USENIX

THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

A First Look at Governments' Enterprise Security Guidance

Kimberly Ruth and Raymond Buernor Obu, *Stanford University;* Ifeoluwa Shode, *Fisk University;* Gavin Li, *Stanford University;* Carrie Gates, *FS-ISAC;* Grant Ho, *University of Chicago;* Zakir Durumeric, *Stanford University*

https://www.usenix.org/conference/usenixsecurity25/presentation/ruth

This artifact appendix is included in the Artifact Appendices to the Proceedings of the 34th USENIX Security Symposium and appends to the paper of the same name that appears in the Proceedings of the 34th USENIX Security Symposium.

August 13–15, 2025 • Seattle, WA, USA

978-1-939133-52-6

Open access to the Artifact Appendices to the Proceedings of the 34th USENIX Security Symposium is sponsored by USENIX.

USENIX Security '25 Artifact Appendix: 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

A Artifact Appendix

A.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.

A.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 our paper (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 our paper (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.

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

A.2.2 How to access

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

A.2.3 Hardware dependencies

None.

A.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 managemment, are provided in the repository's README.

*Work done while visiting Stanford University.

A.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.

A.3 Set-up

A.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
   ```

A.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.

A.4 Evaluation workflow

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

A.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).

A.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 §A.3.1.
Execution: Run `tree_similarity.ipynb` as described in §A.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.

A.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/.