En Unión y Libertad: Subnational Strategies for Hosting Government Services

Esteban Carisimo[1], Mariano G. Beiró[2], Lukas De Angelis Riva[3], Mauricio Buzzone[3], and Fabián E. Bustamante[1]

[1] Northwestern University {esteban.carisimo, fabianb}@northwestern.edu
[2] Universidad de San Andrés & CONICET mbeiro@udesa.edu.ar
[3] Universidad de Buenos Aires {ldeangelis, mbuzzone}@fi.uba.ar

Abstract. We present the first empirical study of subnational hosting strategies, using Argentina's 24 provinces as a case. Starting from official landing pages, we analyze ≈1.2k domains (collected Oct 2023 – Apr 2024), classifying serving networks by operational control (sovereign, domestic third-party, global) and examining authoritative DNS and HTTPS deployment. We relate these choices to 31 demographic, economic, technological, and political covariates – associations only, not causal claims. We find substantial heterogeneity: some provinces operate sovereign infrastructure; others rely on domestic incumbents or outsource to global providers. Federal capacity is rarely used, with provinces favoring bespoke or repurposed networks (including utility backbones). Legacy telecom footprints remain strong predictors of hosting choice even within a shared national umbrella. We also observe frequent splits between hosting and nameservers and uneven HTTPS hygiene. Taken together, the study offers a reusable measurement template and benchmarks that make sovereigntyperformance trade-offs measurable below the nation level.

En unión y libertad

Argentina's National Motto

1 Introduction

Accessing government services, from paying taxes and renewing licenses to scheduling health appointments, depends on reliable, Internet-facing infrastructure [79,35]. Citizens increasingly expect commercial-grade performance and availability [20], yet public agencies face additional constraints: legal data custody [28], long-term archival [82], and institutional accountability through proactive disclosure [29]. Choosing a hosting strategy that balances performance with control forces governments to navigate trade-offs between engineering realities (e.g., scalability, latency, resilience) and policy obligations (e.g., data protection, durable archiving, open publication) [79,35].

Many agencies rely on global clouds and CDNs for scalability, built-in DDoS mitigation, and broad footprints, at the cost of visibility and control over server placement, data residency, and network interconnections. Running infrastructure in-house preserves autonomy but requires specialized staff, capital investment, and long-term maintenance. Prior work has examined how national governments approach these trade-offs and documented sharp differences in outsourcing rates and provider choice [62].

This paper shifts the lens to subnational governments (provinces and states). Operating under the same federal legal and infrastructural umbrella, these jurisdictions manage their own digital portals — property registries, benefits, public health alerts — and independently choose where and how to host them. Local infrastructure, political priorities, and budgets can diverge sharply from national norms.

Argentina is a federal country with 23 provinces and one autonomous city (CABA), each with constitutional authority over its digital infrastructure. Spanning 2.8 million km^2 — from dense urban corridors to sparsely connected Andean and Patagonian regions — it offers a natural setting to observe decentralized choices under a shared national framework. Despite a national fiber backbone (REFEFO) [12] and a federally operated data center, provinces are free to choose their own hosting arrangements.

Study design. Building on Kumar et al. [62], we adapt their nation-level approach to the subnational tier. Starting from each jurisdictions official landing page, we crawl to depth seven and analyze >1,200 domains (Oct 2023Apr 2024), classifying serving ASNs by operational control*sovereign* (Sovereign), *local third-party* (Local), or *global third-party* (Global). We also examine authoritative DNS and HTTPS deployment and relate these choices to 31 demographic, economic, technological, and political covariates (associations only). Government services comprise domains under official provincial portals (e.g., *.gob.ar, *.gov.ar) and first-party subdomains linked from them; we exclude social media, SaaS dashboards, and ad/analytics endpoints. Crawls ran from three Argentine vantage networks, and provider classifications were stable across vantage (§3).

Why Argentina – and why it generalizes. Beyond its scale and diversity, Argentina mirrors a broader Latin American and Caribbean (LAC) pattern in which subnational governments deliver a growing share of public services under heterogeneous capacities and budgets. In LAC, the subnational share of consolidated public spending rose from ≈13% to ≈25% between 1985–2015, and in federations such as Argentina and Brazil it exceeds 40% — placing provinces at the critical edge of digital service delivery [87]. Argentina combines (i) *autonomy* — provinces and agencies retain genuine purchasing and operational control; (ii) *contrast* — a federal backbone (ARSAT's REFEFO) [12] coexists with local and global third-party providers; and (iii) *observability* — government transparency initiatives publicly listing all digital resources make infrastructure choices measurable at scale [29]. This makes it a compelling laboratory for infrastructure choices under shared national policy but uneven local capacity, with lessons for other federal and quasi-federal contexts.

By surfacing these patterns, we show how subnational governments, often overlooked in Internet measurement, shape the geography of public-sector infrastructure. Our results provide both a methodological template and policy-relevant benchmarks for assessing digital sovereignty below the national level. To anchor our claims, we add a lightweight latency sanity check across hosting classes and a cache-locality audit that verifies whether global providers serve from in-country PoPs. *Our goal is not only to characterize one country, but to offer a reusable template that researchers and agencies can apply to states, regions, or municipalities elsewhere.*

This paper makes three contributions. First, it offers the first empirical study of hosting strategies by subnational governments, revealing substantial variation in choices made under shared national policies, infrastructure, and legal frameworks. Second, it provides an annotated map of provincial dependencies by analyzing over 1,200 government domains across 24 jurisdictions, classifying serving providers (Sovereign, Local, Global), authoritative DNS operators, and TLS configurations, and relating these to 31 socioeconomic and political variables. Third, it uncovers evidence of legacy and localization effects, including the imprint of pre-Internet telecom monopolies, the repurposing of utility/education/telco ASNs for government services, and the limited use of federally funded capacity (e.g., ARSAT), raising questions about coordination, performance, and control.

Our key findings are as follows. Heterogeneous strategies: 11 jurisdictions rely primarily on Local, 7 on Global, and the remainder on Sovereign for URLs/bytes,[4] with domains showing similar splitsthere is no single national template. Underuse of federal capacity: only 8 of 24 provinces use ARSAT at all, never exceeding 2% of bytes. Province-operated footprints: 15 provinces serve content from province-run or affiliated networks, with byte shares ranging from ≈1.594%, indicating wide variation in operational maturity. Incumbent imprint: hosting patterns align with historical Telecom and Telefnica regions, where incumbents still retain substantial shares. DNS and TLS gaps: DNS is mostly outsourced (Local 41%, Global 35%); 48% of domains use different providers for web hosting and authoritative DNS; and in 14/24 jurisdictions domains employ neither anycast nor multi-AS redundancy. HTTPS is enforced on 64% of domains, yet 19% remain HTTP-only, with failures often due to certificate-chain or hostname errors.

2 Subnational Digital Infrastructures: Policy, Constraints, and the Argentine Case

We outline national digital programs, noting converged front-end standards but decentralized back-end hosting. We highlight the role of subnational governments in transaction-heavy services, summarize four constraints on their infrastructure choices, and position Argentina as a regional example. We close by showing how these approaches extend to other federations.

[4] Bytes denotes retrieved content size from our crawl, not network traffic volume.

2.1 National Digital Programs

Many national digital programs emphasize unifying the front-end of government websites — through design systems, accessibility rules, and UX guidelines [93,3,4,5,6,1,9,7,36,2,8,53]. By contrast, decisions about hosting, traffic delivery, and data residency are often decentralized or weakly enforced. Some countries have experimented with stronger, sovereignty-motivated approaches (e.g., government-operated platforms or localization mandates), but these vary in scope and longevity.[5] In short, while national frameworks set common front-end baselines, the infrastructure layer frequently remains a patchwork of agency, vendor, and subnational choices. We describe these programs to contextualize observable hosting outcomes; we do not evaluate policy efficacy or compliance.

2.2 Why Subnational Governments Matter

Subnational governments in LAC deliver most education, health, security, and local administrative services. Their share of consolidated public spending grew from 13% to 25% between 1985 and 2015 [87]. Argentina sits in the high-decentralization tier: provinces and municipalities account for roughly 40% of public expenditure [87,85]. Provinces therefore run transaction-heavy, always-on digital servicestax portals, licensing, appointmentswhere performance directly affects citizen engagement. Here, government services denotes public-sector domains and first-party subdomains linked from official portals; regulated private providers are excluded.

At the same time, subnationals often face binding constraints: limited and rotating IT staff, fiscal dependence on transfers, and weak cross-jurisdiction ICT governance [87]. These pressures create incentives to outsource to large platforms and CDNs, trading capital costs for operating expenses — but sometimes sacrificing control, auditability, or data residency. We therefore treat hosting as an *organizational choice under constraints*. Our focus complements national-level studies by showing how local capacity, incumbent footprints, federal offerings, and market structure jointly shape outcomes.

2.3 Challenges for Subnational Digital Services

Policy evidence and practice point to four recurring constraints that shape hosting choices for provinces and municipalities:

Institutional capacity. Teams are small and rotate frequently, with documented hiring freezes (2018–2019), difficulties attracting and retaining skilled

[5] Illustrative examples include: (i) the UK's GOV.UK PaaS, which centralized hosting for some agencies (decommissioned in Dec. 2023); (ii) Russia's personal-data localization rules requiring storage of residents' data within national borders; (iii) India's *MeghRaj* (GI Cloud) framework with in-country data requirements; and (iv) Europe's *Gaia-X* initiative for federated, interoperable infrastructure to bolster data sovereignty. These are illustrative, not endorsements; scopes and timelines differ.

staff, and reliance on temporary contracts, eroding portfolio governance and continuity [77,31,87].

Fiscal constraints. Vertical fiscal imbalance and reliance on intergovernmental transfers limit the predictability of capital expenditures. Cloud models shift costs from capital to operational expenditures ("pay-as-you-go") and reduce anticipatory budgeting; an IDB-cited study finds an in-country data center model could be $\approx 54\%$ costlier than leveraging global cloud providers, and that IXPs and cache placement can reduce bandwidth costs and latency [85,84,35]. We therefore treat budget exposure as context, not as measured spend, and later weight outcomes using domain/URL/bytes aggregates as proxies for footprint and content heft.

Fragmented procurement & governance. Policy recommendations encourage a transition from agency-by-agency purchasing toward a "government-as-one-client" model. Multiple, disconnected platforms hinder reuse, and e-procurement maturity remains uneven across LAC [77,80,14].

Risk management & resilience gaps. Siting, power redundancy, and regularly tested disaster-recovery and business-continuity procedures are often under-specified, despite guidance to provision reserve capacity, diversify and duplicate critical systems, and run exercises [78,91,95]. Operationally, we map these risks to two measurable dimensionsauthoritative DNS anycast/AS-level diversity and failover/DR posturewhich we quantify in §5.

These measured constraints then interact with existing telecom footprints and federal programs (e.g., ARSATs REFEFO backbone [12]), helping explain the heterogeneous hosting choices we document across provinces.

Recent events underscore the stakes: in May 2024, extreme flooding in Rio Grande do Sul (Brazil) led the state IT company (PROCERGS) to shut down its data center, taking multiple state services offline until gradual restoration [16,10,39]. Beyond that case, comparative guidance highlights that exposure to hydrological risk, placement of backup power, reserve capacity, and the frequency of continuity testing are first-order design choices for subnational digital infrastructure [78,91].

2.4 Argentina in Regional Perspective

Argentina intensified its national digital policy in 2016 with the *Plan de Modernización del Estado*, which authorized the National Office of Information Technologies (ONTI) to establish technical standards and offer guidance throughout the public sector [83]. ONTI subsequently published and updated accessibility and technology standards — e.g., the 2019 *Disposición ONTI 6/2019* on Web Content Accessibility Guidelines (WCAG) 2.0 and the *Decálogo Tecnológico* — that continue to shape how public websites are developed and maintained [73,72]. However, these policies largely leave infrastructure choices (where and how to host, use of clouds and CDNs, resilience practices) to each agency or jurisdiction. As a federation of 23 provinces and one autonomous city — with independent administrative and budget processes — Argentina mirrors broader LAC reforms that emphasize clearer assignment of responsibilities and intergovernmental transfers while calling for stronger subnational management capacity and digitalization [87]. Federal infrastructure such as ARSAT's *Red Federal de Fibra Óptica* (REFEFO) provides important backbone capacity, but does not itself determine provincial hosting strategies [12]. These policies set design baselines while leaving hosting to agencies and jurisdictions; our analysis observes the resulting deployment choices, not compliance.

2.5 Argentina as a Case Study

Why Argentina – and why it generalizes. Three features make Argentina a natural case study: (1) federative structure, (2) articulation between federal and subnational governments in digital policies, and (3) availability of federally funded digital infrastructures. **Federative structure**, it is a federation of 23 provinces and one autonomous city with independent administrative and budget processes; provinces and municipalities execute on the order of 40% of consolidated public spending, in line with a broader regional shift in LAC where the subnational share rose from 13% to 25% between 1985 and 2015 [85,87]. This places provinces at the front line of transaction-heavy digital services. **Policy articulation: centralized design, decentralized infrastructure** the policy design separates front-end standardization from back-end infrastructure discretion, creating cross-provincial variation in hosting models, cloud and CDN reliance, and resilience practices that can be observed and measured [77]. **Federally funded digital backbones**, federal infrastructure such as ARSAT's *Red Federal de Fibra Óptica* (REFEFO) provides backbone capacity without prescribing provincial hosting strategies, interacting with heterogeneous provincial telecom footprints and local IXP and cache opportunities [12,35]. From this landscape, we derive testable hypotheses: (H1) Incumbent fixed-line footprints predict provider choice; (H2) Federal capacity (e.g., ARSAT) sees limited uptake relative to local/global options; (H3) Capacity constraints correlate with platform outsourcing; (H4) Cache locality moderates the sovereignty – performance trade-off. Sections 5 – 6 test these associations.

How the case generalizes. Argentina presents a set of features shared by both Latin American countries and nations outside the region. The federal division of powers, centralized design guidance with decentralized control of hosting infrastructure, and the coexistence of a public backbone alongside domestic and global cloud and CDNs create a set of characteristics that can be extended to other countries. These features closely resemble those found in other Latin American and Caribbean federations with extensive territories (e.g., Brazil, Mexico) and in unitary states (e.g., Colombia, Chile). However, this is not exclusive to Latin America, as similar features are shared by OECD countries (e.g., U.S. states, Canadian provinces, and Spain's autonomous communities). We then examine how capacity constraints, capital and operational expenditure trade-offs, fragmented procurement, and resilience and continuity planning — shape Argentina's subnational digital hosting, and derive actionable design implications for any jurisdiction where subnational entities choose hosting under national policy frameworks [87,77,35].

3 Methodology and Dataset

This section describes our methodological framework (§ 3.1), the process for compiling and scraping government domains (§ 3.2), and the non-network variables used to explain hosting decisions (§ 3.3).

3.1 Methodology

Our methodology builds on the work of Kumar et al. [62], who analyzed national government hosting across 61 countries. We adapt and extend their approach to focus on subnational entities.

Scraping approach. We start from each jurisdiction's official landing page and crawl up to depth 7, following the depth criteria applied in Kumar et al. [62] and Singanamalla et al. [89]. Because public sites often embed third-party content (e.g., YouTube, X), we restrict the dataset to *first-party* resources via a multi-step filter: (i) allowlist government domains at the eTLD+1, (ii) retain only internal URLs, and (iii) drop known third-party embeds/trackers. Table 1 details these rules. We canonicalize URLs (lowercasing, query-param normalization), deduplicate exact resources, honor `robots.txt`, and justify depth-7 as a coverage/tractability trade-off with a saturation check.

Table 1. Steps used to identify government-owned domains.

Approach	Description
Government TLDs	Matches `.gob.ar` or `.gov.ar` domains
Domain Matching	Checks if hostname matches known gov domains
TLS SAN Matching	Identifies gov domains via certificate SANs

Government domain identification. We first label as government resources those with domains under known government top-level domains (TLDs), as proposed by Singanamalla et al. [89]. In Argentina, these TLDs include both `.gob.ar` and `.gov.ar`[6].

For domains outside the `.gob.ar`/`.gov.ar` TLDs (e.g., state-owned firms like `ypf.com`), we apply manual and certificate-based heuristics to confirm government ownership before resolving to IPs and mapping to ASNs.

[6] While `.gob` is the natural extension for "government" (gobierno) in Spanish (Argentina's official language), many Spanish-speaking countries also use `.gov` as a legacy convention influenced by the United States.

Vantage points. We fetch pages from in-country vantage points to reflect domestic path selection, cache affinity, and CDN mapping. We deploy three vantage points in the Buenos Aires metropolitan area: one on a large university (albeit running on three VMs) and two residential nodes in major access ISPs – Telecom Argentina (AS7303) and Telecentro (AS27747) – which serve approximately 27% and 10% of the user base, respectively [63]. This placement captures the CDN choices applied to Argentina-based users, including in multi-CDN settings. Our scraping did not include vantage points in every province, as we assume that hosting strategies will not change on a per-province basis.

Attribution and metrics. We record the server IP address for every fetched resource and attach per-resource metrics. IP addresses are mapped to origin Autonomous Systems (ASes) using CAIDA's prefix-to-AS dataset [22]. For each URL, we log bytes served — the on-disk size retrieved by the crawler for first-party HTML and same-origin static assets during rendering — not traffic volume; we report URLs, domains (eTLD+1), and bytes per provider to capture structure, scope, and content heft. To contextualize provider reliance, we report three aggregates per site, each highlighting a different aspect of government hosting: (i) distinct serving URLs — capturing application surface and fragmentation; (ii) distinct domains (eTLD+1) — capturing administrative scope and delegation; and (iii) total bytes served — capturing content and storage footprint. All aggregates are broken down by AS/provider. Together, these measures reveal the structural dependence of government sites on specific networks and platforms.

Provider taxonomy (control, not location). Our approach groups the serving ASN for each domain into three categories based on operational control: (1) Government/SOE (Sovereign), (2) Third-Party Local (Local), and (3) Third-Party Global (Global). Sovereign covers ASNs dedicated to government use or run by state-owned enterprises (e.g., ARSAT). Local includes commercial ASNs registered and primarily operating in Argentina. Global refers to multinational networks with an international footprint. These labels reflect ownership and controlnot physical server location. For instance, ARSAT remains Sovereign even with broad peering, while global CDNs with local PoPs remain Global. Classification is supported by registry and documentary sources.

Our provider taxonomy combines documentary verification with registry checks. First, government and state-owned networks were manually classified by searching official sources (laws/decrees, corporate registries, and agency pages) that establish public ownership or control. Second, we label as Local firms incorporated in Argentina, i.e., companies registered under Argentine jurisdiction. Third, we label as Global global platforms that operate across multiple continents (e.g., multi-region cloud/CDN providers). Within our measurement corpus, we did not identify foreign providers serving Argentina's government content other than these Global.

3.2 Dataset

We compiled an official seed list of landing pages for Argentina's federal government and its 24 subnational jurisdictions. For the federal branch, our starting point was the executive-branch CSV that lists ministries and secretariats (e.g., Ministry of Economy), federal agencies (e.g., ENACOM), and state-owned enterprises (e.g., Nucleoelctrica Argentina S.A.)[7]. We complemented this automatically harvested list with a manual review of the legislative and judicial branches (Senate, Chamber of Deputies, Supreme Court, Public Prosecutor), using official organograms to locate each branch's top-level portal.

At the provincial level, we followed the same approach: most provinces publish organograms or directories that link to ministries, secretariats, and agencies. Where a consolidated list was unavailable, we extracted candidate URLs from official organizational charts and then manually verified coverage by spot-checking each province's directory pages. State ownership of companies is generally traceable from these official disclosures — for example, the federal CSV identifies YPF as dependent on the Ministry of Economy and the Secretary of Energy. Provinces likewise publish their lists of publicly owned companies: for example, the Province of Buenos Aires documents its ownership of power-generation firms such as Centrales de la Costa Atlántica S.A.[8]

Terminology. A *landing page* is a top-level portal listed in an official organogram. A *domain* is counted at the eTLD+1 (e.g., `salta.gob.ar`). *URLs* are distinct first-party pages discovered during the crawl. *Bytes* denote the on-disk size retrieved for first-party HTML and same-origin static assets during rendering (not traffic volume).

We crawled **202** federal and **1,446** subnational landing pages between Oct. 2023 and Apr. 2024, yielding **57,730** first-party URLs across **1,244** eTLD+1 domains (8.8 GB). Figure 1 shows landing-page coverage by province.

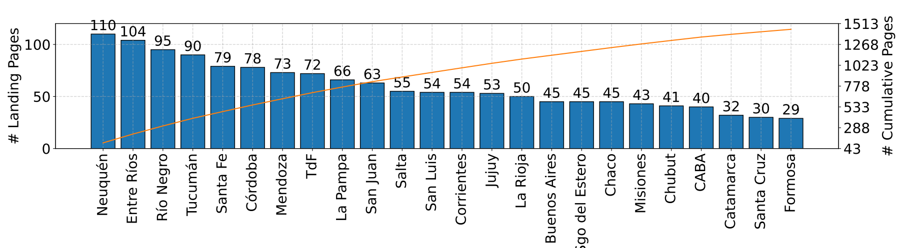

Fig. 1. Landing pages collected per province (bars) and cumulative share (line).

Table 2 summarizes where content lives by suffix class. We apply the same control-based taxonomy to *authoritative DNS operators* (sovereign/local/global)

---

[7] Latest disclosure of Argentina's executive branch composition: https://mapadelestado.jefatura.gob.ar/back/api/datos.php?db=m&id=9&fi=csv. Accessed Feb. 15, 2025.

[8] Latest disclosure of the Province of Buenos Aires state-owned enterprises: https://www.gba.gob.ar/infraestructura/empresas. Accessed Dec. 2, 2025.

Table 2. Share of URLs, domains, and bytes by suffix class.
Categories mix second-level domains under `.ar` (e.g., `gob.ar`) and global TLDs (e.g., `.com`). Percentages may not sum to 100 due to rounding.

Metric	gob.ar	gov.ar	com.ar	.com	edu.ar	org.ar	.ar	.org	.net	.tv	.app
URLs	46.5	32.1	11.1	3.1	2.6	2.3	1.1	0.6	0.2	0.0	0.0
Domains	47.1	34.1	11.9	0.0	2.3	2.0	1.0	0.8	0.3	0.1	0.1
Bytes	48.2	27.2	12.7	2.9	2.5	0.6	1.6	3.9	0.4	0.0	0.0

and evaluate HTTPS deployment and certificate chain/hostname quality across all domains.

At the time of writing, the Milei administration is restructuring the executive branch (e.g., ministry mergers; AFIP $\rightarrow$ ARCA). To gauge URL stability, we compare two snapshots of *federal top-level portals and agencies* (Aug. 23, 2023, pre-election; Sep. 3, 2025, mid-term): the Jaccard distance is 0.35 (similarity 0.65), with **158/243** URLs unchanged and **85/243** changed (new entities, mergers, decommissions, or rewrites). Thus, while the federal web evolved, a *majority* of portals remained stable across the period.

We release the seed list, suffix-class rules, and per-province aggregates upon request to support replication.

3.3 Contextual Variables

To explain variation in hosting choices, we compile 31 contextual variables across four categories:

We organize these non-network variables into four categories: demographic, economic, technological, and governmental. The dataset includes 31 variables for the federal government and 24 provincial jurisdictions, with a complete version available in Appendix B. All data were collected from authoritative sources, including Argentina's communications regulator (ENACOM), the National Institute of Statistics and Censuses (INDEC), and several United Nations agencies.

The following paragraph provides a brief overview of the variables within each pillar.

Demographic variables. This pillar includes 11 indicators: inequality (Gini coefficient) [76], Human Development Index (HDI) [86], average years of schooling [52], share of the population with tertiary education [44], total population [45], urbanization rate [46], migration rate [33], life expectancy [51], illiteracy rate [32], poverty rate [49], and financial inclusion [13]. Most of these variables come from INDEC's national census and routine surveys, supplemented by data from United Nations sources and provincial agencies. We use these variables to test whether each jurisdiction's social profile influences the government's choice of global, local, or in-house hosting strategies.

Economic variables. This pillar tracks 8 indicators: GDP per capita [19], unemployment rate [50], public expenditure as a share of GDP [81], public debt as a share of GDP [60], fiscal balance (tax result) as a share of GDP [43], inflation rate [48], and public-sector employment [56]. This compilation was obtained from several authoritative sources, including INDEC, academic research centers, the Congressional Budget Office, and national ministries. We use these variables to test whether a province's economic performance affects its government's content-delivery strategies.

Technological variables. This pillar covers 9 indicators: overall Internet-penetration rate [34]; fixed-line Internet access rate [47]; mobile Internet access rate [47]; household Internet-penetration rate [47]; enterprise Internet-penetration rate [47]; average connection speed [34]; number of peering facilities [23]; number of registered ASes[9]; and the share of government domains among the country's 100k most-visited websites [37]. We use these metrics to test whether the strength of a jurisdiction's Internet infrastructure influences its preference for global, local, or government-run hosting strategies.

Governmental variables. This pillar tracks three indicators: (i) the democracy score, derived from the Varieties of Democracy (V-Dem) methodological guidelines [92]; (ii) congress diversity, measured as the entropy of party representation in each provincial legislature; and (iii) political diversity, defined as the number of active political parties [24]. We use these variables to test whether a jurisdiction's political profile influences its preference for global, local, or in-house hosting strategies.

4 Serving Infrastructures

In this section, we examine the hosting choices of Argentina's subnational governments. We focus on their preference for self-hosting versus third-party providers (§4.1); their use of federal- or province-owned infrastructure (§4.2); the prevalence of legacy fixed-line incumbents in content delivery (§4.3); and their reliance on local hosting providers (§4.4).

4.1 Government-Run vs. Third-Party Hosting

We quantify, for each jurisdiction, the fraction of *URLs*, *domains*, and *bytes* served by Sovereign (sovereign), Local (third-party domestic), and Global (third-party foreign) providers (Fig. 2). Using *bytes* as the primary metric and a *plurality* decision rule, 11 provinces predominantly rely on Local, 7 on Global, and the remainder on Sovereign; *domains* show a similar split.[10] There is no single national template; for illustration, CABA serves 38.1% of domains, 33.5% of URLs, and 48.9% of bytes via Sovereign.

Distributions vary widely across jurisdictions. For *URLs*, Local spans 0–84% (mean 42.3%), Global 2.2–84.5% (31.2%), and Sovereign 0–76.7% (26.5%). For *domains*, Local spans 0–84% (44.1%), Global 5.9–66.7% (29.2%), and Sovereign

[9] Computed with a script that combines LACNIC delegation files, RDAP records, and geolocation APIs from OpenStreetMaps, Geoapify, and OpenCageData.

[10] Replacing plurality with a $> 50\%$ threshold leaves conclusions qualitatively unchanged.

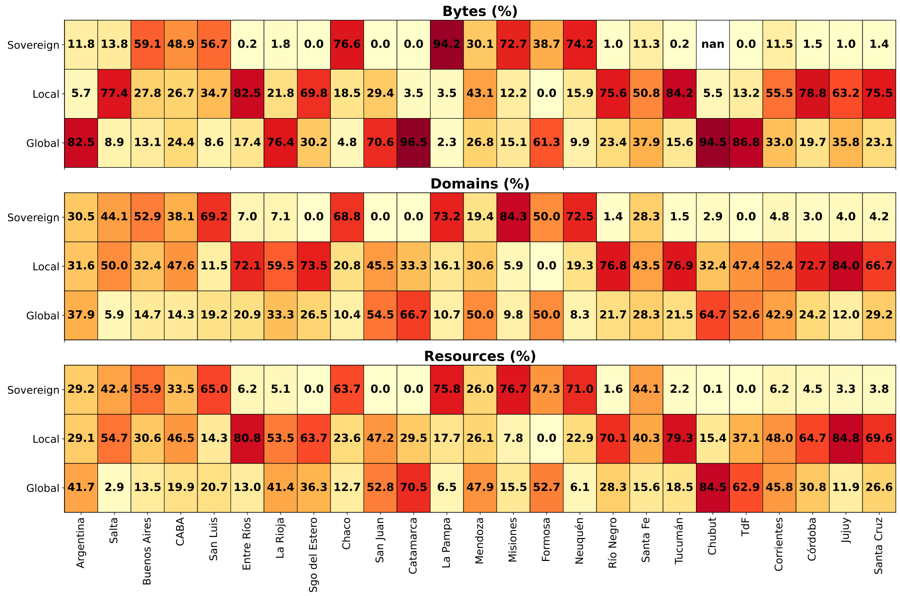

Fig. 2. Hosting choices by jurisdiction: share of URLs, domains, and bytes served via Sovereign, Local, and Global.

0–84.3% (26.7%). For *bytes*, Local spans 0–84.2% (39.0%), Global 2.3–96.5% (36.8%), and Sovereign 0–94.2% (24.5%).

Sanity checks (summary). (i) *Latency*: from three Buenos Aires area vantage points, median TTFB/ping differs across classes as expected (Global ≈ Local when served from in-country PoPs; Sovereign varies with provincial backhaul). (ii) *Cache locality*: for each domain in Global, we resolve/trace to confirm in-country CDN PoPs when present.

4.2 Federal- and Province-Operated Infrastructures

ARSATs REFEFO backbone and the Benavídez data center are widely reachable, yet provincial uptake is limited. At the federal tier, ARSAT is the #2 network by domains, URLs, and bytes; at the provincial tier, only 8/24 jurisdictions use ARSAT at all — typically for a single domain — and it never exceeds 2% of bytes (max: Jujuy). This pattern suggests organizational or procurement factors dominate raw reachability.

Argentinas federal government owns ARSAT, founded to operate geostationary satellites [55] and later expanded to build the Federal Optical Fiber Network (REFEFO), a country-wide backbone [12], and to run a national data center. The Benavídez facility (inaugurated in 2012 [11]) connects directly to REFEFO, providing paths to all provinces. Given this federally funded capacity, we examine whether governments leverage ARSATs network and data center to host federal- and province-level websites.

Table 3. Top networks serving the federal government by URLs, domains, and bytes.

	URLs		Domains		Bytes	
ASN–ASName	#	Rank	#	Rank	MB	Rank
13335–Cloudflare	1,561	1	12	1	1,618.0	1
52361–ARSAT	**410**	**2**	**10**	**2**	**83.5**	**2**
8075–Microsoft	231	8	9	3	20.8	7
16814–NSS	369	3	7	4	16.2	12
27823–Dattatec	321	4	7	4	16.7	11
266700–HCDN	300	5	4	8	50.2	3
6121–Ministry of Justice	136	12	2	10	34.2	4
7049–Silica Networks	168	9	2	10	21.3	5

At the federal level (Table 3), ARSAT ranks second by all three metrics (10 domains; 410 URLs; ∼83.5 MB). At the provincial level, the picture changes markedly.

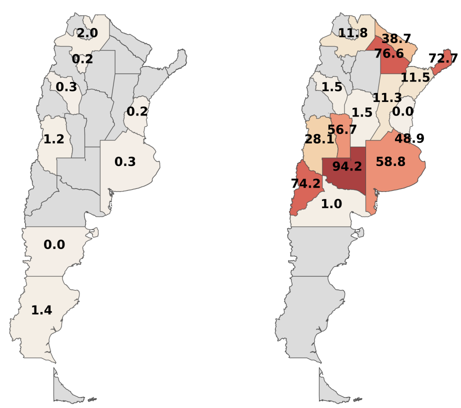

(a) federal-owned infrastructure (b) province-owned infrastructure

Fig. 3. Share of provincial bytes delivered via (a) federal-owned (ARSAT) and (b) province-owned networks.

As shown in Fig. 3a, only a third of provinces (8/24) use ARSAT to host content, and it is never used to serve more than 2% of bytes (Jujuy). Each of these provinces relies on ARSAT for a single domain; by number of URLs, the largest is Mendoza (115 URLs). Despite uneven coverage, a reduced measurement campaign using 60 in-country RIPE Atlas vantage points found the 90th-percentile probe-to-Benavídez RTT ≤ 40 ms; suggesting that latency is unlikely to be the primary driver of hosting decisions.

Turning to province-operated infrastructures, Fig. 3b shows that 15 provinces serve content from their own or affiliated networks. Across 23 such networks, byte shares range widely — from La Pampa at 94.2%, Neuquén at 74.2%, and San Luis at 56.7% to La Rioja at 1.5%.

Table 4. Fraction of bytes served by each province-owned network.

ASN–ASName	Province	%	Role
27967–Buenos Aires Gov	Buenos Aires	58.8	Gov
269906–Chaco Digital	Chaco	0.0	Telco
52373–ECOM Chaco	Chaco	76.6	Telco
3449–University of Buenos Aires	CABA	0.5	Education
269766–Council of Magistracy	CABA	1.6	Gov
267816–Public Prosecutor's Office	CABA	3.5	Gov
52318–ASIGCBA	CABA	43.3	Gov
270044–Corrientes Telecom	Corrientes	11.5	Telco
262150–EPEC	Córdoba	1.5	Utility
272834–ENER	Entre Ríos	0.0	Utility
263791–REFSA	Formosa	38.7	Utility
52308–Waters of the Colorado	La Pampa	94.2	Utility
28048–Internet for All	La Rioja	1.5	Telco
264808–Ministry of Finance	Mendoza	12.3	Gov
27894–Judiciary Branch	Mendoza	15.8	Gov
263774–Marandú	Misiones	72.7	Telco
262222–OPTIC	Neuquén	74.2	Gov
267820–NEUTICS	Neuquén	0.0	Gov
265687–ALTEC	Río Negro	1.0	Telco
264848–Gov. Financial Service	Salta	10.9	Gov
265651–Ministry of Finance	Salta	0.9	Gov
52440–University of La Punta	San Luis	56.7	Education
52293–Santa Fe Gov	Santa Fe	11.3	Gov

Zooming in (Table 4), six provincial governments operate eleven dedicated networks for content delivery (e.g., AS27967 in Buenos Aires; AS52318 in CABA), while others repurpose pre-existing province-owned infrastructures: four leverage utility networks (e.g., AS262150 in Córdoba), five rely on telecom networks (e.g., AS263774 in Misiones), and two use educational networks (e.g., AS52440 in San Luis). Reuse can deliver control at low capex but may inherit design constraints; we evaluate resilience proxies (authoritative-DNS anycast/AS diversity and failover posture) in §5.

4.3 Incumbents Fingerprints

In the early 1990s, Argentina privatized the state telco ENTEL [69,67] and split the country into two service regions — north (upper) and south (lower) — with the City of Buenos Aires divided between them [68]. The northern region was awarded to Telecom Argentina (then a Telecom Italia/France Télécom JV) [70]; the southern region to Telefónica de Argentina [71]. Although Internet service (commercial from 1994 [75]) was not regulated as a public service, fixed-line incumbents quickly leveraged their plant and customer bases to dominate last-mile access.

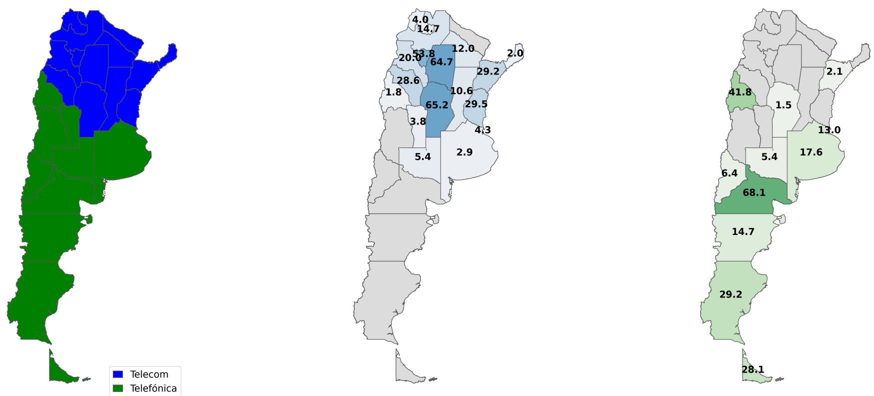

(a) Legacy split (b) Telecom Argentina (AS7303) (c) Telefónica Argentina (AS10834)

Fig. 4. Share of provincial *domains* hosted by each incumbent. Panel (a) shows the 1990 legacy split; (b)–(c) show present-day provincial shares.

Present-day government hosting patterns track the historical Telecom/Telefónica demarcation. Assigning each province to its legacy incumbents service area, the incumbent's share of *domains* hosted within its home zone mirrors the legacy split. Using *bytes* yields the same qualitative result, with province-level variation (e.g., Santa Cruz 29.2→7.3%, Córdoba 65.2→75.3%). We treat this as an *association*, not a causal claim.

Two mechanisms are consistent with this pattern: inherited access/backhaul footprints and institutional familiarity (procurement/sales channels). Server *location* may differ (e.g., CDN PoPs) and does not determine class; we account for cache locality and latency when interpreting performance.[11]

4.4 The Role of Local Providers

Beyond the two fixed-line incumbents, a handful of domestic networks account for most provincial outsourcing — e.g., Claro (AS11664; also AS19037/AMX), NSS/iPlan (AS16814), and Dattatec/DonWeb (AS27823). We use *bytes* as the primary metric, sort the heatmap by each ASNs *aggregate* byte share across

[11] We observed anecdotal evidence of locally hosted government servers (e.g., Río Gallegos for Santa Cruz). Geolocation for these hosts was inconsistent across databases, and we lacked in-province active measurements; we therefore do not rely on these anecdotes in our analysis.

provinces, and annotate the top five. Single-province providers are common (11 provinces), but only four exceed 5% of bytes; we list these explicitly in Table 5 and discuss operational risk in §5.

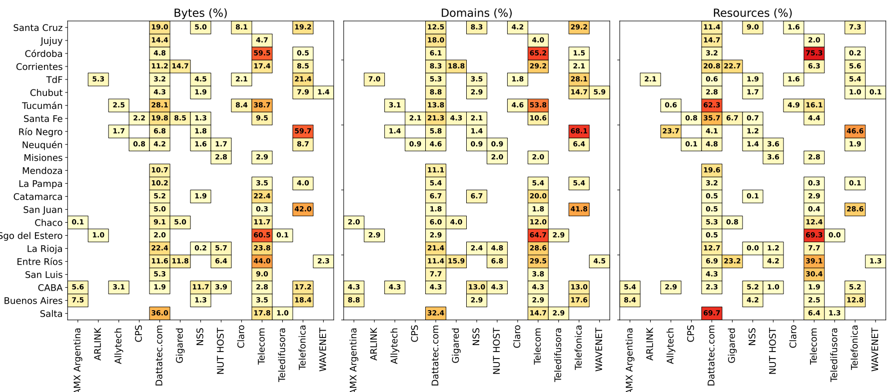

Fig. 5. Share of URLs, domains, and bytes served by each *local* network that serves more than one province. Rows (ASNs) are sorted by total byte share across provinces.

Figure 5 highlights local providers with multi-province reach. Besides Telecom Argentina (AS7303) and Telefnica Argentina (AS10834), major domestic networks include Claro (AS11664/AS19037) and NSS/iPlan (AS16814). Regional patterns emerge: Gigared (cable along the Paran River) serves Chaco, Corrientes, Santa Fe, and Entre Ros, while Dattatec/DonWeb (AS27823) appears in 22 provinces and reaches up to 69.7% of URLs and 36.0% of bytes in Salta. Results are similar when ranking by domains instead of bytes.

We also examine providers dedicated to a *single* province. This preference appears in 11 provinces; among 18 such networks, only four exceed 5% of bytes: AS52312 (TV MUSIC HOUSE, Jujuy), AS264738 (Sebastián Souto, Santa Cruz), AS27879 (IyT, Mendoza), and AS61449 (RESEARCH, Misiones). Several others are very small and, in some cases, are personal or boutique ASNs (e.g., 265781 — Pala Pablo Federico; 264689 — Luciano Gabriel Chersanaz). Appendix C provides per-AS details (Table 10).

Table 5. Single-province providers exceeding 5% of bytes.

ASN–ASName	Province	Bytes (%)
52312–TV MUSIC HOUSE	Jujuy	>5
264738–Sebastián Souto	Santa Cruz	>5
27879–IyT	Mendoza	>5
61449–RESEARCH	Misiones	>5

Takeaways. Hosting strategies vary: most use local third-party providers, some rely on global CDNs, and others have government-run footprints. Federally owned capacity (ARSAT, REFEFO) appears underutilized, with only 8 of 24 jurisdictions using ARSAT and never exceeding 2% of bytes, suggesting procurement or operational issues. Fifteen provinces serve content through province-operated networks, often using utility or telco ASNs, resulting in a large variation in bytes (1.5% to 94%) and affecting resilience and maintainability. Legacy incumbents — Telecom and Telefonica — influence provincial domain hosting, aligned with historical service regions. Future analysis should consider economic factors in hosting decisions, such as hosting costs and budgets, and foreign currency exposure in countries with volatile exchange rates.

5 Domain Name Infrastructures and Certificates

We examine two further components of the delivery stack: authoritative DNS (§5.1) and HTTPS (§5.2).

5.1 Name Servers Choices

Domain management introduces additional choicesthird-party DNS, anycast, multihoming, and load balancing. We focus on three questions: (1) how often governments outsource nameservers, (2) whether they deploy redundancy, and (3) how DNS choices align with web-hosting choices.

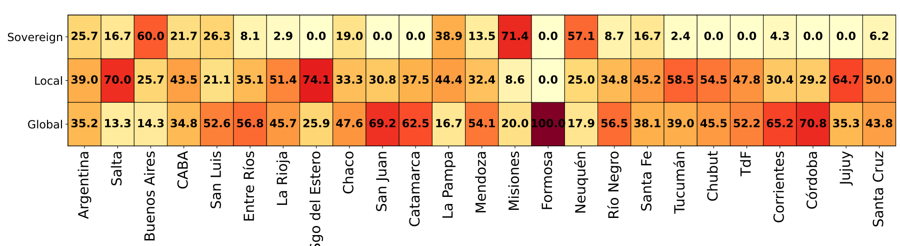

Fig. 6. Nameserver outsourcing strategies across Argentina's jurisdictions, classified by sovereign, local, and global providers.

Outsourcing Nameservers Management. We map which provinces manage or delegate their authoritative name servers. Figure 6 shows the share of domains using sovereign, local, or global providers. On average, 23.5% use sovereign platforms, 41.1% local providers, and 35.4% global operators. Self-hosting is uncommon: only Buenos Aires, Misiones, and Neuquén use sovereign NSs for most domains, and eight provinces never do so.

Mix and Match. We then examine whether provinces separate web-hosting and DNS providers. Many do, since some vendors offer only one service. Figure 7

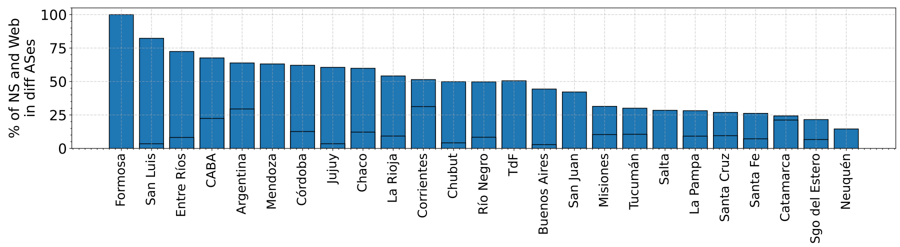

Fig. 7. Percentage of different hosting infrastructures for web and nameservers across Argentinian jurisdictions.

reports the share of domains whose web and name servers sit in different ASes; the average (median) split is 48% (50%).

Motivations vary. Formosa and Entre Ros, which host most content on sovereign or local networks, outsource DNS to global providers. CABA adopts a hybrid model, operating its own NSs while advertising a third-party NS as apparent failover. These cases illustrate how provinces balance local control with external resilience.

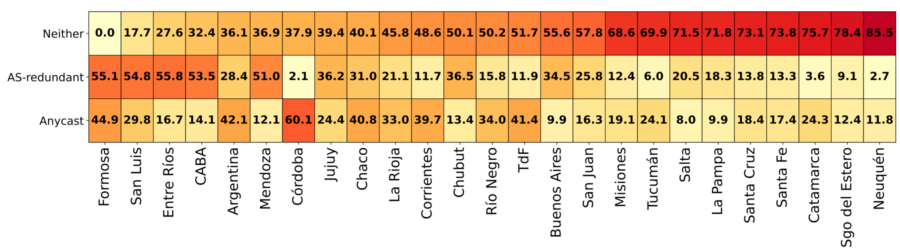

Fig. 8. Percentage distribution of nameserver redundancy mechanisms across Argentinian jurisdictions, classified into Anycast, AS-level redundancy, and no redundancy.

Redundacy Strategies. We also measure redundancy in authoritative DNS via two mechanisms: anycasted NS prefixes and placement of NS records across multiple ASes. Figure 8 shows the share of domains using anycast, AS-level redundancy, or neither; domains may employ both, so totals can exceed 100%.

Redundancy practices vary widely. Córdoba records the highest anycast adoption, yet every province uses anycast for at least some domains. AS-level diversification is also common, with pronounced peaks in Mendoza, Entre Ríos, and the Autonomous City of Buenos Aires (CABA). Despite these examples, 14 of the 24 jurisdictions place more than half of their domains in the "neither" category, indicating that a substantial portion of provincial namespaces still lack basic DNS resilience.

5.2 HTTPS Adoption

We next assess HTTPS adoption across the 24 provinces and the federal government. We classify domains into four categories: HTTP-only; HTTPS-only; both with redirect (80443); and both without redirect.

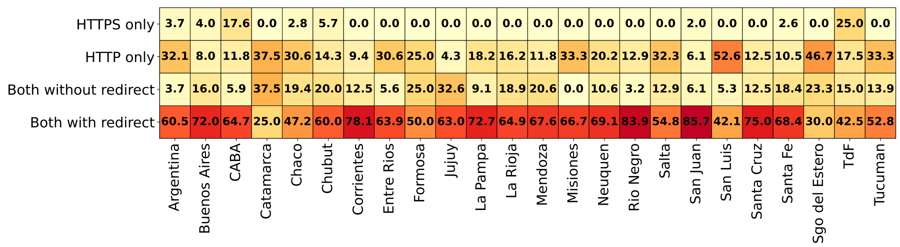

Fig. 9. Percentage distribution of HTTPS adoption across Argentinian jurisdictions, classified into HTTPS only, HTTP only, both with and without redirect.

Figure 9 presents a heatmap categorizing HTTPS adoption across the federal government and the 24 Argentine provinces. Our findings reveal that the "Both with redirect" configuration is the most prevalent, with an average adoption rate of 63.7% across the 25 jurisdictions, ranging from 25% to 89.4%. In contrast, the proportion of domains using only HTTP is lower at 18.5% on average. However, there is a significant disparity: 7 provinces have more than one-quarter of their domains operating as HTTP only, and 3 provinces exceed one-third.

Table 6. Counts and percentages of HTTPS connection errors.

Error Type	**Count**	**%**
Invalid or missing CA chain	51	37.0
Hostname mismatch	28	20.3
Max retries	21	15.2
Self-signed certificate	12	8.7
Read timed out	12	8.7
Expired certificate	10	7.2
SSL wrong version number	2	1.4
Connection reset by peer	1	0.7
Self-signed in chain	1	0.7

To understand HTTPS gaps (Table 6), we group failures into availability errors (e.g., "Max retries," "Read timed out"HTTPS unavailable) and certificate errors, which dominate. Certificate issues include broken chains, hostname mismatches, self-signed or expired certificates, incorrect SSL versions, and invalid roots.

These patterns suggest difficulty renewing certificates and maintaining services, reflecting a tendency to treat deployments as one-off launches rather than systems requiring ongoing operational upkeep.

6 Explainability Analysis

Sections 3–5 describe our measurement pipeline (seed and crawl construction, provider classification by operational control, and DNS/TLS hygiene checks). Here, we examine which *non-technical covariates* — demographic, economic, technological, and governance variables (see § 3.3) — are statistically associated with the observed hosting shares. We fit linear regression models to quantify these associations across provinces. While these models reveal consistent patterns in the data, they do not prove causality due to unobserved confounders and the limitations of provincial-level aggregation.

6.1 Model specification and model selection

We fit one model per provider class (Sovereign, Local, Global) and metric (bytes, domains, URLs). Let $y_{p,c}^{(m)} \in [0,1]$ denote the fraction served by class c, for province p and metric m. The specification of the model is:

$$y_{p,c}^{(m)} = \alpha_c^{(m)} + X_p^\top \beta_c^{(m)} + \varepsilon_{p,c}^{(m)}, \qquad (1)$$

where X_p is the vector of covariates for province p (standardized to mean 0 and variance 1), $\beta_c^{(m)}$ is the vector of coefficients for class c and metric m, $\alpha_c^{(m)}$ is an intercept term, and $\varepsilon_{p,c}^{(m)}$ is an error term.

In our linear regression framework, each coefficient $\beta_c^{(m)}$ indicates the change in hosting share associated with a one-unit change in the corresponding predictor, holding all other variables constant. We make no causal claims: coefficients should be interpreted as conditional correlations given the included covariates. Potential confounding, omitted variables, and reverse causality may remain.

6.2 Model fit

Given that the number of samples is low (n=24 provinces), we constrain the complexity of the model via forward stepwise selection [94]: features are added one at a time, choosing the one that most improves predictive performance in a leave-one-out cross-validation (LOOCV) framework, minimizing LOOCV mean squared error (MSE). In LOOCV, n-1 samples are used to train the model, and the held-out sample is used to evaluate performance; this process is repeated for each sample, and then the results are averaged. The procedure stops when no further improvement is possible. Then, the model is refit on the selected features to obtain final coefficient estimates.

The predictive capacity of the model is assessed using Pearson's correlation between observed and predicted hosting shares in a LOOCV (leave-one-out cross-validation) framework. We report these values in Table 7: we observe that sovereign-share models exhibit the highest out-of-fold correlations, while global/local shares are harder to predict.

Table 7. LOOCV Pearson correlation (r) between observed shares and predictions, by provider type and metric.

	Metric		
Provider	**Bytes**	**Domains**	**URLs**
Global	0.35	0.25	0.20
Local	0.37	0.20	0.22
Sovereign	0.76	0.83	0.70

6.3 Key associations

The statistical significance of each coefficient is assessed via its t-statistic, which is computed as the ratio of the estimated coefficient to its standard error. A higher absolute value of the t-statistic indicates a stronger relationship between the predictor and the hosting share, pointing out that the observed association is less likely to have occurred by chance. In Table 8, we report the features retained after forward stepwise selection for each model, along with their t-statistics, sorted by their importance, and including only those features with $|t| > 2$ (i.e., significant at approximately the 5% level under a two-tailed test).

Table 8. Features retained after forward stepwise selection. Entries show t-statistics from the *refit* on the selected set, including only those features with $|t| > 2$.

Provider Type	**Bytes**	**Domains**	**URLs**
Global	—	Internet speed (−2.1)	Internet speed (−2.0)
Local	Mobile Internet (−2.3)	Mobile Internet (−2.1)	Mobile Internet (−2.3)
Sovereign	Household Internet (−6.0) Illiteracy (5.6) Public employment (3.9) Inequality (−3.6) Internet speed (2.9)	Household Internet (−4.4) Tertiary education (−3.9) Illiteracy (3.2) Tax result (−3.1) GDP per capita (2.8) HDI (2.8) Congress diversity (2.8) Internet speed (3.0)	Household Internet (−4.1) Illiteracy (4.4) Tax result (−3.0) GDP per capita (2.5) Inequality (−2.1)

We also use SHAP (SHapley Additive exPlanations) [65] to explain the individual predictions of each model. SHAP computes feature importance by measuring how much each feature contributes to moving a prediction away from the baseline (mean) prediction, conditional on other included features. SHAP is based on cooperative game theory, specifically the concept of Shapley values, which fairly distribute the "payout" (prediction) among the "players" (features) based on their contributions. Analyzing the SHAP explanations offers a robustness check against the t-statistics derived from our model's fitted regression coefficients.

Figure 10 shows SHAP summaries for three selected models. In this plot, each point represents a single observation from the dataset. The horizontal position of a point indicates the SHAP value (contribution to the prediction), with positive values on the x-axis increasing predictions and negative values decreasing them. The color of each point represents the feature's value for that observation. Features are ranked vertically in the plot by their mean absolute SHAP value, with the most important features at the top.

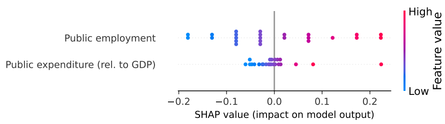

(a) Global providers (bytes)

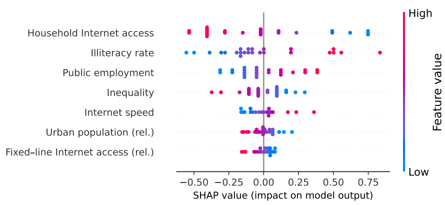

(b) Sovereign providers (bytes)

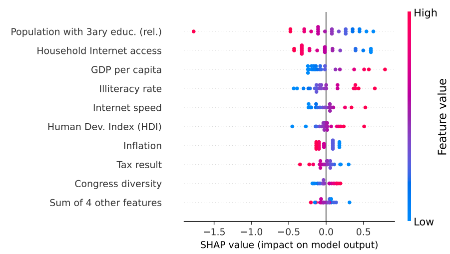

(c) Sovereign providers (domains)

Fig. 10. SHAP summaries for three selected models. Features are ordered by mean absolute SHAP (average impact on predictions across provinces). Each point is a province; color encodes the direction of the features contribution.

We observe in panel (b) that lower household Internet access is *associated* with higher local shares; panel (a) points out that higher public-employment levels are *associated* with greater global shares. For sovereign shares, higher illiteracy and lower tertiary attainment are *associated* with larger sovereign shares. These SHAP-based patterns align with the t-statistics in Table 8: household Internet emerges as the dominant predictor for sovereign hosting, while illiteracy and education variables also appear with consistent signs and magnitudes. However, we note that these patterns are correlational and may proxy for underlying factors (e.g., fiscal capacity, institutional maturity) not explicitly modeled here.

7 Dicussion

As with any study, our work has limitations. Below, we summarize the key limitations in concrete terms and offer steps to address them in follow-on work.

Scope of observation. Our dataset is restricted to pages reachable from the portal seeds and to content crawled up to seven levels deep. We therefore do not capture non-web infrastructure such as VDI or remote desktop systems, internal admin portals, private APIs, backend databases, cloud control planes, enterprise CDNs, log systems, or data-residency clauses. Any conclusions should be read as statements about the public web footprint only, not the complete content delivery approach of each government.

Portal-seed and discovery bias. Our list compiles public information from official portals and published organograms; we then scrape services linked from those pages. Services that are unlinked, contractor-hosted microsites, campaign pages, or payment gateways might not have been covered. To reduce this bias, future work should combine portal crawling with passive DNS data, registrar zone lists, certificate transparency logs, and coordinated disclosures from agencies.

Measurement design and proxy limits. Our analysis of DNS records, TLS observations, anycast detection, and AS mappings offers valuable indirect insights into how and where content is served. However, these are not direct metrics of user experience, uptime, or redundancy as per contracts. The accuracy of inferences can depend on the locations of the vantage points and the timing of the experiments. To mitigate these limitations, increasing probe diversity, extending sampling over time, and correlating with client-side measurements would reduce these vulnerabilities.

Aggregation and sample size. Province-level aggregates can hide important agency-level heterogeneity: different ministries or services within the same province often make distinct hosting choices. Statistical analyses based on only 24 provinces also have limited degrees of freedom and low statistical power; thus estimated effect sizes should be treated as exploratory

Missing contextual covariates and generalizability. We do not examine procurement records, costs, vendor selection rationales, or political drivers that motivate hosting choices. Without those covariates, causal interpretations are constrained. Institutional contexts also vary across countries, making our findings may not generalize without careful consideration of market structure and regulatory differences.

8 Related Work

Governments increasingly shape Internet infrastructure and online services — from granting access [26] and imposing politically motivated shutdowns [17] to delivering e-government portals [42,90,41] and enforcing privacy rules [38,30]. The measurement community has tracked these shifts accordingly.

Closest to our setting, Kumar et al. [62] map hosting choices for *61* countries, highlighting diverse mixes of on-premises and third-party providers. Jansen et al. [54] analyze government infrastructures in six states with tense neighbor relations. Hsiao et al. [42] study external dependencies of public-facing sites in the G7, and Jonker et al. [57] show that many Russian sites were domestically hosted pre-invasion and later lost Western providers — a pattern consistent with our observations. Reliance on single third-party DNS operators is documented by Sommese [90] and Houser [41]; other work examines HTTPS adoption [89], privacy risks [88], and cookie practices [38]. More broadly, consolidation across hosting, DNS, CDNs, and certificate authorities has gained notable attention from various perspectives far beyond government services [40,61,59,58,64].

To our knowledge, this is the first study to quantify *subnational* diversity in hosting choices within a single country under a shared legal and infrastructural framework, while also incorporating DNS and HTTPS hygiene and linking outcomes to non-technical covariates. Finally, our Argentina-focused analysis contributes to recent efforts to deepen understanding of Latin American network infrastructure [27,25,21,66,74,18,15].

9 Conclusions

Subnational governments shape public-sector infrastructure in ways that defy a single national template. In Argentina, we observe wide diversity – some provinces invest in sovereign footprints, others lean on domestic incumbents or global platforms – with patterns that reflect historical legacies and institutional autonomy as much as technical capacity. We report *associations*, not causal effects, and contextualize sovereigntyperformance trade-offs with latency and cache-locality sanity checks.

Beyond this case, we distill a reusable measurement template (domain discovery → provider classification by operational control → DNS/HTTPS hygiene and resilience proxies) and subnational benchmarks others can replicate. Limits include portal-seed bias and that bytes measure retrieved content size, not traffic. A natural next step is a second country (or two) and deeper performance/cost slices to test which invariants travel.

Acknowledgements. This work was supported in part by NSF grants CNS-2246475 and CNS-2107392. The views expressed are solely the authors'. We thank the reviewers and our shepherd for their helpful feedback.

References

1. Digital service standard. www.digital.gov.au
2. Estándares digitales y guías. digital.gob.cl
3. An introduction to federal website standards. digital-gov-static-prod.app.cloud.gov
4. Nz government web standards. www.digital.govt.nz
5. Section 508 of the rehabilitation act. www.section508.gov
6. Web standards. webstandards.ca.gov
7. Government technology standards and guidance. www.gov.uk (2016)
8. Guía de estándares de calidad e interoperabilidad de los datos abiertos del gobierno de colombia. herramientas.datos.gov.co (2020)
9. Digital accessibility best practices guide. www.gov.br (2023)
10. Agência Brasil: Centro de dados da procergs é invadido pela água e desligado; há risco de interrupção de serviços. agenciabrasil.ebc.com.br (05 2024)
11. ARSAT: Arsat en el tiempo. www.arsat.com.ar (2025)
12. ARSAT: Red federal de fibra óptica. www.arsat.com.ar (2025)
13. Banco Central de la República Argentina: Informe de inclusión financiera - feb 2020. www.bcra.gob.ar (2020)

14. Benavides, J.L., M'Causland Sánchez, M.C., Flórez Salazar, C., Roca, M.E.: Public procurement in latin america and the caribbean and idb-financed projects: A normative and comparative study. Tech. rep., IDB (2016)
15. Berenguer, S.S., Carisimo, E., Alvarez-Hamelin, J.I., Pintor, F.V.: Hidden internet topologies info: Truth or myth? In: LANCOMM (2016)
16. Bertholdo, L.M., Paredes, R., de Lima Marin, G., Loureiro, C.A., Kashwakura, M.K., de Botelho Marcos, P.: Analyzing the effect of an extreme weather event on telecommunications and information technology: Insights from 30 days of flooding. In: Proc. of PAM (2025)
17. Bischof, Z.S., Pitcher, K., Carisimo, E., Meng, A., Bezerra Nunes, R., Padmanabhan, R., Roberts, M.E., Snoeren, A.C., Dainotti, A.: Destination unreachable: Characterizing internet outages and shutdowns. In: Proc. of ACM SIGCOMM (2023)
18. Bischof, Z.S., Rula, J.P., Bustamante, F.E.: In and out of cuba: Characterizing cuba's connectivity. In: Proc. of IMC (2015)
19. Bolsa de Comercio de Rosario: Cómo fue el desempeño del pib per cápita. www.bcr.com.ar (2025)
20. Boston Consulting Group: Trust imperative 4.0: Genaithe trust multiplier for government. Tech. rep. (2024)
21. Brito, S.H.B., Santos, M.A., Fontes, R.d.R., Perez, D.A.L., Rothenberg, C.E.: Dissecting the largest national ecosystem of public internet exchange points in Brazil. In: Proc. of PAM (2016)
22. CAIDA: Routeviews prefix to as mappings dataset (pfx2as) for ipv4 and ipv6. www.caida.org (2008)
23. CAIDA: Peeringdb dump - 1 jan 2025. publicdata.caida.org (2025)
24. Cámara Nacional Electoral: Listado de partidos políticos reconocidos. www.argentina.gob.ar (2025)
25. Carisimo, E., Fiore, J.M.D., Dujovne, D., Pelsser, C., Alvarez-Hamelin, J.I.: A first look at the latin american ixps. ACM SIGCOMM Computer Communication Review (mar 2020)
26. Carisimo, E., Gamero-Garrido, A., Snoeren, A.C., Dainotti, A.: Identifying ases of state-owned internet operators. In: Proc. of IMC (2021)
27. Carisimo, E., Kumar, R., Wang, C.J., Klein, S., Bustamante, F.E.: Ten years of the venezuelan crisis - an internet perspective. In: Proc. of ACM SIGCOMM (08 2024)
28. Congreso de la Nación Argentina: Ley 25.326 (proteccin de datos personales). argentina.gob.ar (2000)
29. Congreso de la Nación Argentina: Ley 27.275 (acceso a la informacin pblica). argentina.gob.ar (2016), boletín Oficial
30. Dabrowski, A., Merzdovnik, G., Ullrich, J., Sendera, G., Weippl, E.: Measuring cookies and web privacy in a post-gdpr world. In: Proc. of PAM (2019)
31. Dener, C., Ghunney, L.E., Johns, K., Nii-Aponsah, H.: Govtech maturity index: The state of public sector digital transformation. Tech. rep., World Bank (2021)
32. Dirección General de Estadística y Censos, GCBA: Indicadores de analfabetismo 2014. www.estadisticaciudad.gob.ar (2014)
33. Dirección Nacional de Población: Análisis de la movilidad residencial interna en argentina. www.argentina.gob.ar (2023)
34. ENACOM: Acceso a internet fija - 2022 t2. indicadores.enacom.gob.ar (2022)
35. García Zaballos, A., Iglesias Rodríguez, E.: Cloud Computing: Opportunities and Challenges for Sustainable Economic Development in Latin America and

the Caribbean. IDB (2018), https://publications.iadb.org/publications/english/document/Cloud-Computing-Opportunities-and-Challenges-for-Sustainable-Economic-Development-in-Latin-America-and-the-Caribbean.pdf
36. de Gobierno Digital, D.: Guía para el diseño de interfaces web institucionales. digital.gob.cl (2019)
37. Google Chrome User Experience Report: Crux - top sites lists. github.com (2025)
38. Gotze, M., Matic, S., Iordanou, C., Smaragdakis, G., Laoutaris, N.: Measuring web cookies in governmental websites. In: Proc of. ACM Web Sciences (2022)
39. Governo do Estado do Rio Grande do Sul: Procergs reativa data center e prevê retorno de serviços a partir de 27 de maio. estado.rs.gov.br (05 2024)
40. Habib, R., Ruth, K., Akiwate, G., Durumeric, Z.: Formalizing dependence of web infrastructure. In: Proc. of ACM SIGCOMM (2025)
41. Houser, R., Hao, S., Cotton, C., Wang, H.: A comprehensive, longitudinal study of government dns deployment at global scale. In: Proc of IEEE DSN (2022)
42. Hsiao, H.C., Kim, T.H.J., Ku, Y.M., Chang, C.M., Chen, H.F., Chen, Y.J., Wang, C.W., Jeng, W.: An investigation of cyber autonomy on government websites. In: Proc. of the WWW (2019)
43. IIEP, Universidad de Buenos Aires: Reporte especial - resultado fiscal de las provincias. iiep.economicas.uba.ar (2024)
44. INDEC: Censo 2022 - educación: población con nivel terciario completo. www.indec.gob.ar (2022)
45. INDEC: Censo 2022 - población total por provincia. censo.gob.ar (2022)
46. INDEC: Población urbana por provincia (cuadro n° 020210). www.indec.gob.ar (2022)
47. INDEC: Acceso a internet - informe 2024. www.indec.gob.ar (2024)
48. INDEC: índice de precios al consumidor - dic 2024. www.indec.gob.ar (2024)
49. INDEC: Informe de pobreza (eph). www.indec.gob.ar (2024)
50. INDEC: Mercado de trabajo - eph 3.er t 2024. www.indec.gob.ar (2024)
51. INDEC: Esperanza de vida al nacer. www.indec.gob.ar (2025)
52. INDEC: Promedio de años de escolaridad. www.indec.gob.ar (2025)
53. ITIF, TicTac: Evaluación de los sitios web del gobierno colombiano. www.ccit.org.co (Sep 2020)
54. Jansen, B., Kadenko, N., Broeders, D., van Eeten, M., Borgolte, K., Fiebig, T.: Pushing boundaries: An empirical view on the digital sovereignty of six governments in the midst of geopolitical tensions. Government Information Quarterly (2023)
55. JGM: Arsat empresa argentina de soluciones satelitales sociedad anónima (arsat). www.argentina.gob.ar (2025)
56. JGM: Información provincial sobre empleo público. mepp.jefatura.gob.ar (2025)
57. Jonker, M., Akiwate, G., Affinito, A., Claffy, k., Botta, A., Voelker, G.M., van Rijswijk-Deij, R., Savage, S.: Where .ru? assessing the impact of conflict on russian domain infrastructure. In: Proc. of IMC (2022)
58. Kashaf, A., Dou, J., Belova, M., Apostolaki, M., Agarwal, Y., Sekar, V.: A first look at third-party service dependencies of web services in africa. In: Proc. of PAM (2023)
59. Kashaf, A., Sekar, V., Agarwal, Y.: Analyzing third party service dependencies in modern web services: Have we learned from the mirai-dyn incident? In: Proc. of IMC (2020)
60. KPMG Argentina: Créditos y depósitos provinciales. assets.kpmg.com (2012)

61. Kumar, R., Asif, S., Lee, E., Bustamante, F.E.: Each at its own pace: Third-party dependency and centralization around the world. Proc. of ACM SIGMETRICS (2023)
62. Kumar, R., Carisimo, E., De Angelis Rivas, L., Buzzone, M., Bustamante, F.E., Qazi, I.A., Beiró, M.G.: Of choices and control: A comparative analysis of government hosting. In: Proc. of IMC (Nov 2024)
63. Labs, A.: Visible asns: Customer populations (est.) argentina. stats.labs.apnic.net (2025)
64. Liu, E., Akiwate, G., Jonker, M., Mirian, A., Savage, S., Voelker, G.M.: Who's got your mail? characterizing mail service provider usage. In: Proc. of IMC (2021)
65. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. Advances in neural information processing systems **30** (2017)
66. Mazzola, F., Marcos, P., Barcellos, M.: Light, camera, actions: characterizing the usage of ixps' action bgp communities. In: Proc. of CoNEXT (2022)
67. MECON: Decreto 62/90. mepriv.mecon.gob.ar
68. MECON: Decreto 60/90. mepriv.mecon.gob.ar (2025)
69. MECON: Ley 23.696. mepriv.mecon.gob.ar (2025)
70. MECON: Origen y ámbito de acción - telecom. mepriv.mecon.gob.ar (2025)
71. MECON: Origen y ámbito de acción - telefónica. mepriv.mecon.gob.ar (2025)
72. Ministerio de Modernización: Decálogo Tecnológico ONTI: Versión 1.1.0. Tech. rep., Presidencia de la Nacin (Argentina), Buenos Aires, Argentina (Nov 2019), publicado el 8 de noviembre de 2019
73. de Modernizacin, M.: Disposicin onti 6/2019 (accesibilidad web 2.0). www.argentina.gob.ar (Sep 2019), boletn Oficial, 30 de septiembre de 2019
74. Müller, L., Luckie, M., Huffaker, B., Claffy, K., Barcellos, M.: Challenges in inferring spoofed traffic at ixps. In: Proc. of CoNEXT (2019)
75. NIC.ar: 30 años de nic.ar. nic.ar (2025)
76. Observatorio Federal Urbano: Coeficiente gini - provincias. ofu.obraspublicas.gob.ar (2025)
77. OECD: Digital Government Review of Argentina: Accelerating the Digitalisation of the Public Sector. Paris (2019), https://doi.org/10.1787/354732cc-en
78. OECD: Good Governance for Critical Infrastructure Resilience (2019), https://doi.org/10.1787/02f0e5a0-en
79. OECD: Good practice principles for public service design and delivery in the digital age. Tech. rep. (2022)
80. OECD: Digital transformation of public procurement: Good practice report. Tech. rep., Public Governance Directorate (2025)
81. Oficina de Presupuesto del Congreso: Gasto público provincial sobre pib. opc.gob.ar (2024)
82. P.E.N.: Decreto 1131/2016 (gde). argentina.gob.ar (2016)
83. P.E.N.: Decreto 13/2016 (modificacin del decreto 357/2002). www.argentina.gob.ar (Jan 2016), sancionado el 5 de enero de 2016; publicado en el Boletn Oficial el 6 de enero de 2016
84. Porto, A.: Transferencias intergubernamentales y disparidades fiscales a nivel subnacional en argentina. Tech. rep., Inter-American Development Bank (2016)
85. Porto, A., Puig, J.P.: On the fiscal behavior of subnational governments. a long-term vision for argentina. Tech. rep., Buenos Aires, Argentina (2022), https://aaep.org.ar/works/works2022/4588.pdf
86. Programa de las Naciones Unidas para el Desarrollo: El mapa del desarrollo humano en argentina. www.undp.org (2023)

87. Radics, A., Eguino, H.: Próximos pasos para la descentralización y gobiernos subnacionales en américa latina y el caribe. Tech. rep., IDB (Oct 2018), https://publications.iadb.org/publications/spanish/document/Proximos-pasos-para-la-descentralizacion-y-gobiernos-subnacionales-en-America-Latina-y-el-Caribe.pdf
88. Samarasinghe, N., Adhikari, A., Mannan, M., Youssef, A.: Et tu, brute? privacy analysis of government websites and mobile apps. In: Proc. of the WWW (2022)
89. Singanamalla, S., Jang, E.H.B., Anderson, R., Kohno, T., Heimerl, K.: Accept the risk and continue: Measuring the long tail of government https adoption. In: Proc. of IMC (2020)
90. Sommese, R., Jonker, M., van der Ham, J., Moura, G.C.M.: Assessing e-government dns resilience. In: Proc. of CNSM) (2022)
91. Swanson, M., Bowen, P., Phillips, A., Gallup, D., Lynes, D.: Contingency planning guide for federal information systems. Tech. rep., NIST (2010)
92. Transparencia Electoral: Mapa de integridad electoral. transparenciaelectoral.org (2025)
93. U.S. General Services Administration: Federal website standards. standards.digital.gov
94. Wasserman, L.: All of statistics: a concise course in statistical inference. Springer Science & Business Media (2013)
95. World Bank: Building resilience against disasters in latin america and the caribbean. www.worldbank.org (2023)

A Ethical Considerations

Our study reveals provincial infrastructure choices while minimizing harm: we limited collection to public, client-facing pages, performed a single, user-like crawl with no logins or vulnerability probing, and kept our load on government servers negligible. Our results are reported only in aggregate — domain-to-provider mappings are omitted to prevent reconstruction of targets — and the compiled seed dataset will be shared only for replication under a non-redistribution agreement. By constraining granularity, omitting sensitive attribution, and controlling data access, we comply with ethical best practices for Internet measurement while enabling independent validation.

B Explanatory Variables

Table 9 summarizes the variables that have been considered to explain the hosting decisions of the provinces, grouped by type and including their sources and descriptions.

Table 9: Summary of variables, sources, and descriptions

Type	Variable	Source
	Inequality	Observatorio Federal Urbano, Ministerio del Interior, Obras Públicas y Vivienda [76]
	HDI	United Nations Development Programme [86]
	Avg. years of schooling	INDEC [52]

Continued on next page

Demographic

Table 9 – *Continued from previous page*

Type	Variable	Source
	Population with tertiary education	INDEC [44]
	Population	INDEC [45]
	Urban population	INDEC [46]
	Migration rate	Dirección Nacional de Población, Ministerio del Interior [33]
	Life expectancy	INDEC [51]
	Illiteracy rate	Dirección General de Estadística y Censos, GCBA [32]
	Poverty	EPH (INDEC) [49]
	Financial inclusion	BCRA [13]
Economic	GDP per capita	Bolsa de Comercio de Rosario [19]
	Unemployment rate	EPH (INDEC) [50]
	Public expenditure	Oficina de Presupuesto del Congreso [81]
	Credit	KPMG Consulting [60]
	Tax result	IIEP – UBA [43]
	Inflation	INDEC [48]
	Public employment	Jefatura de Gabinete de Ministros [56]
Technological	Internet penetration	ENACOM [34]
	Fixed–line Internet access	INDEC [47]
	Mobile Internet access	INDEC [47]
	Household Internet access	INDEC [47]
	Enterprise Internet access	INDEC [47]
	Internet speed	ENACOM
	Peering facilities	CAIDA [23]
	Registered ASes	LACNIC + RDAP + OSM + Geoapify + OpenCageData
	Government websites	Google CrUX [37]
Governmental	Democracy score	Transparencia Electoral [92]
	Congress diversity	—
	Political diversity	Cámara Nacional Electoral [24]

C Hosting Providers Operating in a Single Province

Table 10 presents the 18 hosting providers operating hosting content of a single province.

Table 10. List of local providers serving exactly one province. Notably, TV MUSIC HOUSE JUJUY (AS52312) and Sebastian Souto (AS264738) account for the largest share of served bytes for the governments of Jujuy and Santa Cruz, respectively.

ASN-AS Name	Province	URLs	Dom.	Bytes
52312-TV MUSIC HOUSE JUJUY	Jujuy	64.0	60.0	45.9
264738-Sebastian Souto	Santa Cruz	15.0	8.3	45.7
27879-IyT	Mendoza	11.9	8.3	20.3
61449-RESEARCH	Misiones	2.1	2.0	5.8
262187-Patagonia Green	Entre Ríos	2.6	2.3	4.6
28009-Davitel	Neuquén	5.4	4.6	3.9
272150-TELECOMUNICACIONES	Entre Ríos	2.6	2.3	3.1
27747-Telecentro	CABA	3.8	4.3	2.8
61493-BAEHOST	Mendoza	2.8	5.6	2.7
265781-PALA PABLO FEDERICO	Santa Fe	1.0	2.1	2.1
27983-Red Intercable Digital	TdF	0.7	1.8	1.5
52236-G2K ARGENTINA	Jujuy	1.8	2.0	0.6
264689-Luciano Gabriel Chersanaz	Santa Cruz	3.3	4.2	0.5
22927-Telefonica de Argentina	Mendoza	0.7	5.6	0.4
52351-Integral Insumos	Santa Fe	1.0	2.1	0.4
7049-Silica Networks	Tucumán	1.7	1.5	0.3
265749-RBA	La Rioja	1.4	2.4	0.2
61443-KPMG Argentina	Neuquén	0.6	0.9	0.2