

Characterizing Iran’s Phased National Internet Shutdown in 2025: A Progressive and Distributed Action

Shibo Cui
Tsinghua University
BNRist
Beijing, China
csb24@mails.tsinghua.edu.cn

Mingxuan Liu*
Zhongguancun Laboratory
Beijing, China
liumx@mail.zgclab.edu.cn

Baojun Liu
Tsinghua University
BNRist
Beijing, China
lbj@tsinghua.edu.cn

Haixin Duan
Tsinghua University
Beijing, China
duanhx@tsinghua.edu.cn

Ruixuan Li
Tsinghua University
Beijing, China
lirx25@mails.tsinghua.edu.cn

Chaoyi Lu
Zhongguancun Laboratory
Beijing, China
lucy@zgclab.edu.cn

Jin Zhang
Tsinghua University
Beijing, China
zj24@mails.tsinghua.edu.cn

Zhicheng Wang
Qi An Xin Technology Group Inc.
Beijing, China
wangzhicheng@qianxin.com

Jinghua Bai
Qi An Xin Technology Group Inc.
Beijing, China
baijinghua@qianxin.com

Abstract

In June 2025, the Iranian government executed a nationwide shutdown. This shutdown did not employ traditional large-scale BGP route withdrawals; instead, it relied on service-level restrictions. This shift in shutdown strategy renders existing passive-traffic-based monitoring and network-level active probing systems ill-equipped to capture the event’s fine-grained characteristics.

To address this gap, we develop a service-level shutdown monitoring framework. Leveraging continuous, large-scale active port scanning data of Iran’s entire IPv4 space, we treat the collected 8.65 million results as a statistically representative sample of the nation’s network state. Then, we identify shutdowns by detecting significant drops in service activity against a dynamic baseline computed via an adaptive sliding window. Based on our monitoring results, we reveal that this shutdown was not a monolithic event, but a sophisticated operation with three core properties: phased, progressive, and distributed. Specifically, the operation unfolded in four distinct phases: it began with two complementary localized drills that shifted focus from infrastructure control to information obstruction, escalated into a near-total nationwide blockade, and concluded with a tiered, censorship-oriented recovery. This escalation was progressive, with the blockade’s scope expanding to impact 98 of the top 100 ASes and 49 of the top 50 network services. Furthermore, we find significant heterogeneity in the shutdown’s impact and recovery across different ASes, indicating a distributed enforcement architecture. Beyond these primary characteristics, our analysis reveals deeper consequences. This nationwide shutdown action also caused collateral impacts, such as unexpected

port exposure and traffic surges. Ultimately, we publicly release this aggregated network service dataset covering this action, calling for broader discourse on the mechanisms and impact of network shutdown actions to maintain the resilience and connectivity of the global Internet.

CCS Concepts

• **Networks** → **Network measurement; Network monitoring;**
• **Social and professional topics** → **Network access control; Censorship.**

Keywords

Internet Shutdown; Network Measurement; Network Monitoring

ACM Reference Format:

Shibo Cui, Mingxuan Liu, Baojun Liu, Haixin Duan, Ruixuan Li, Chaoyi Lu, Jin Zhang, Zhicheng Wang, and Jinghua Bai. 2026. Characterizing Iran’s Phased National Internet Shutdown in 2025: A Progressive and Distributed Action. In *Proceedings of the ACM Web Conference 2026 (WWW ’26)*, April 13–17, 2026, Dubai, United Arab Emirates. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3774904.3792699>

Resource Availability:

The dataset of this paper has been made publicly available at <https://doi.org/10.5281/zenodo.17236149>.

1 Introduction

Internet service availability underpins global network interconnection and serves as a critical cornerstone for securing cross-regional cyber interactions in politics, economy, culture, and other fields. The Internet community has made substantial efforts to this end. However, the frequent occurrence of Internet censorship [2, 20, 54] or government-ordered shutdowns [1] continues to underscore the fragility of Internet infrastructure.

In June 2025, Iran’s nationwide network shutdown event attracted widespread international attention. After Israeli strikes on

* Corresponding author.



Iran on June 13 [64], 2025, Iran’s Ministry of Communications announced temporary restrictions on the nation’s Internet [18, 33]. This action intensified on June 17 with a significant decline in traffic, leading to a near-complete service disruption by midday on June 18. Recovery began with partial recovery on June 21, followed by large-scale restoration on the 25th [43].

Iran has a long-standing history of Internet censorship spanning over a decade and has implemented large-scale shutdowns on previous occasions [22]. A notable prior instance occurred on November 16, 2019, when a week-long shutdown was implemented following widespread protests [31]. That event was primarily characterized by *network-level controls*, specifically the large-scale BGP withdrawal of approximately 33% of the country’s IPv4 prefixes, which resulted in a substantial drop in connectivity.

In contrast, the 2025 shutdown did not exhibit similar large-scale IPv4 prefix withdrawals [8]. Instead, the strategy shifted towards the *service-level*, manifesting as a series of fine-grained restrictions targeting specific network services (e.g., DNS, SSH, HTTP, IKEv2). *This evolution from coarse-grained, network-level blocking to fine-grained, service-level restrictions signifies a growing sophistication in shutdown techniques, posing new challenges for shutdown monitoring.*

Existing shutdown monitoring methods exhibit limitations in terms of observation scope, comprehensiveness, and granularity. First, methods based on the routing network control plane (e.g., BGP Eye [53], I-Seismograph [65]) suffer from limited monitoring scope, as no obvious routing announcement anomalies occurred during Iran’s action. Second, passive-traffic-based methods on the data plane (e.g., Internet Background Radiation [10, 19, 26], Cloudflare Radar [16]) have comprehensiveness biases in their observations due to reliance on service types covered by their data sources. Finally, existing network-level active probing methods for network connectivity monitoring (e.g., Trinocular [44]) mainly focus on the liveness of target /24 network segments via ICMP ping [9], limited to accurately and sensitively characterize service-level anomalies at a fine granularity. *Thus, for this novel shutdown event, a more fine-grained shutdown monitoring method is required.*

A Service-Level Shutdown Monitoring Framework. To capture the nuances of modern shutdown strategies, we propose a novel service-level shutdown monitoring framework. Our approach leverages high-frequency, large-scale port scan network service data from HUNTER¹, a commercial cyberspace search engine indexing nearly 30 billion network assets, which continuously probes thousands of ports across Iran’s entire IPv4 space from globally distributed vantage points. The core of our methodology is an algorithm that treats the time series of observed active network services within discrete time intervals (10 minutes) as a direct proxy for the nation’s network state. To detect disruptions, we employ a three-day sliding window that is adaptive—it automatically excludes previously identified shutdown periods to ensure the baseline of normal activity is not skewed. A shutdown is flagged if the Service Blocked Ratio (SBR)—the proportional drop in network activity relative to the baseline—exceeds a statistically calibrated critical threshold. This framework enables us to identify and quantify service-level disruptions with high precision and temporal granularity.

¹The search engine is available at <https://hunter.qianxin.com/>

Main Findings. Through service-level monitoring of 8.65 million active network services in Iran during June 2025, we reconstructed the nation’s network shutdown event. Our analysis reveals that this shutdown was not a one-off event, but a sophisticated operation characterized by three core properties: **phased, progressive, and distributed.**

- *Phased Shutdown Actions.* The shutdown was a calculated operation unfolding in four phases that revealed a clear strategic evolution. It began with two complementary localized drills that shifted focus from infrastructure control to information obstruction. This escalated into a near-total nationwide blockade, which then transitioned into a tiered and censorship-oriented recovery.
- *Progressive Rollout and Recovery.* The shutdown’s logic was consistently progressive, governing both its rollout and recovery: the blockade systematically escalated from impacting 36 to 98 (of the top 100) ASes and 12 to 49 (of the top 50) services, while the recovery was tiered and incomplete, leaving a majority of both restricted.
- *Distributed Enforcement Architecture.* We observe significant heterogeneity in the initiation and recovery times of the shutdown across Iran’s network operators. The evidence points toward a distributed enforcement architecture, where blocking policies are applied at multiple, distributed points within the national network, resulting in the observed asynchronicity.

Furthermore, we reveal two *deeper consequences*. First, analysis of data from OONI [54] shows the shutdown operated in tandem with Iran’s existing censorship mechanism, creating a multi-layered “defense-in-depth” control system in which anomaly rates on surviving traffic soared past 75%. Second, the event caused unexpected collateral impacts: during the shutdown, our analysis identified significant DNS traffic surges, including a 115% peak increase in DNS queries, while the hasty recovery led to anomalous port exposure, evidenced by a 22% surge in exposed SSH servers.

Contribution. From our monitoring results, our analysis reveals that the June 2025 Iran event demonstrates *an evolution in its national shutdown strategies*: a shift from coarse, network-level blocking to a sophisticated paradigm of fine-grained, service-level control, in which blocking can be selectively applied to specific services. In light of this evolution, the service-level shutdown monitoring framework we introduce in this paper is a critical first step toward comprehensive, service-level aware monitoring of modern network connectivity. Furthermore, we have publicly released the unique network service dataset used in this work to support related research.² Finally, our work calls for more efforts to collectively maintain the resilience and connectivity of the global Internet.

2 Background and Related Work

In this section, we first review existing work on Internet shutdown and censorship. We then outline established techniques for Internet shutdown monitoring.

2.1 Internet Shutdown & Censorship

Internet shutdown and *censorship* represent distinct blocking strategies on a spectrum of impact, both employed by governments to control network connection. Censorship typically involves the selective blocking of specific news media websites [62], social media

²The dataset is available at <https://doi.org/10.5281/zenodo.17236149>.

platforms [14], or circumvention tools [41], whereas a shutdown aims to sever a region’s connectivity to the global Internet [35, 60].

Extensive research has been dedicated to Internet censorship [27, 34, 45, 50, 56, 61, 63], facilitated by numerous public measurement platforms like OONI (volunteer-based) [54], ICLab (VPN-based) [13], and Censored Planet (remote measurement-based) [57]. Previous studies have identified Iran’s primary censorship techniques, including DNS poisoning, HTTP blockpage injection, HTTPS connection disruption, and UDP traffic dropping [6, 11, 52]. More recent work has highlighted that censorship in Iran is not monolithic, observing inconsistencies across different Internet Service Providers (ISPs) [4].

Techniques for implementing nationwide network shutdowns span multiple levels, from physical-level disruptions like fiber cuts to network-level BGP route withdrawals [25, 40]. Concerning the June 2025 event, multiple studies have observed a key anomaly: a large-scale blockade was achieved while BGP route advertisements remained globally stable [7, 8]. This phenomenon marks a shift from network-level to service-level control.

2.2 Internet Shutdown Monitoring

Precisely identifying and characterizing Internet shutdowns is a critical endeavor for the research community, essential for bolstering the resilience of the global Internet [39]. Prominent real-time shutdown monitoring systems (such as Cloudflare Radar [16], IODA [30], NetBlocks [36], Kentik [3], and ThousandEyes [55]) often employ a hybrid of following techniques.

Passive-traffic-based monitoring. This includes *control plane* methods that analyze BGP updates for routing anomalies (e.g., BGP Eye [53], I-Seismograph [65] and [21], [23], [12]), and *data plane* methods that inspect traffic from sources like NetFlow traffic [49], BitTorrent traffic [15], CDN logs [46], user’s HTTP/DNS traffic [16], Internet Background Radiation [10, 19, 26] and long-running TCP connections [51] to detect connectivity changes.

Network-level active probing. This involves sending probes to actively assess the liveness of hosts at the network’s edge. For example, Trinocular [44] detects shutdowns by actively probing all /24 IPv4 prefixes with ICMP echo requests (i.e., “pings”) in a short cycle, assessing the reachability of approximately 5 million blocks every 11 minutes. Subsequent work [9] improves the accuracy of such active probing methods for sparsely populated address blocks by proposing a new Full Block Scanning (FBS) algorithm.

However, these established techniques share a common limitation in their observational scope. They primarily focus on *network-level* liveness or passively observe a narrow subset of *service-level* traffic (e.g., Cloudflare Radar’s reliance on HTTP/DNS). Consequently, they are ill-equipped to perform a fine-grained analysis of the recent Iran’s shutdown, which was defined by the temporal evolution of blocking strategies against dozens of mainstream network services (e.g., DNS, SSH, HTTP, IMAP, IKEv2). Our work addresses this critical gap by introducing a methodology for *service-level shutdown monitoring*.

3 Service-level Monitoring Framework

In this section, we illustrate our methodology for monitoring and characterizing Iran’s nationwide Internet shutdown in 2025. As

shown in Figure 1, we first introduce our *Network Service Data Collector*, and then detail the *Service-level Shutdown Monitor*.

3.1 Network Service Data Collector

We build the *Network Service Data Collector* to reconstruct Iran’s network connectivity from an external perspective, which serves as the basis for service-level Internet shutdown monitoring. This collector is built upon data from HUNTER, a cyberspace search engine maintained by Qi An Xin, a well-known cybersecurity firm. This engine periodically conducts large-scale, active measurements on the entire active IPv4 address space, covering over 6,000 ports to ensure comprehensive observation of various network services. The data collection process operates in three key steps:

① **Randomize Measurement Queue.** To minimize the impact on the target network, the engine’s scheduler randomizes the measurement queue (a list of $\langle ip, port \rangle$ pairs). It uses a custom cryptographic algorithm called *Blackrock*, which was designed for the well-known, open-source network scanner MASSCAN [24]. This algorithm encrypts an index to create random permutations of the IPv4 address pool, shuffling measurement tasks for contiguous IP addresses and ports.

② **Multi-protocol Probes.** The scan engine uses a globally distributed infrastructure with over 400 Vantage Points (VPs) located in data centers across Tokyo, Singapore, and Frankfurt, among other places. These VPs conduct multi-protocol probes using more than 200 protocol templates, including HTTP, SSH, DNS, and SNMP, to accurately identify the services running on each port. For each port, the scan engine attempts to probe using the default network service assigned by IANA [29], as well as some of the most prevalent services such as HTTP(S), RDP, etc.

③ **Data Ingestion and Storage.** The Vantage Points collect scanning results in real time. The data, which includes IP, port, and the detected network service name, is then streamed through a Kafka cluster and written into the ClickHouse database for later analysis.

Data Overview. In June 2025, we collected data on 8,650,874 (8.65M) active network services in Iran. This dataset covers 490 Iranian autonomous systems (ASes) and encompasses 192 distinct services, such as HTTP(S), SSH, DNS, SNMP, and IKEv2, across 6,130 ports.

3.2 Service-Level Shutdown Monitor

Based on our extensive collection of network service data, we design a novel *Service-level Shutdown Monitor*. The process for detecting network shutdowns involves three key steps:

④ **Aggregate Network Service Data and Formalize Network Activity.** The first step is to construct a time series that quantitatively represents Iran’s network activity. To achieve this, we divide the study period of June 2025 into a series of contiguous, non-overlapping intervals (t_1, t_2, \dots) , each with a duration of $\delta t = 10$ min. For each interval t , we aggregate the observed network services to obtain the **Observed Network Activity** (n_t), defined as the total number of active services detected at t . Next, we introduce the two-step process for constructing n_t .

Search Space. We first define the *Network Service Search Space* (\mathcal{S}), which contains all potential targets of our monitoring. It is the set of all tuples $s = \langle ip, port \rangle$ composed of all publicly routable IPv4 addresses in Iran (approximately 10^7 , estimated from the daily count

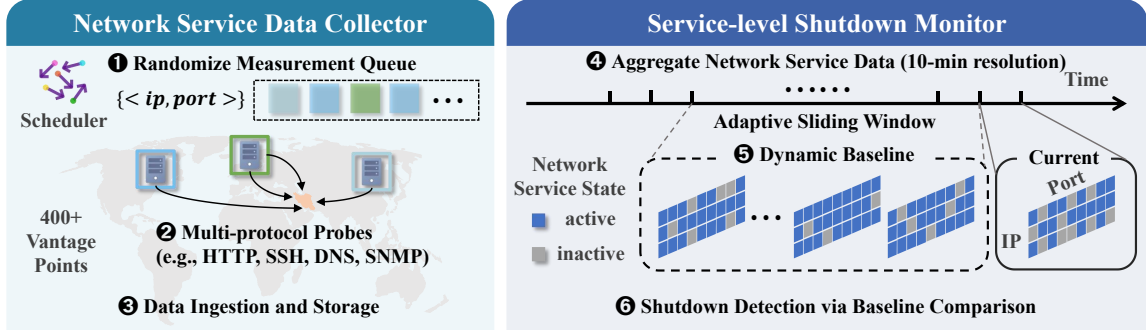


Figure 1: Workflow of Service-level Shutdown Monitoring Framework.

of reachable IPv4 addresses in June from the public routing dataset RouteViews [58] project) and the roughly 6,000 ports scanned by the cyberspace search engine. Let the total size of this space be $|\mathcal{S}| = K \approx 6 \times 10^{10}$. At any given time t , the ideal state of this space is described by a *Network State Function* ($\Phi(t)$), a K -dimensional binary vector where each component is ‘1’ for an active service and ‘0’ otherwise. The true number of active services in the network, or the *Overall Network Activity* (N_t), is the sum of all these components:

$$N_t = \sum_{j=1}^K \phi_j(t) \quad (1)$$

Observed Network Activity as a Reliable Proxy. Since a complete census to determine the true network activity (N_t) is infeasible within short intervals, our method instead estimates it by applying principles from sampling theory [17] to a large-scale, randomized sample. Specifically, we measure the *Observed Network Activity* (n_t), defined as the number of active services detected within a probed subset $\mathcal{P}_t \subseteq \{1, 2, \dots, K\}$:

$$n_t = \sum_{j \in \mathcal{P}_t} \phi_j(t) \quad (2)$$

The randomization of the measurement queue (Step 1) ensures this sample is statistically representative of the entire population. This allows us to use the sample proportion, $\hat{p}_t = n_t/|\mathcal{P}_t|$, as the Maximum Likelihood Estimator (MLE) and a consistent estimator for the true proportion $p_t = N_t/K$. The estimator for the total population, $\hat{N}_t = (K/|\mathcal{P}_t|) \cdot n_t$, is directly proportional to our observable metric n_t , as long as the sampling rate $|\mathcal{P}_t|$ remains relatively stable. This assumption is well-founded, as the cyberspace search engine we use maintains a constant scanning bandwidth, leading to a stable probe volume over time. Thus, fluctuations in n_t serve as a reliable proxy for changes in the true activity N_t —in essence, allowing our measurements of the part to reliably represent the whole.

5 *Dynamic Baseline Calculation based on Adaptive Sliding Window.* To detect anomalous drops in network activity, we establish a **Dynamic Baseline** (B_t), representing the expected normal state for any given time interval t . We compute this baseline using a 3-day sliding window that is specifically designed to be adaptive. To obtain the “health” baseline, adaptive means the sliding window automatically excludes any data points previously flagged as part of a shutdown. This prevents past incidents from skewing the baseline and ensures it always reflects a healthy network state. To maintain

a consistent sample size for the calculation, the window backfills with older, non-shutdown data if necessary.

The baseline B_t is then defined as the mean of the most recent L healthy observations:

$$B_t = \text{mean}\{n_i \mid i \in \mathcal{W}_t^*\} \quad (3)$$

where \mathcal{W}_t^* is the set of time indices for the L most recent, non-shutdown data points prior to t . The window length L is set to contain three days of observations (i.e., $L = 432$ for 10-minute intervals).

6 *Shutdown Detection via Baseline Comparison.* With the dynamic baseline B_t established, the next step is to compare the current network activity, n_t , against this baseline to formally detect shutdown events. To quantify the magnitude of such deviations, we introduce the **Service Blocked Ratio (SBR)**. We use the SBR as a critical metric, since our study focuses on a deliberate shutdown event where the primary cause of service unavailability is widely understood to be active blocking or filtering. The metric itself, however, quantifies the overall drop in observed service activity relative to the baseline, regardless of the specific underlying cause. The SBR at time t is calculated as:

$$\text{SBR}_t = \frac{B_t - n_t}{B_t} = 1 - \frac{n_t}{B_t} \quad (4)$$

A shutdown is then flagged at time t if its SBR exceeds a pre-determined critical threshold, θ_{shutdown} . We set this threshold to $\theta_{\text{shutdown}} \approx 10\%$. This value was statistically derived by modeling a historical “healthy” baseline of network activity to identify an extreme lower bound that represents a statistically significant drop. The detailed methodology for this calibration, including the statistical modeling and data validation, is presented in Appendix B.

Shutdown Identification for Iran’s Action. We applied our service-level shutdown monitoring framework to the network service dataset for Iran in June 2025. The process began with an initialization phase using the first three days of data (June 1-3), which served as a verified nominal baseline. This initial data was used to both learn the critical shutdown threshold ($\theta_{\text{shutdown}} \approx 10\%$) and establish the initial state of the 3-day adaptive sliding window. With the framework initialized, we then proceeded to iteratively apply our detection method to the remainder of the month (June 4-30). Finally, our monitor produced a classification for each consecutive 10-minute interval, labeling it as either *Shutdown* or *Normal*.

4 Characterization of Iran Shutdown

In this section, we first reveal the shutdown’s service-level nature and four-phase timeline from a macro-perspective. Subsequently, we conduct a micro-perspective analysis at the AS and service levels, systematically uncovering the Iran shutdown’s progressive and distributed characteristics. Next, we show how existing censorship systems remained active, operating as a complementary layer to the main blockade. Finally, we assess the event’s aftermath, revealing significant collateral impact, including the exposure of network ports and anomalous surges in DNS traffic.

4.1 Overview of Iran’s Shutdown

Finding 1. Multiple data sources confirm that Iran’s June 2025 Internet shutdown was a service-level action.

Service-level Shutdown Confirmation. We analyzed BGP routing data from RIPE RIS [47] and RouteViews [58] in June 2025, particularly during the four-phase shutdown. Analysis from the routing perspective confirms this shutdown was not a network-level action. Throughout June, the number of announced Iranian IPv4 prefixes remained highly stable (44k–46k /24-equivalents), indicating no systematic BGP withdrawals. This strategic shift enhances stealth by evading BGP monitoring systems and enables greater flexibility for rapid service restoration.

4.2 Phased Action

Finding 2. Iran’s Internet shutdown in June 2025 was not a monolithic event but a distinct four-phase action.

Four Phased Action. Based on our monitoring results, we first reveal this shutdown event’s phased nature. As shown in Figure 2, we categorize the shutdown into four phases: two *Localized Shutdown Drills (LSD)*, followed by a *Nationwide Shutdown (NS)* and *Nationwide Shutdown Recovery (NSR)*. Since our service-level monitor operates at a 10-minute temporal resolution, we reconstruct a fine-grained timeline of Iran’s shutdown and cross-validate our findings against multiple third-party monitoring reports.

Phase LSD-1. Our monitor captured the first drop in network activity below the critical threshold at 08:50 on June 13³, marking the beginning of LSD-1. This phase lasted for 2 days, 3 hours, and 40 minutes, during which the maximum drop in network service activity reached 18.64%. This finding is corroborated by Cloudflare Radar [16], which reported an anomaly in user HTTP traffic during nearly the same period (07:15–09:45).

Phase LSD-2. Subsequently, on June 17, several major Iranian operators (e.g., MCI, Irancell) reportedly experienced connection disruptions [43]. We identified another significant drop in activity at 14:00, signaling the start of LSD-2, which continued for 11 hours and 50 minutes with a maximum activity drop of 12.69%. This start time perfectly matches that reported by Cloudflare Radar for shutdown events across multiple ASes.

Phase NS. The nationwide shutdown began on June 18. Our data shows that network activity began with a preliminary decline at 06:30 before dropping sharply around 12:50, ultimately entering the

NS phase. Our identified timeline is largely consistent with the start times for the full-scale shutdown reported by Cloudflare (12:50) and IODA [30](approximately 13:30).

Phase NSR. At 02:00 on June 21, we observed a significant recovery in network activity, a timing consistent with IODA’s report. However, our service-level data further reveals a novel insight: *following this initial recovery, the overall network service activity did not stabilize but instead entered a trend of slow, continuous decline.* This suggests that the shutdown was still being intensified in certain regions or at the service level. We therefore use this time point as a demarcation between two phases of differing shutdown intensity, defining the subsequent period as NSR. Finally, network connectivity was largely restored on June 25.

Finding 3. The two localized shutdowns were a planned technical drill, evidenced by their selective and complementary blocking test of different network services.

Localized Shutdowns as Drills. Our unique service-level perspective reveals a critical pattern: the blocking strategies in the first two localized shutdowns (LSD-1 and LSD-2) were highly selective and targeted different service categories. We hypothesize that they were deliberate drills, rather than indiscriminate outages. An analysis of Iran’s top 20 most-used network service ports provides evidence for this hypothesis. During LSD-1, only 6 of the top 20 ports experienced a Service Blocked Ratio (SBR) over 10%. The restrictions were heavily concentrated on 5 foundational and circumvention services, including DNS (port 53, 37%), SNMP (port 161, 35%), NTP (port 123, 26%), IKEv2 (port 500, 34%), and L2TP (port 1701, 34%). The strategy then evolved in LSD-2, showcasing a clear complementarity. The blockade expanded to 12 of the top 20 ports, with the focus shifting decisively to the application layer. This phase impacted 6 services essential for user applications and databases, such as RDP (port 3389, 25%), FTP (port 21, 18%), MySQL (port 3306, 20%), and core web services on ports 80 (11%) and 443 (13%). Collectively, the two drills impacted a total of 15 unique ports within the top 20 with an SBR of over 10%, systematically testing the state’s blocking capabilities against the vast majority of Iran’s critical network services (see Appendix C for a detailed breakdown).

4.3 Progressive Action

Finding 4. The shutdown exhibited a progressive rollout at the AS-level, with both its escalation and subsequent restoration unfolding in distinct, gradual phases.

AS Level Strategy. To analyze the shutdown’s evolution across networks, we focused on the top 100 ASes in Iran (covering 93.67% of observed services). As illustrated in Figure 3, we classified the state of each AS in each phase based on its Service Blocked Ratio (SBR). An AS was generally labeled *Affected* for an SBR over 10%; to capture recovery nuances in the NSR phase, we introduced a *Partially Recovered* category (SBR 10–30%) and raised the *Affected* threshold to above 30%. Notably, Arvancloud (AS205585) operates its CDN service using an Anycast deployment with over 40 global Points of Presence (PoPs) [5], and was thus unaffected.

³All times in this paper are denoted in Coordinated Universal Time (UTC).

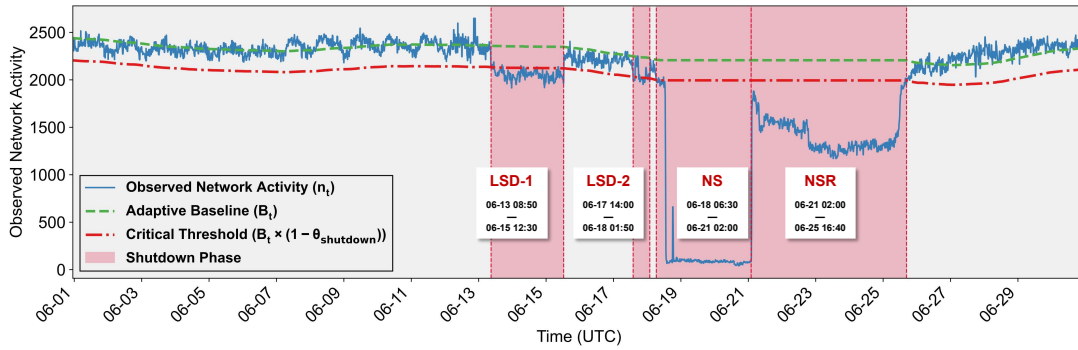


Figure 2: Iran’s Four-Phase Shutdown Timeline: Two Localized Drills, a Nationwide Shutdown and Subsequent Recovery.

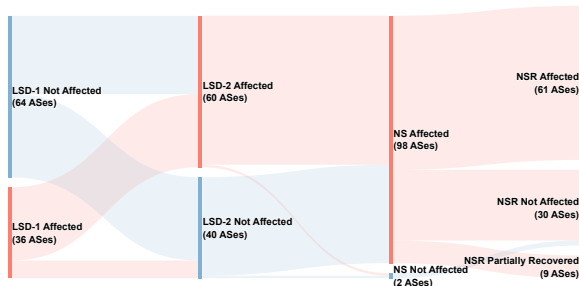


Figure 3: State Transition of the Top 100 ASes (by Network Service Volume) During Four Shutdown Phases.

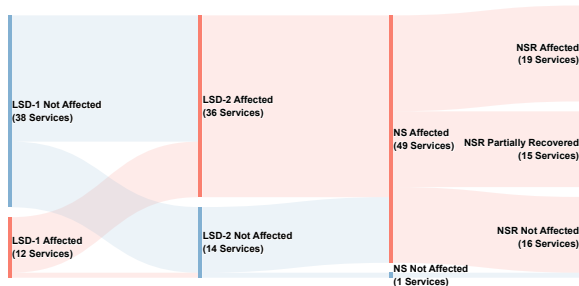


Figure 4: State Transition of the Top 50 Network Services (by Network Service Volume) During Four Shutdown Phases.

The results reveal a clear pattern of *progressive escalation*. The number of affected ASes grew systematically from 36 in LSD-1 to 60 in LSD-2, culminating in the near-total blockade of 98 ASes during NS. This progressive nature was mirrored in the non-uniform recovery during the NSR phase, where only 30 ASes returned to normal, 9 were partially restored, and a majority of 61 remained under the full effects of the blockade.

Finding 5. The shutdown’s service-level strategy evolved through four phases: from infrastructure control, to information obstruction, to a total blockade, and finally to a censorship-oriented recovery.

Service Level Strategy. To investigate the granularity of the service-level blocking strategy, we analyzed Iran’s top 50 network services by volume (e.g., HTTP, DNS, RIP, IKEv2), which collectively represent 99.77% of all services. Following a methodology similar to our AS-level analysis, we classified each service as *Affected*, *Not Affected*, or *Partially Recovered* based on its SBR (Service Blocked Ratio) during each shutdown phase.

The results reveal a clear progressive pattern in both the service-level shutdown’s escalation and its subsequent recovery. As shown in Figure 4, the number of affected services grew systematically from 12 in LSD-1 to 36 in LSD-2, peaking at 49 in NS. GPRS Tunneling Protocol (GTP)—essential for mobile core networks—remained largely unaffected. This progressive pattern was mirrored in the non-uniform recovery during the NSR phase, where services diverged into three distinct groups: 16 fully recovered, 15 partially restored, and the largest group, 19, still blocked. This tiered restoration strategy confirms that a fine-grained, progressive control logic governed the entire event, from initiation to conclusion.

Service Level Evolution. Our service-level analysis reveals a clear strategic evolution through four distinct phases: an initial targeted disruption of core infrastructure and tunneling services; a comprehensive escalation to the application layer; a total blockade with strategic exemptions; and a final, tiered, and censorship-oriented recovery. To demonstrate this, we categorized the top 50 services in Iran by function into 12 classes, such as *Web Access* and *Core Infrastructure Services*. As shown in Figure 5, we then quantified the impact on each category across the shutdown phases using *Service Blocked Ratio* (SBR).

LSD-1: Targeted Disruption of Infrastructure and Tunneling Services. This initial phase was highly precise, targeting *Core Infrastructure Services* (e.g., DNS[38, 59], 37%; NTP[37, 48], 26%) and *Tunneling Services* for network connectivity (L2TP, 34%; IKEv2, 34%). By disabling domain resolution, time synchronization, and communication tunnels, this phase achieved the primary objective of establishing comprehensive infrastructure control.

LSD-2: Escalation to the Application Layer. Building on these initial disruptions, the strategy escalated sharply as the number of affected services surged from 12 to 36. The blockade expanded from infrastructure to the application layer, impacting *Web and Content Access* (HTTP, 14%; HTTPS, 11%), *Email Services* (IMAP, 30%; POP3, 32%), and *File Transfer and Sharing Services* (FTP, 26%; SMB, 31%). This




Figure 5: Service Blocked Ratio of Top 50 Network Services During Four Shutdown Phases.

signaled a definitive shift in control objectives from infrastructure control to the obstruction of information flow.

NS: Total Blockade with Strategic Exemption for Mobile Infrastructure. During this phase, nearly all mainstream services (49/50) were paralyzed, with blocking ratios universally exceeding 85%. The sole exception was the GPRS Tunneling Protocol (GTP). This exemption was a key strategic decision that preserved the operational capacity of the internal mobile core network, likely for essential management and specific state communications.

NSR: Tiered and Censorship-Oriented Recovery. The recovery was a strategically controlled, tiered operation with two clear objectives. First, the restoration was tiered by service function: *Core Infrastructure Services* (DNS, 6%; NTP, 22%) recovered quickly, ostensibly to restore basic connectivity. In stark contrast, services central to information dissemination, such as email and databases, remained heavily suppressed (SBRs of 50-62%). This tiered approach demonstrates a continued strategic distinction, prioritizing the restoration of core infrastructure while still maintaining the *obstruction of information flow*. Second, the recovery was censorship-oriented. The restoration of unencrypted HTTP (17%) was far more advanced than that of encrypted HTTPS (35%). This prioritizes traffic that is easier to inspect—a finding corroborated by our censorship analysis in Section 4.5. This differentiated recovery reveals a clear strategy to reshape the post-shutdown information environment by maximizing censorship capabilities.

4.4 Distributed Deployment

Finding 6. The significant cross-AS heterogeneity in the timing of the shutdown’s initiation and recovery reveals its reliance on a distributed enforcement system.

Distributed Nature of Shutdown Enforcement. We conduct a more detailed analysis of the actions taken by different ASes during the shutdown event, as shown in Figure 6. First, we found that the initiation of the blockade was highly asynchronous. The heatmap clearly shows that the LSD-1 phase selectively impacted specific providers like Respina (AS42337) and Fanap (AS206065), and the transition to a full blockade was staggered by hours, and even days, across different ASes. Second, the NSR recovery was a non-uniform and dynamic process. The heatmap reveals a complex mosaic of recovery patterns, including “recover-then-re-block” cycles (e.g., AS48147), which are indicative of dynamic, independent policy adjustments. Taken together, these patterns point to a sophisticated hybrid control architecture. Taken together, these patterns point to a sophisticated hybrid control architecture underpinned by a distributed enforcement layer.

From these observations, we argue that while the highly synchronous onset of the NS phase suggests a centralized command capable of issuing nationwide directives, this macro-level view belies a more complex reality at the enforcement level. It is the cross-AS heterogeneity in the blockade’s application and removal that reveals the system’s more critical feature: its distributed nature. Reports [28, 43] indicate Iranian censorship infrastructure relies on a hybrid distributed deployment of Deep Packet Inspection (DPI) boxes—some state-controlled at large ISPs, others independently installed by various providers. Such a deployment strategy is consistent with Iran’s active pursuit of a tiered Internet structure [42] through its National Information Network (NIN) initiative.

4.5 Complementary Measure

Finding 7. Censorship persisted during the shutdown, operating as a complementary layer to the service-level shutdown to form a defense-in-depth system.

Censorship-based blocking. We analyzed public data from the Open Observatory of Network Interference (OONI) [54] to investigate censorship during the shutdown. OONI’s volunteer-run probes offered a unique internal perspective on censorship dynamics amidst the blockade. However, the service-level shutdown severely impacted OONI’s measurement capabilities. OONI’s daily test volume in Iran, which totaled 121,333 in June 2025, plummeted to under 200 during the shutdown’s peak (June 19-20), as macro-level policies prevented most probes from operating. Despite the reduced overall volume, censorship intensity on the “surviving” traffic increased. OONI’s measured anomaly rate, typically 40-60% pre-shutdown, consistently rose to 50-75% during the shutdown phases. This demonstrates that the censorship system remained active, catching residual traffic that bypassed the broader blockade. Censorship was therefore not replaced by the shutdown but acted as its critical supplement, creating a multi-layered, defense-in-depth control system (see Appendix D for a visualization of these trends).

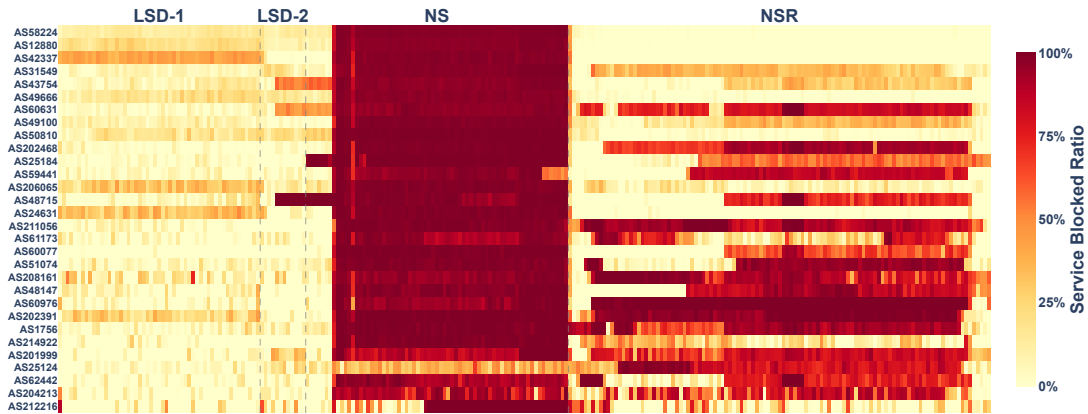


Figure 6: Top 30 AS's Service Blocked Ratio During Four Shutdown Phases.

4.6 Potential Collateral Impact

Finding 8. Iran's Internet shutdown in June 2025 caused unexpected collateral impacts, such as anomalous port exposure and traffic surges.

Port Exposure. We compared the network service activity during the three days post-shutdown (June 26-28) against a three-day baseline pre-shutdown (June 10-12). Our analysis reveals that the hasty recovery process led to the anomalous exposure of numerous ports. The most alarming trend was a sharp increase in active SSH services (port 22), which surged by 22% from 148,222 to 181,077. These new services were primarily concentrated in Iranian cloud and hosting providers (e.g., AbrArvan, AS202468) and the national telecom operator (TCI, AS58224). We argue that this surge in port 22 exposure creates a dual risk: revealing critical servers and expanding the attack surface for brute-force attacks. Furthermore, we observed high growth in the activity of web administration panels (e.g., cPanel/WHM on ports 2082/2083) and common web development ports (e.g., 8080, 3000). Over 90% of the new services on ports 8080 and 3000 were plaintext HTTP.

DNS Traffic Surge. We analyzed June data from a large public DNS recursive operator. The dataset covered two dimensions: 19.6 billion queries originating from 311 Iranian ASes, and 780 million queries from global ASes for Iran's country-code top-level domain (ccTLD), .ir. During the shutdown, DNS query traffic experienced an anomalous surge. Queries originating from Iran increased by an average of 36%, with a peak of 115%, while global queries for .ir domains peaked at an 87% increase. We hypothesize that the increase in DNS query volume may result from two factors: emerging DNS queries caused by network service switching and reconfiguration due to the shutdown, and users' frequent attempts to send network requests.

5 Lessons Learned

Evolution of Internet Shutdown Strategy. The June 2025 shutdown in Iran demonstrates a clear evolution in shutdown strategies from coarse-grained, network-level control towards fine-grained, service-level control. Our findings reveal the state's capability to

selectively impose blockades on specific ASes or service categories. This was exemplified during the LSD-1 phase, where the blockade impacted only 36 of the top 100 ASes and was precisely targeted at core infrastructure and tunneling services, showcasing a more sophisticated and granular control capability.

Towards Service-Level Aware Monitoring. Existing shutdown monitoring techniques, from *passive-traffic-based* to *network-level active probing*, lack the comprehensive service-level perspective needed to capture the strategic evolution of blocking policies seen in the Iran shutdown. We therefore argue that future shutdown monitoring systems should integrate more diverse service-level signals to effectively detect and characterize such events. Our *service-level active probing* methodology establishes a framework for this approach, enabling a fine-grained characterization of the June 2025 Iran Internet shutdown. To foster community research, we have released our aggregated dataset of Iranian network services used in this study to help build long-term monitoring capabilities for maintaining network availability and stability.

6 Conclusion

Understanding the mechanisms of modern Internet shutdowns is fundamental to ensuring global Internet resilience. In this paper, we develop and apply a service-level shutdown monitoring framework that analyzes network service data from active probing. Using this framework, we perform the first fine-grained analysis of the June 2025 Iran shutdown. Our findings reveal the action was not a monolithic blackout but a sophisticated operation defined by three core properties: phased, progressive, and distributed. Our study, along with the public network service dataset we provide, encourages further exploration by the Internet community into the mechanisms of such advanced shutdowns and the development of more adaptive Internet shutdown monitoring techniques.

Acknowledgments

We thank all anonymous reviewers for their valuable and constructive feedback. This work is supported in part by the National Key Research and Development Program of China (No. 2023YFB3105600), Zhongguancun Laboratory and the National Natural Science Foundation of China (62102218).

References

- [1] Access Now. 2025. *Lives on hold: internet shutdowns in 2024*. Technical Report. Access Now.
- [2] Giuseppe Aceto, Alessio Botta, Pietro Marchetta, Valerio Persico, and Antonio Pescapé. 2018. A comprehensive survey on internet outages. *Journal of Network and Computer Applications* 113 (2018), 36–63.
- [3] Leon Adato. 2024. *Using Kentik NMS to Identify Network Outages*. Retrieved October 5, 2025 from <https://www.kentik.com/blog/using-kentik-network-monitoring-system-to-identify-network-outages/>
- [4] Abdulrahman Alaraj and Eric Wustrow. 2025. Proxies as Sensors: Measuring Censorship of Refraction Networking in Iran. In *Proceedings of the 20th ACM Asia Conference on Computer and Communications Security (ASIA CCS '25)*. Association for Computing Machinery, New York, NY, USA, 759–772. doi:10.1145/3708821.3733879
- [5] Arvancloud. 2025. *Content Delivery Network (CDN), Fast and Secure Content Delivery*. Arvancloud. Retrieved October 6, 2025 from <https://www.arvancloud.ir/en/products/cdn>
- [6] Simurgh Aryan, Homa Aryan, and J. Alex Halderman. 2013. Internet Censorship in Iran: A First Look. In *3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI 13)*. USENIX Association, Washington, D.C. <https://www.usenix.org/conference/foci13/workshop-program/presentation/aryan>
- [7] Arash Aryapour. 2025. Architectural and Traffic Analysis of the June 2025 Iran Internet Shutdown. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5528119
- [8] Arash Aryapour. 2025. Iran's Stealth Internet Blackout: A New Model of Censorship. arXiv:2507.14183 [cs.NI] <https://arxiv.org/abs/2507.14183>
- [9] Guillermo Baltra and John Heidemann. 2020. Improving coverage of internet outage detection in sparse blocks. In *International Conference on Passive and Active Network Measurement*. Springer, 19–36.
- [10] Karyn Benson, Alberto Dainotti, KC Claffy, and Emile Aben. 2012. Gaining insight into as-level outages through analysis of internet background radiation. In *Proceedings of the 2012 ACM conference on CoNEXT student workshop*. 63–64.
- [11] Kevin Bock, Yair Fax, Kyle Reese, Jasraj Singh, and Dave Levin. 2020. Detecting and evading {Censorship-in-Depth}: A case study of {Iran's} protocol whitelister. In *10th USENIX Workshop on Free and Open Communications on the Internet (FOCI 20)*.
- [12] Matthew Caesar, Lakshminarayanan Subramanian, and Randy H Katz. 2003. *Towards localizing root causes of BGP dynamics*. Computer Science Division, University of California.
- [13] Calipr Networking Group. [n. d.]. *ICLab*. CICS UMass Amherst. Retrieved October 5, 2025 from <https://iclab.gitlab.io/>
- [14] Kelvin Chan and Barbara Ortutay. 2025. *Iran asks its people to delete WhatsApp*. AP News. Retrieved October 4, 2025 from <https://apnews.com/article/iran-whatsapp-meta-israel-d9e6fe43280123c9963802e6f10ac8d1>
- [15] David R Choffnes, Fabián E Bustamante, and Zihui Ge. 2010. Crowdsourcing service-level network event monitoring. In *Proceedings of the ACM SIGCOMM 2010 Conference*. 387–398.
- [16] Cloudflare. 2025. *About Cloudflare Radar*. Cloudflare. Retrieved October 4, 2025 from <https://radar.cloudflare.com/about>
- [17] William Gemmill Cochran. 1977. *Sampling techniques*. John Wiley & Sons.
- [18] Cody Combs. 2025. *Iran internet disrupted as country imposes restrictions*. Retrieved September 15, 2025 from <https://www.thenationalnews.com/future/technology/2025/06/13/iran-internet-outage-israel/>
- [19] Alberto Dainotti, Claudio Squarcella, Emile Aben, Kimberly C Claffy, Marco Chiesa, Michele Russo, and Antonio Pescapé. 2011. Analysis of country-wide internet outages caused by censorship. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. 1–18.
- [20] Ronald Deibert, John Palfrey, Rafal Rohozinski, and Jonathan Zittrain. 2008. *Access denied: the practice and policy of global Internet filtering*. The MIT Press.
- [21] Shivani Deshpande, Marina Thottan, Tin Kam Ho, and Biplab Sikdar. 2009. An online mechanism for BGP instability detection and analysis. *IEEE transactions on Computers* 58, 11 (2009), 1470–1484.
- [22] Freedom House. 2023. *Freedom on the Net 2023: Iran Country Report*. Technical Report. Freedom House. <https://freedomhouse.org/country/iran/freedom-net/2023>
- [23] Kristin Glass, Richard Colbaugh, and Max Planck. 2010. Automatically identifying the sources of large Internet events. In *2010 IEEE International Conference on Intelligence and Security Informatics*. IEEE, 108–113.
- [24] Robert David Graham. 2014. *Masscan: Mass IP port scanner*. GitHub. <https://github.com/robertdavidgraham/masscan>
- [25] Deborah Grey. 2024. *Fibre optic cables vandalised in France*. W.Media. Retrieved October 5, 2025 from <https://w.media/fibre-optic-cables-vandalised-in-france/>
- [26] Andreas Guillot, Romain Fontugne, Philipp Winter, Pascal Merindol, Alistair King, Alberto Dainotti, and Cristel Pelsler. 2019. Chocolate: Outage detection for internet background radiation. In *2019 Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 1–8.
- [27] Nguyen Phong Hoang, Sadie Doreen, and Michalis Polychronakis. 2019. Measuring {I2P} censorship at a global scale. In *9th USENIX Workshop on Free and Open Communications on the Internet (FOCI 19)*.
- [28] Freedom House. 2024. *Iran: Freedom on the Net 2024 Country Report*. <https://freedomhouse.org/country/iran/freedom-net/2024>
- [29] Internet Assigned Numbers Authority (IANA). 2025. *Service Name and Transport Protocol Port Number Registry*. <https://www.iana.org/assignments/service-name-port-numbers/service-names-port-numbers.xhtml>
- [30] Internet Outage Detection and Analysis (IODA). 2025. *IODA: Internet Outage Detection and Analysis*. Retrieved October 5, 2025 from <https://ioda.inetintel.cc.gatech.edu/>
- [31] Internet Society. 2019. *Iran | 16 November 2019 - 24 November 2019*. Internet Society Pulse. Retrieved October 4, 2025 from <https://pulse.internetsociety.org/en/shutdowns/shutdown-194/>
- [32] Internet Society Pulse. [n. d.]. *About*. Internet Society Pulse. Retrieved October 5, 2025 from <https://pulse.internetsociety.org/about>
- [33] IRANWIRE. 2025. *IRAN IMPOSES INTERNET RESTRICTIONS AFTER ISRAELI STRIKES*. Retrieved September 15, 2025 from <https://iranwire.com/en/news/142043-iran-imposes-internet-restrictions-after-israeli-strikes/>
- [34] Zubair Nabi. 2013. The anatomy of web censorship in Pakistan. In *3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI 13)*.
- [35] NDTV News Desk. 2025. *Internet Cut For 48 Hours In UP's Bareilly Amid 'I Love Muhammad' Posters Row*. NDTV. Retrieved October 5, 2025 from <https://www.ndtv.com/india-news/internet-cut-for-48-hours-in-ups-bareilly-amid-i-love-muhammad-posters-row-9384013>
- [36] NetBlocks. 2025. *About NetBlocks*. Retrieved October 4, 2025 from <https://netblocks.org/about>
- [37] Gunter Ollmann. 2016. *NTP: The Most Neglected Core Internet Protocol*. Retrieved January 23, 2026 from https://circleid.com/posts/20161205_ntp_the_most_neglected_core_internet_protocol
- [38] Zakiyaameen Pachapuri. 2024. *DNS security: Fortifying the core of Internet infrastructure*. Retrieved January 23, 2026 from <https://www.catchpoint.com/blog/dns-security-fortifying-the-core-of-internet-infrastructure>
- [39] Ramakrishna Padmanabhan and Alberto Dainotti. 2021. Internet Outages: How much of a problem are they? CAIDA. <https://www.caida.org/workshops/wombir/2101/slides/wombir2021-paper37.pdf>
- [40] Ramakrishna Padmanabhan, Alberto Dainotti, Nima Fatemi, Arturo Filastò, Maria Xynou, and Simone Basso. 2019. *Iran's nation-wide Internet blackout: Measurement data and technical observations*. OONI. Retrieved October 4, 2025 from <https://ooni.org/post/2019-iran-internet-blackout/>
- [41] Ramakrishna Padmanabhan, Arturo Filastò, Maria Xynou, Ram Sundara Raman, Kennedy Middleton, Mingwei Zhang, Doug Madory, Molly Roberts, and Alberto Dainotti. 2021. A multi-perspective view of Internet censorship in Myanmar. In *Proceedings of the ACM SIGCOMM 2021 Workshop on Free and Open Communications on the Internet*. 27–36.
- [42] Pana. 2023. An Internet for the Few: Iran's Digital Segregation Plans Edge Closer to Reality. *filter.watch* (24 May 2023). <https://filter.watch/2023/05/24/an-internet-for-the-few-irans-digital-segregation-plans-edge-closer-to-reality/>
- [43] Ania M. Piotrowska. 2025. *Nym report on Iran's recent Internet blackouts (June 2025)*. Retrieved October 4, 2025 from <https://nym.com/blog/nym-iran-report-2025>
- [44] Lin Quan, John Heidemann, and Yuri Pradkin. 2013. Trinocular: Understanding internet reliability through adaptive probing. *ACM SIGCOMM Computer Communication Review* 43, 4 (2013), 255–266.
- [45] Reethika Ramesh, Ram Sundara Raman, Matthew Bernhard, Victor Ongkowijaya, Leonid Evdokimov, Anne Edmundson, Steven Sprecher, Muhammad Ikram, and Roya Ensafi. 2020. Decentralized control: A case study of russia. In *Network and Distributed Systems Security (NDSS) Symposium 2020*.
- [46] Philipp Richter, Ramakrishna Padmanabhan, Neil Spring, Arthur Berger, and David Clark. 2018. Advancing the art of internet edge outage detection. In *Proceedings of the Internet Measurement Conference 2018*. 350–363.
- [47] RIPE Network Coordination Centre (RIPE NCC). 2001–2025. *Routing Information Service (RIS)*. <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/>
- [48] Teemu Ryttilahti, Dennis Tatang, Janosch Köpper, and Thorsten Holz. 2018. Masters of time: An overview of the NTP ecosystem. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 122–136.
- [49] Dominik Schatzmann, Simon Leinen, Jochen Kögel, and Wolfgang Mühlbauer. 2011. Fact: Flow-based approach for connectivity tracking. In *International Conference on Passive and Active Network Measurement*. Springer, 214–223.
- [50] Nehad Selaiha. 2013. The fire and the frying pan: Censorship and performance in egypt. *TDR* (2013), 20–47.
- [51] Anant Shah, Romain Fontugne, Emile Aben, Cristel Pelsler, and Randy Bush. 2017. Disco: Fast, good, and cheap outage detection. In *2017 Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 1–9.
- [52] Jonas Tai, Karthik Nishanth Sengottuvelavan, Peter Whiting, and Nguyen Phong Hoang. 2025. {IRBlock}: A {Large-Scale} Measurement Study of the Great Firewall of Iran. In *34th USENIX Security Symposium (USENIX Security 25)*. 705–722.

- [53] Soon Tee Teoh, Supranamaya Ranjan, Antonio Nucci, and Chen-Nee Chuah. 2006. BGP eye: a new visualization tool for real-time detection and analysis of BGP anomalies. In *Proceedings of the 3rd International Workshop on Visualization for Computer Security (Alexandria, Virginia, USA) (VizSEC '06)*. Association for Computing Machinery, New York, NY, USA, 81–90. doi:10.1145/1179576.1179593
- [54] The OONI Project. 2025. *OOONI Data: Real-time and historical data on internet censorship around the world*. Retrieved October 4, 2025 from <https://explorer.ooni.org/>
- [55] ThousandEyes. [n. d.]. *ThousandEyes Internet Outages*. ThousandEyes. Retrieved October 5, 2025 from <https://www.thousandeyes.com/outages/>
- [56] Elisa Tsai, Ram Sundara Raman, Atul Prakash, and Roya Ensafi. 2024. Modeling and Detecting Internet Censorship Events. In *Network and Distributed System Security Symposium (NDSS)*.
- [57] University of Michigan. 2018. *Censored Planet*. University of Michigan. Retrieved October 5, 2025 from <https://censoredplanet.org/#/about>
- [58] University of Oregon. 1995–2025. Route Views Project. <https://www.routeviews.org/>.
- [59] Roland van Rijswijk-Deij, Mattijs Jonker, Anna Sperotto, and Aiko Pras. 2015. The Internet van Names: A DNS Big Dataset. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (London, United Kingdom) (SIGCOMM '15)*. Association for Computing Machinery, New York, NY, USA, 91–92. doi:10.1145/2785956.2789996
- [60] Hilary Whiteman. 2025. *A two-day internet blackout shows even the Taliban can't turn back time*. CNN. Retrieved October 5, 2025 from <https://www.cnn.com/2025/10/03/asia/afghanistan-internet-shutdown-intl-hnk-dst>
- [61] Diven Xue, Reethika Ramesh, Valdik S S, Leonid Evdokimov, Andrey Viktorov, Arham Jain, Eric Wustrow, Simone Basso, and Roya Ensafi. 2021. Throttling Twitter: an emerging censorship technique in Russia. In *Proceedings of the 21st ACM internet measurement conference*. 435–443.
- [62] Maria Xynou and Arturo Filastò. 2022. *New blocks emerge in Russia amid war in Ukraine: An OONI network measurement analysis*. OONI. Retrieved October 4, 2025 from <https://ooni.org/post/2022-russia-blocks-amid-ru-ua-conflict/>
- [63] Tarun Kumar Yadav, Akshat Sinha, Devashish Gosain, Piyush Kumar Sharma, and Sambuddho Chakravarty. 2018. Where the light gets in: Analyzing web censorship mechanisms in india. In *Proceedings of the Internet Measurement Conference 2018*. 252–264.
- [64] Kim Zetter. 2014. *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon*. Crown, New York.
- [65] Mingwei Zhang, Jun Li, and Scott Brooks. 2017. I-Seismograph: Observing, Measuring, and Analyzing Internet Earthquakes. *IEEE/ACM Transactions on Networking* 25, 6 (2017), 3411–3426. doi:10.1109/TNET.2017.2748902

A Ethics and Limitations

A.1 Ethics Considerations.

Our service-level monitoring of network shutdowns relies on data from the routine, periodic scans of a commercial cyberspace search engine. The scanning platform itself employs established harm-reduction techniques, such as randomized task scheduling, to strictly limit its impact on target networks to a controllable range. The scanning platform's vantage points are all hosted in cloud data centers, with the provider's informed consent to conduct its measurement activities. For transparency and accountability with external networks, the platform is identifiable via PTR records on these vantage points and a custom User-Agent in its web probes. Both methods provide project details and an opt-out mechanism. Notably, we did not commission any new or targeted scans for this research. Consequently, our work introduced no additional measurement traffic to Iran's network. Furthermore, our supplementary datasets are sourced from established platforms such as OONI [54], RIPE RIS [47], and RouteViews [58]. These organizations have extensive experience, their own ethics review processes, and publish data that does not contain sensitive or personally identifiable information. Similarly, the dataset we will publish with this study has been aggregated to the autonomous system level. This process removes all specific IP addresses, ensuring that our released data does not contain any personally identifiable information. While some impacts may still exist, our study enables a fine-grained understanding

of a major network shutdown event and supports the monitoring and mitigation of shutdown effects. Thus, we believe its benefits outweigh the potential ethical risks. Our service-level monitoring of network shutdowns relies on data from the routine, periodic scans of a commercial cyberspace search engine. The scanning platform itself employs established harm-reduction techniques, such as randomized task scheduling, to strictly limit its impact on target networks to a controllable range.

A.2 Limitations.

Our study characterizes the connectivity between Iran and the global Internet. This external perspective means our measurements could be influenced by factors independent of the Iran shutdown, such as outages or censorship targeting our vantage points. To mitigate this potential confound, we strategically select measurement nodes in countries with robust network infrastructure and high Internet freedom, such as Japan and Singapore. During our study period in June 2025, we observed no such outages or censorship affecting our vantage points.

B Calibration of the Shutdown Threshold

This appendix details the statistical process for calibrating the critical shutdown threshold, θ_{shutdown} . The core logic is to define abnormal states using the extreme lower bound of normal (i.e., non-shutdown) network conditions.

To establish a model for normal network behavior, we first selected a period of n_t data from June 1-3, 2025, which was verified to be a "healthy" baseline free of anomalies by cross-referencing public reports. The verification sources include IODA, Cloudflare Radar, and Internet Society Pulse [32]. The time series exhibited significant over-dispersion (variance greater than the mean), likely due to correlations between network services. We therefore fitted a Negative Binomial distribution, a standard approach for modeling such count data.

The model provided an excellent goodness-of-fit to the data (p-value = 0.89), as shown in Figure 7. From this validated model, we define the threshold based on an extreme lower bound of normal activity, specifically the 0.01st percentile ($Q_{0.0001}$). The critical drop ratio, θ_{shutdown} , is the proportional drop from the historical mean (μ_{hist}) to this percentile:

$$\theta_{\text{shutdown}} = 1 - \frac{Q_{0.0001}}{\mu_{\text{hist}}} = 1 - \frac{2119}{2344} = 9.61\% \approx 10\% \quad (5)$$

C Detailed Port-Level Data for Localized Shutdowns

Figure 8 is a detailed visualization of the Service Blocked Ratio for the top 20 network service ports during the two localized shutdown drills (LSD-1 and LSD-2), providing the primary evidence for their distinct and complementary targeting strategies as discussed in the main text. The figure illustrates the different targeting strategies of the two phases. LSD-1 (blue) primarily impacted infrastructure and VPN ports (e.g., 53, 161, 1701), while LSD-2 (red) shifted focus to user-facing application ports (e.g., 80, 443, 3389), highlighting the complementary and deliberate nature of the drills.

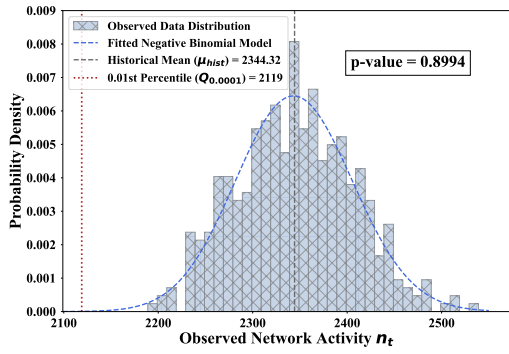


Figure 7: Distribution of n_t during a verified non-shutdown period, used for threshold calibration. The histogram of observed data is overlaid with the fitted Negative Binomial curve.

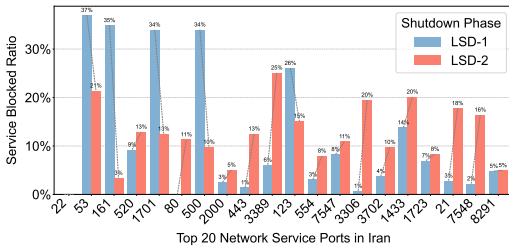


Figure 8: Service Blocked Ratio for the top 20 network service ports during the two localized shutdown drills (LSD-1 and LSD-2).

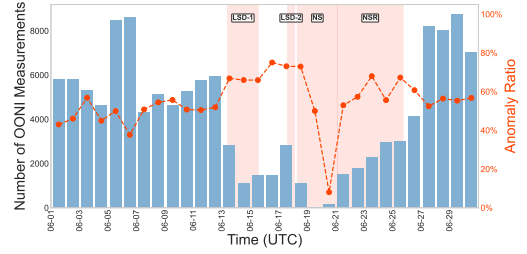


Figure 9: OONI Measurement Volume (blue bars, left y-axis) and Anomaly Ratio (orange line, right y-axis) in Iran for June 2025.

D Detailed OONI Measurement Data and Anomaly Analysis

We analyzed OONI test data from Iran in June 2025, which covers 121,333 queries across 33 ASes. This section provides the detailed data visualization supporting our analysis of censorship during the shutdown, as discussed in the main text. Figure 9 illustrates the concurrent trends of plummeting OONI measurement volume and rising anomaly rates during the shutdown phases, which provides the core evidence for Finding 6. The figure demonstrates that as the shutdown intensified, particularly during NS, the number of successful measurements plummeted. Concurrently, the anomaly rate for the surviving traffic remained consistently high or increased, supporting the conclusion that censorship mechanisms were still active. The rate for June 19-20 is not statistically representative due to low sample size, but the overall trend holds.