Our tree similarity code requires two specially-formatted input data files: one capturing the first three levels of our content taxonomy, the other capturing the remaining two levels. Both data files are provided in our artifact materials, as `Roots of shrubs.xlsx` and `v2 Shrubs to trees.xlsx` respectively.

H.3 Set-up

H.3.1 Installation

Download all files from the artifact repository and place them together into the same directory. To install dependencies for our tree similarity code:

1. Install `pygraphviz`. On Ubuntu LTS:

```
sudo apt-get install graphviz graphviz-dev
pip install pygraphviz
```

Non-Ubuntu installation instructions and troubleshooting tips are linked from the repository’s README.
2. We suggest `uv` for dependency management. To install `uv`:

```
curl -LsSf https://astral.sh/uv/install.sh | sh
```
3. Install the remaining Python dependencies listed in `requirements.txt`. Using `uv`:

```
uv sync
```

Or, if running from within a Python `venv`:

```
uv sync --active
```

H.3.2 Basic Test

For our dataset of government guidance document links, open and view `Guidance document links.xlsx`. The file should contain one row of data for each of the 41 countries we included in our search.

For our tree similarity code, create an `out/` directory in the same location as the code, then open and run `tree_similarity.ipynb` in Jupyter Lab. Key functions take the two input file names (defaulting to `Roots of shrubs.xlsx` and `v2 Shrubs to trees.xlsx`) as parameters. The code should generate the following output files in the `out/` directory:

1. `pairwise_tree_differences.csv`: Raw (non-normalized) tree distances between pairs of documents.
2. `doc_tree_to_full_tree_cost.csv`: Raw (non-normalized) tree distance from each individual document’s tree to the total taxonomic tree, used for computing coverage numbers.
3. `tree_*.png` and `full_tree.png`: Visualizations of each document’s content tree as well as of the total taxonomic tree.
4. `venn_us_au_uk.png`: Venn diagram showing overlap in controls from three key allied countries.

H.4 Evaluation workflow

Below we outline the major claims of our work (§H.4.1) and describe the steps to validate these major claims (§H.4.2).

H.4.1 Major Claims

- (C1): Governments commonly present guidance in a wide array of differently scoped documents. This is illustrated in Table 1 (Section 3.2) and supported by experiment (E1).
- (C2): Governments vary widely in the breadth and amount of implementation detail they consider “essential” guidance for companies. This is illustrated in Figure 2 (Section 4.3) and supported by experiment (E2).
- (C3): There is little consensus among governments—even between close allies—about what controls are most essential for companies to implement. This is illustrated in Figures 3, 4, 5, and 6 (Section 4.4), and supported by experiments (E3), (E4), and (E5).
- (C4): Governments present implementation details for key controls inconsistently, even exhibiting direct contradictions. This is described in Section 4.5 and supported by experiment (E6).

H.4.2 Experiments

- (E1): [2 human-hours] Verify guidance availability.

Preparation: Open and inspect `Guidance document links.xlsx`.

Execution: Inspect the data for the following ten document categories (spreadsheet column labels parenthesized): General Purpose (F), Timely (AI), Size-Specific (AJ), Critical Infra (AU), Sector Specific - Any (AK), Tech Specific - Any (AV), Threat Model Specific - Any (H), Mitigation Specific - Any (S), For Individuals (BJ), Professional Development (BK).

Results: Compare nonempty spreadsheet cells with filled circles in Table 1 to validate table data. Observe that a majority of rows in the table have guidance available (a filled circle or a number greater than zero) for at least six of these ten guidance document categories.

- (E2): [20 human-minutes + 5 compute-minutes] Validate content trees.

Preparation: Follow installation instructions in §H.3.1.

Execution: Run `tree_similarity.ipynb` as described in §H.3.2.

Results: Inspect the files titled `tree_*.png` in the `out/` directory. Note that each file name contains an identifying two-letter country code. Compare with Figure 2.

- (E3): [20 human-minutes] Verify pairwise country differences.

Preparation: Open `pairwise_tree_differences.csv` from the `out/` directory generated from running `tree_similarity.ipynb` in experiment (E2).

Open v2 Shrubs to trees.xlsx and go to the pairwise_tree_differences tab.

Execution: Verify that the data in pairwise_tree_differences.csv matches that in the pairwise_tree_differences tab. Note that rows may not appear in the same order.

Results: Compare the “Normalized” column numbers in the pairwise_tree_differences tab with the corresponding numbers in Figure 3. Observe that most pairs of country documents differ by over 50%.

(E4): [10 human-minutes] Verify distributions of content coverage.

Preparation: Open v2 Shrubs to trees.xlsx.

Execution: Inspect the tabs labeled themes_count, subthemes_count, and controls_count.

Results: Compare themes_count with Figure 4, and subthemes_count and controls_count with Figure 5. Observe that only a minority of content attains widespread consensus to include.

(E5): [5 human-minutes] Validate allied-countries results.

Preparation/Execution: Open venn_us_au_uk.png

from the out/ directory generated from running tree_similarity.ipynb in experiment (E2).

Results: Compare with Figure 6.

(E6): [30 human-minutes] Validate variance in implementation details.

Preparation: Open v2 Shrubs to trees.xlsx.

Execution: Inspect the tabs labeled with control names, especially Backup Operation, Vulnerability Remediation, and Password Authentication.

Results: Observe the data sparsity in the tables, where each attribute value typically appears in a minority of the guidance documents. Note lines 42-43 in the Password Authentication tab, which is an example of contradictory advice.

H.5 Version

Based on the LaTeX template for Artifact Evaluation V20231005. Submission, reviewing and badging methodology followed for the evaluation of this artifact can be found at <https://secartifacts.github.io/usenixsec2025/>.