Quantifying DNS Resilience

A Statistical Framework for Global Internet Infrastructure

Hümay Zehra Özer

ICANN86 Nextgen

The Invisible Spine of the Internet

DNS is the invisible backbone of our digital world. It often goes unnoticed when working seamlessly, but its disruption brings digital life to a standstill.

Critical Dependency

Modern systems rely entirely on DNS. However, current tools only monitor failures after they happen (Reactive).

Strategic Mission

ICANN ensures a secure and stable internet. To do this, we need to see hidden, latent risks before a crash occurs.

The Research Gap

How can we use statistical models to predict future failure risks and engineer proactive resilience?

Global Consequences

The Case of the Dyn DNS Outage

A single point of technical failure can cause massive ripple effects across the entire globe.

Service Loss

Major platforms became inaccessible. This proved that we cannot just watch the system; we must estimate risks early.

Economic Impact

Billions were lost. Reactive monitoring only reports the loss, but it cannot prevent the economic damage.

Interconnectivity

Failures spread instantly. We need a statistical risk model to see how a small local error becomes a global crisis.

Beyond Lookups: The Ecosystem Chain

Root Servers

The authoritative starting point for every resolution process.

TLD Servers

Registry operators directing queries to specific domains like .com or .org.

Recursive Resolvers

The "librarians" of the web, finding and catching the final answers.

“If one link in this chain fails, the whole internet suffers. We need a statistical way to predict which link is weak.”

Engineering-Inspired DNS View

Complex systems must be built on measurable engineering pillars to ensure longevity and stability.

Reliability

The probability that a component performs its required function. We can calculate this using statistical data.

Fault Tolerance

The system's power to survive a crash. Our model will predict this survivability.

Redundancy

Having backup server nodes. We need statistics to find the optimum number of backups.

Designing for Failure: k-out-of-n

DNS is architected to avoid "Single Points of Failure" (SPOF) through a multi-layered redundancy model.

N

The total number of available server nodes. Our model analyzes if this number is enough for heavy traffic.

K

The minimum active nodes needed to keep the system operational. If active nodes fall below k, the system crashes.

R

The reliability of a node. We use real data to calculate this risk dynamically.

The Logic: The system remains operational as long as *k* out of *n* nodes are active. This provides a buffer against simultaneous outages and surges.

The Statistical Risk Model

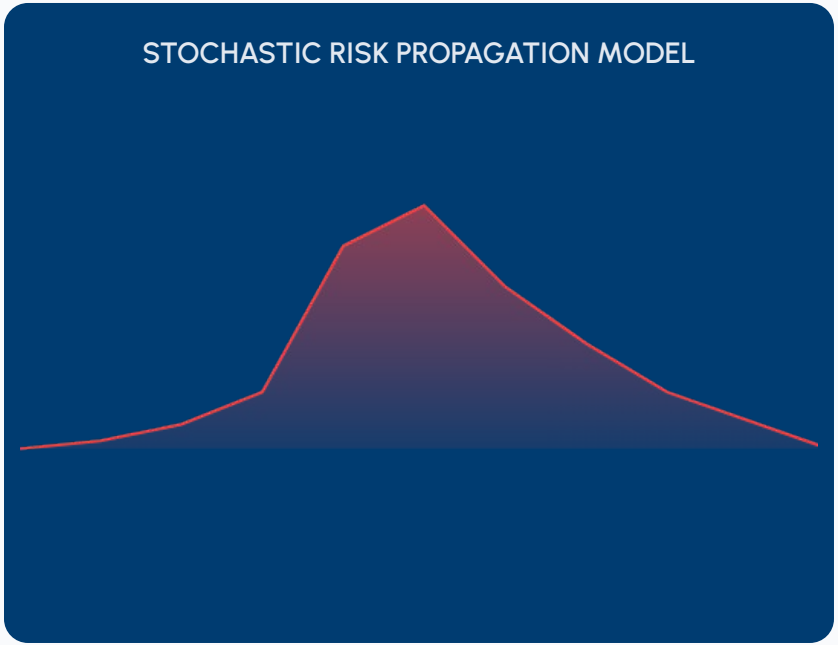

Predictive Power: We use stochastic processes to calculate estimate how a small local failure spreads through the network.

Dynamic Estimation: Instead of just measuring uptime, this model gives ICANN a proactive "Resilience Score" before a crisis happens.

From Monitoring to Prediction

Feature	Current Monitoring Focus	Predictive Resilience Perspective
Temporal Focus	What is happening now?	What could happen next?
Operational Nature	Reactive (Fixing Errors)	Predictive (Managing Risk)
Analytical Method	Observation & Telemetry	Statistical Resilience Estimation
Outcome Goal	Uptime Maintenance	Resilience Engineering

From Measurement to Governance

Policy Support

Data-Driven Policies: Transforming complex statistical results into clear, objective facts for ICANN regulations.

SSR Goals

Proactive SSR Targets: Helping ICANN move from dynamic data monitoring to long-term resilience engineering.

Stakeholders

Community Action: Empowering operators and policy makers with simple, actionable risk metrics.

Integrating Existing ICANN Tools

RSSAC: Monitoring the Root Server system's performance and stability indicators.

DAAR: Domain Abuse Activity Reporting to analyze patterns in malicious DNS use.

ITHI: Identifier Technologies Health Indicators for ecosystem wellness.

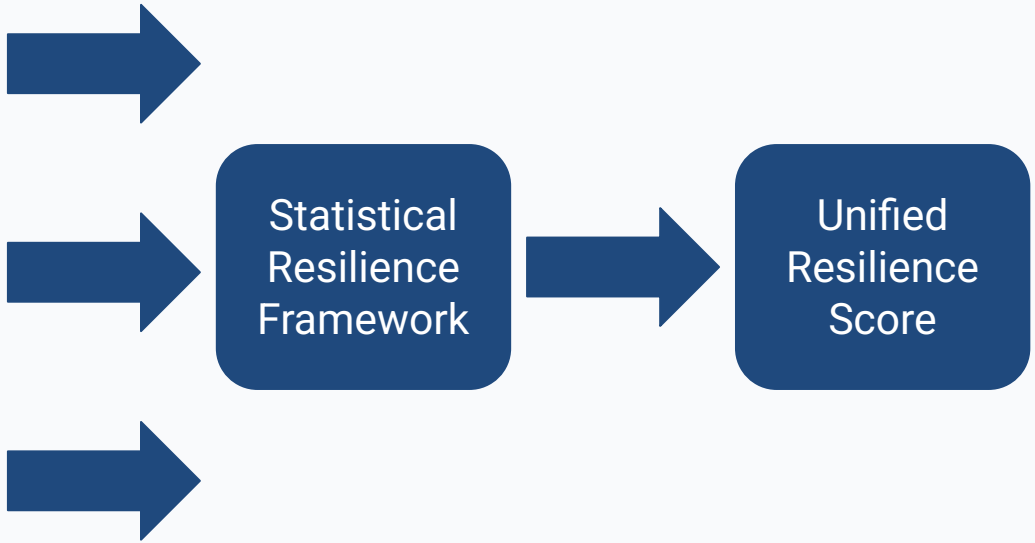

Our framework serves as the "analytical glue" that connects these isolated data points into a unified resilience score.

Why Resilience Metrics Matter

Operators

Capacity planning and load analysis to stay robust during spikes.

Security Teams

Identifying weak spots and potential "failure zones" before attacks occur.

Policy Makers

Making informed, risk-based decisions for global internet governance.

A Resilient Future

"We cannot manage what we do not measure — and we cannot protect what we do not understand."

Thank You!

Questions?

huemay.oezer@students.unibe.ch | ICANN86 Nextgen

backup slide

Mathematical Foundation

Quantifying system-wide reliability through the k-out-of-n probability distribution.

$$R_{sys} = \sum_{i=k}^{n} \binom{n}{i} r^i (1-r)^{n-i}$$

n: Total available root server nodes.

k: Minimum operational threshold for stability.

r: Probability of a single node performing successfully.

backup slide

Assumptions & Boundaries

Defining the scope and technical constraints of the current statistical framework.

Statistical Independence

Initial modeling assumes component failures are independent events to baseline structural risks.

Data Sync Frequency

Model accuracy is dynamically constrained by the polling intervals of RSSAC and ITHI source APIs.

Weight Uniformity

Starting parameters treat node vulnerabilities as uniform before environmental stress factors are applied.

backup slide

Strategic Roadmap

Scaling the predictive framework into a production-ready ecosystem for ICANN governance.

Phase 1: ML Integration

Transitioning from static probability to dynamic AI models for real-time node degradation alerts.

Phase 2: Live Dashboard

A high-fidelity interface providing the ICANN community with a centralized "Resilience Score."

Phase 3: Stress Tests

Introducing empirical parameters to evaluate infrastructure against simulated massive traffic anomalies.