# Comparing Government Hosting Strategies in China and Mexico Through Network Tracing and Geolocation

Samarth Arul
samartharul@u.northwestern.edu
Northwestern University
Evanston, IL, USA

## Abstract

Governments worldwide rely on varying web hosting strategies, with some prioritizing domestic infrastructure while others depend on third-party global providers. Understanding these hosting models is crucial in assessing cross-border dependencies, cybersecurity risks, and digital sovereignty. This study builds on the findings of Kumar et al. (IMC 2024) by analyzing government hosting practices in two contrasting countries, one with a strong domestic hosting preference (China) and another relying heavily on third-party, cross-border cloud providers (Mexico).

We conduct an empirical analysis using traceroute measurements to examine network paths between government websites and their hosting infrastructures. By geolocating IP addresses along these paths, we identify cross-border dependencies, revealing whether a country's government services traverse foreign networks. Our methodology includes selecting a representative set of 20 government-related domains from each country, performing traceroute-based path tracing, and visualizing the routing topology using geolocation mapping techniques.

Our results highlight key differences in government hosting practices, showcasing varying degrees of reliance on foreign infrastructure and intermediaries. The analysis provides insights into the implications of these dependencies, such as potential security concerns and data jurisdiction issues. Furthermore, we discuss the challenges encountered during internet-scale measurements, including data accuracy limitations and network dynamics.

## Keywords

Network Measurement, Government, Third-party, Centralization, Hosting, Geolocation, Cross-border Dependency

**ACM Reference Format:**
Samarth Arul. 2025. Comparing Government Hosting Strategies in China and Mexico Through Network Tracing and Geolocation. In . ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

## 1 Introduction

The internet forms the backbone of modern government services, enabling citizens to access critical information, public records, and administrative resources. However, the infrastructure supporting these digital services varies significantly between countries. Some governments prioritize domestic hosting to maintain data sovereignty and reduce external dependencies, while others rely heavily on global cloud providers for scalability and cost efficiency. The choice of hosting strategy carries implications for security, resilience, and cross-border dependencies.

Recent work by Kumar et al. [1] provides a comprehensive analysis of government hosting strategies worldwide, categorizing nations based on their reliance on domestic infrastructure versus third-party international providers. Their study highlights the varying degrees of external dependencies in government hosting, raising concerns about data jurisdiction, geopolitical risks, and service reliability. Understanding these hosting models is essential in assessing the broader implications of digital sovereignty and infrastructure resilience.

In this paper, we extend the insights of Kumar et al. [1] by conducting an empirical study of cross-border dependencies in government web hosting. We analyze two countries with contrasting hosting strategies—one with a high reliance on domestic infrastructure and another that predominantly uses foreign cloud providers. Using traceroute measurements, we map the network paths between government websites and their hosting locations, geolocating IP addresses along the routes to identify cross-border dependencies. Our analysis provides a detailed examination of how government traffic flows across different jurisdictions, shedding light on potential security and policy concerns.

By combining internet-scale measurement techniques with geolocation analysis, we aim to contribute to the ongoing discussion on government infrastructure resilience and digital autonomy. These findings could offer valuable insights for policymakers, researchers, and network operators seeking to assess and mitigate cross-border risks in government hosting.

## 2 Methodology

To analyze cross-border dependencies in government web hosting, we conducted a measurement study using traceroute-based path tracing and geolocation analysis. This section outlines the dataset selection, country selection criteria, path tracing methodology, and data visualization techniques.

### 2.1 IP Address Dataset

For this study, we utilized the dataset provided in the `vantage_domain_ip_server_map.csv` file. This dataset contains a curated list of government-related domains from multiple countries, along with their associated IP addresses and geolocation data. From this dataset, we selected a representative sample of 20 government


*Conference'17, Washington, DC, USA*

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
https://doi.org/10.1145/nnnnnnn.nnnnnnn

resources per country, ensuring an even distribution of domestically and internationally hosted content. This selection enabled us to analyze variations in hosting strategies and identify potential cross-border dependencies.

## 2.2 Country Selection

We selected China and Mexico as our two case study countries due to their contrasting government hosting strategies, as highlighted in Kumar et al. [1].

China exemplifies a strong domestic hosting model, with the vast majority of government services hosted within national infrastructure. The country enforces strict internet regulations, including the Great Firewall, which limits reliance on foreign hosting providers. This approach enhances data sovereignty and security but also raises concerns about network centralization. Nevertheless, a small subset of government resources are hosted internationally, most prominently in Japan.

In contrast, Mexico demonstrates a more diverse hosting strategy, with a significant portion of government services hosted on international platforms, particularly by U.S.-based cloud providers. This reliance on foreign infrastructure introduces potential vulnerabilities related to data jurisdiction, cross-border routing, and geopolitical risks.

By comparing these two countries, we aim to illustrate how government hosting strategies affect internet routing, external dependencies, and potential cybersecurity concerns.

## 2.3 Path Tracing and Geolocation

To conduct path tracing, we used RIPE Atlas probes via AquaLab's Custom RIPE Atlas API Platform. RIPE Atlas is a global network measurement platform that allows researchers to perform traceroutes from geographically distributed vantage points. Using the provided tool, we executed traceroutes from in-country probes to the selected government websites in China and Mexico. This allowed us to capture the sequence of intermediate hops that packets traverse when reaching their destination.

The traceroute data was processed using a custom Python script, which extracts key routing information, including IP addresses, round-trip times, and ASN (Autonomous System Number) mappings. This structured data provided insight into network paths and potential cross-border traffic flow.

For geolocation, we employed the IP geolocation API `ipinfo.io`, which returns the latitude, longitude, city, and country associated with each IP address. The geolocation results allowed us to map the physical locations of intermediate hops, highlighting instances where government traffic was routed through foreign networks. In cases where API queries failed, we supplemented missing data with MaxMind's GeoLite2 database.

## 2.4 Data Visualization

To present our findings, we generated interactive maps using the Folium library. Each traceroute was visualized as a network path, with nodes representing IP addresses and edges illustrating packet flows between them. Different colors were assigned to each traceroute path from another. The resulting maps provided a clear depiction of cross-border dependencies, revealing patterns of international traffic exchange and foreign infrastructure reliance.

# 3 Results

This section presents our findings on government hosting strategies in China and Mexico, focusing on their hosting profiles, cross-border dependencies, and key observations derived from traceroute analysis.

## 3.1 Hosting Profiles

Our analysis revealed stark contrasts between the hosting practices of China and Mexico. As expected, China maintained a strong preference for domestic hosting, with the majority of government services hosted within its national infrastructure. However, a small subset of services was found to be hosted in Japan, indicating limited reliance on foreign infrastructure.

In contrast, Mexico exhibited a far more distributed hosting strategy, with a significant portion of its government services relying on foreign infrastructure. The majority of traffic was routed through the United States, underscoring Mexico's dependence on U.S.-based cloud providers and network intermediaries. Additionally, some government services were found to be hosted in Brazil, Israel, and even as far as Australia, further emphasizing the extensive cross-border dependencies in Mexico's hosting approach.

These results align with the trends identified by Kumar et al. [1], where countries with strong digital sovereignty policies, such as China, tend to prioritize domestic hosting, while others, like Mexico, rely more heavily on global cloud infrastructure.

## 3.2 Cross-Border Dependencies

Our traceroute-based path tracing allowed us to identify several instances of cross-border dependencies:

- **China's Limited Cross-Border Routing:** While China primarily hosts its government services within domestic infrastructure, we observed cases where traffic was routed through Japan. This suggests that despite China's stringent digital sovereignty policies, certain government-related services still depend on international connectivity. - **Mexico's Extensive Foreign Reliance:** The majority of Mexico's government services were hosted or routed through the United States, indicating a strong dependency on U.S. infrastructure. This reliance introduces potential risks related to data jurisdiction and service continuity in the event of geopolitical or policy changes. - **Additional Hosting in Brazil, Israel, and Australia:** A smaller fraction of Mexico's government content was hosted in Brazil, Israel, and Australia. These cross-border dependencies highlight the diverse hosting choices made by the Mexican government, which may be influenced by factors such as cost, service availability, or international partnerships.

These findings demonstrate the varying levels of control governments exercise over their internet infrastructure, with China exhibiting high domestic hosting concentration and Mexico relying on a globally distributed hosting strategy.

### 3.3 Figures

To further illustrate these hosting patterns and cross-border dependencies, we present the following figures.

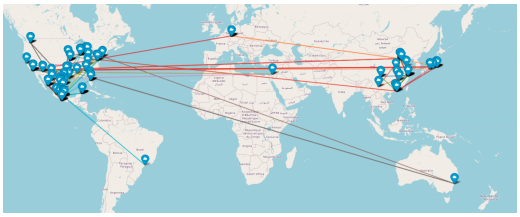

**Figure 1: All traceroute paths collected in this study, showing government network traffic routes and cross-border dependencies.**

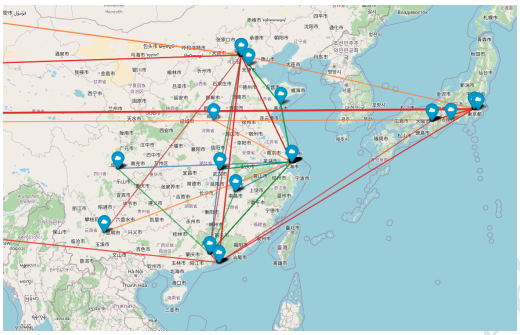

**Figure 2: Close-up of China's government hosting and network paths.**

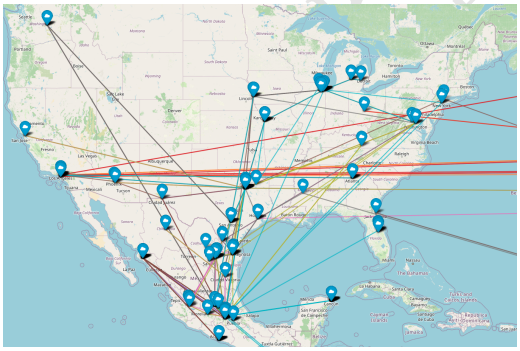

**Figure 3: Close-up of Mexico's government hosting and network paths.**

These figures provide a clear depiction of government internet traffic flows and the extent of reliance on foreign hosting providers. The visualizations highlight the stark contrast between China's primarily domestic hosting model and Mexico's extensive foreign dependencies.

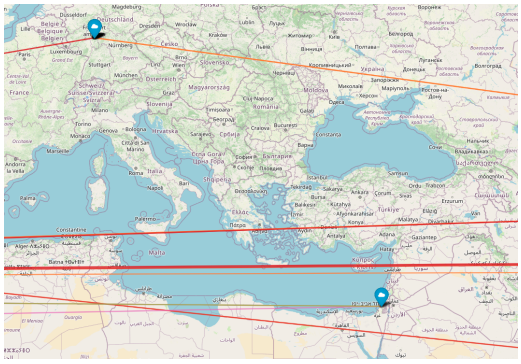

**Figure 4: Cross-border dependencies: Network paths passing through Israel and Germany.**

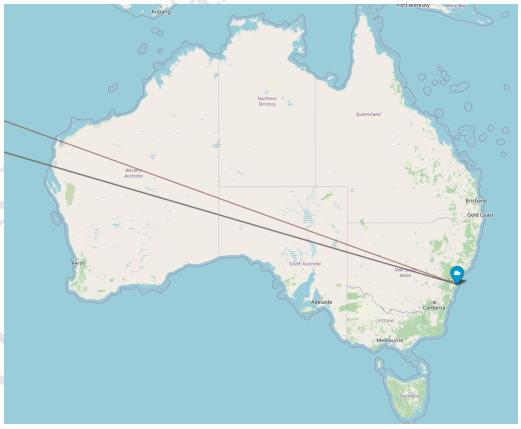

**Figure 5: Cross-border dependency with Australia.**

## 4 Conclusion

This study analyzed the hosting strategies of China and Mexico, focusing on their government web infrastructure and cross-border dependencies. Our findings confirm that China strongly prioritizes domestic hosting, with only minimal reliance on international infrastructure. Conversely, Mexico exhibits a highly distributed hosting model, with significant reliance on the United States and additional dependencies on Brazil, Israel, and Australia.

Our results highlight key policy considerations regarding digital sovereignty, security, and resilience. Countries like China maintain control over their digital infrastructure to minimize foreign dependencies, whereas countries like Mexico rely on external providers, introducing potential risks related to data jurisdiction and service continuity.

**Reflections on measurement challenges:** This study encountered a small number of challenges inherent to internet-scale measurements. Traceroute results varied due to network dynamics, and some IP addresses were unavailable. Geolocation data was often incomplete or inaccurate, requiring supplemental lookups that were not always successful. Routing asymmetry further complicated analysis, as forward and return paths did not always align. These challenges highlight the complexities of measuring global internet

infrastructure and underscore the need for robust methodologies in assessing government hosting strategies.

Future research could extend this analysis by incorporating more countries and tracking hosting trends over time. Our study provides a foundation for understanding government hosting models and their implications for digital sovereignty, cybersecurity, and cross-border data governance.

## 5 Acknowledgements

I would like to express my sincere gratitude to Professor Fabian Bustamante and Teaching Assistant Kedar Thiagarajan for their instruction and support throughout this project. This research was conducted as part of *COMP SCI 445: Internet-scale Experimentation* at Northwestern University, and I appreciate the opportunity to explore large-scale network measurement techniques within this academic framework.

## References

[1] Kumar et al. 2024. Of Choices and Control - A Comparative Analysis of Government Hosting. *Proceedings of the 2024 ACM Internet Measurement Conference (IMC '24)* (Nov. 2024).