

SCION: The Trusted Connectivity Foundation for AI and the Post-Quantum World

Adrian Perrig
Network Security Group

ETH zürich

SCION

Exciting Opportunities of a Next-Generation Secure Internet



For AI

SCION + Hummingbird offer high-quality connectivity for demanding AI use cases



Against AI attacks

AI-based attacks require fundamentally secure defenses

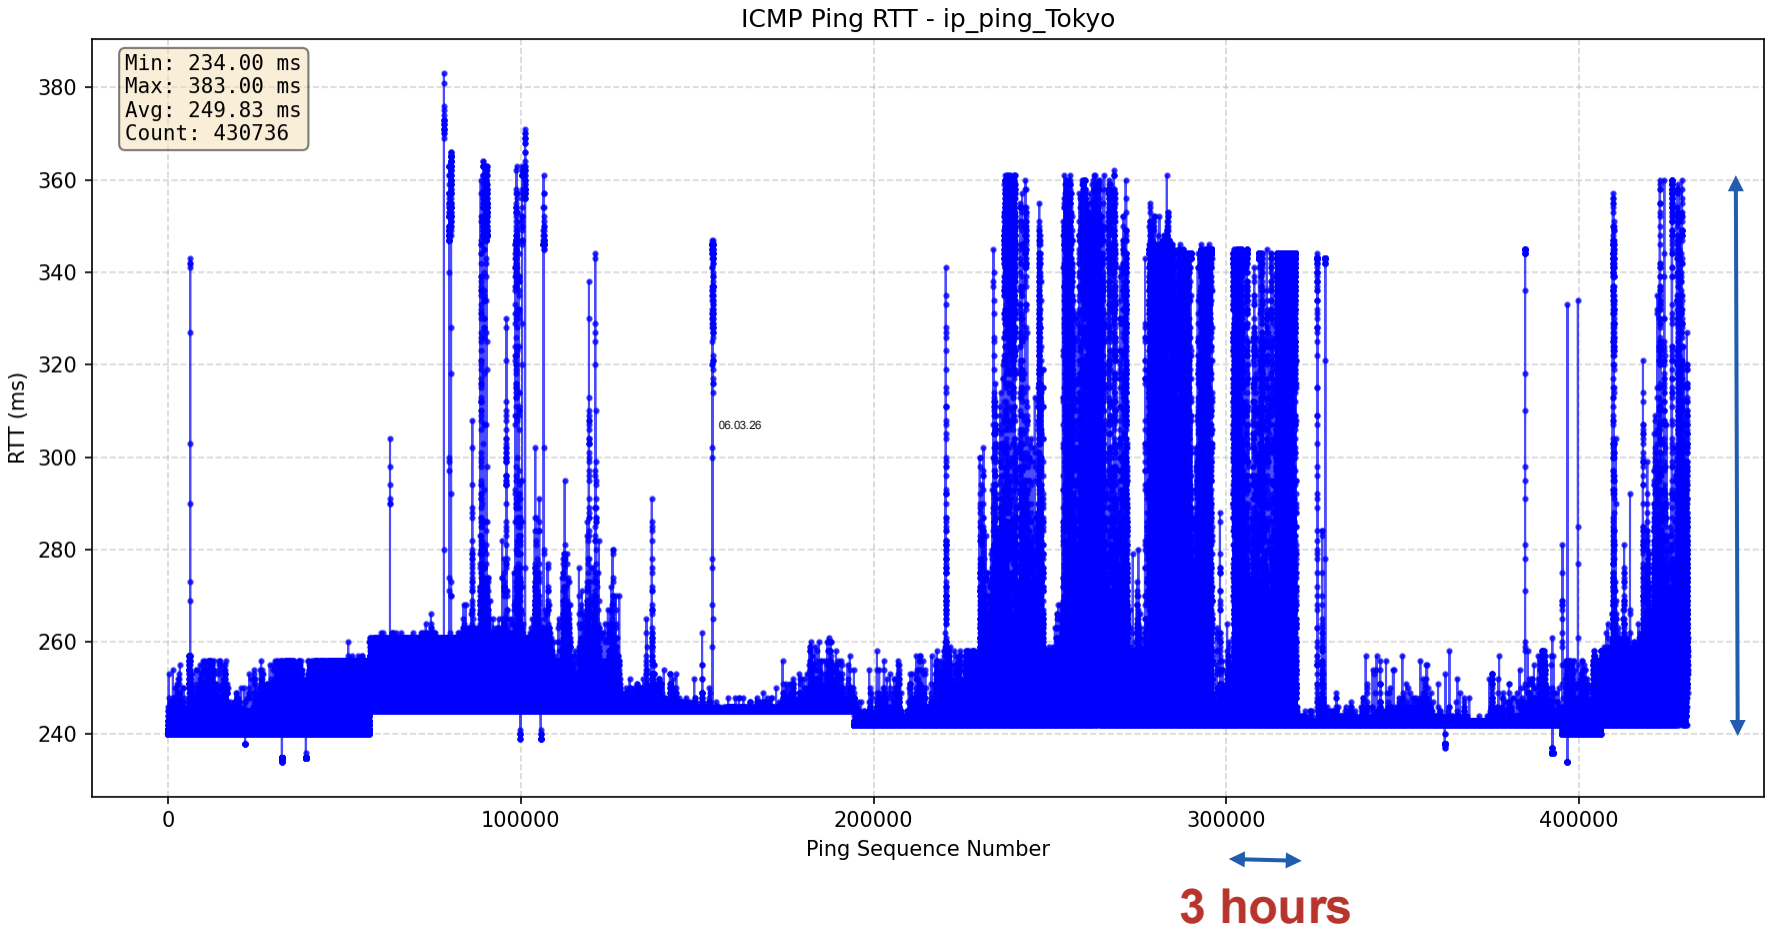


Post Quantum

Rapid progress of quantum computers requires a quantum-resilient network infrastructure

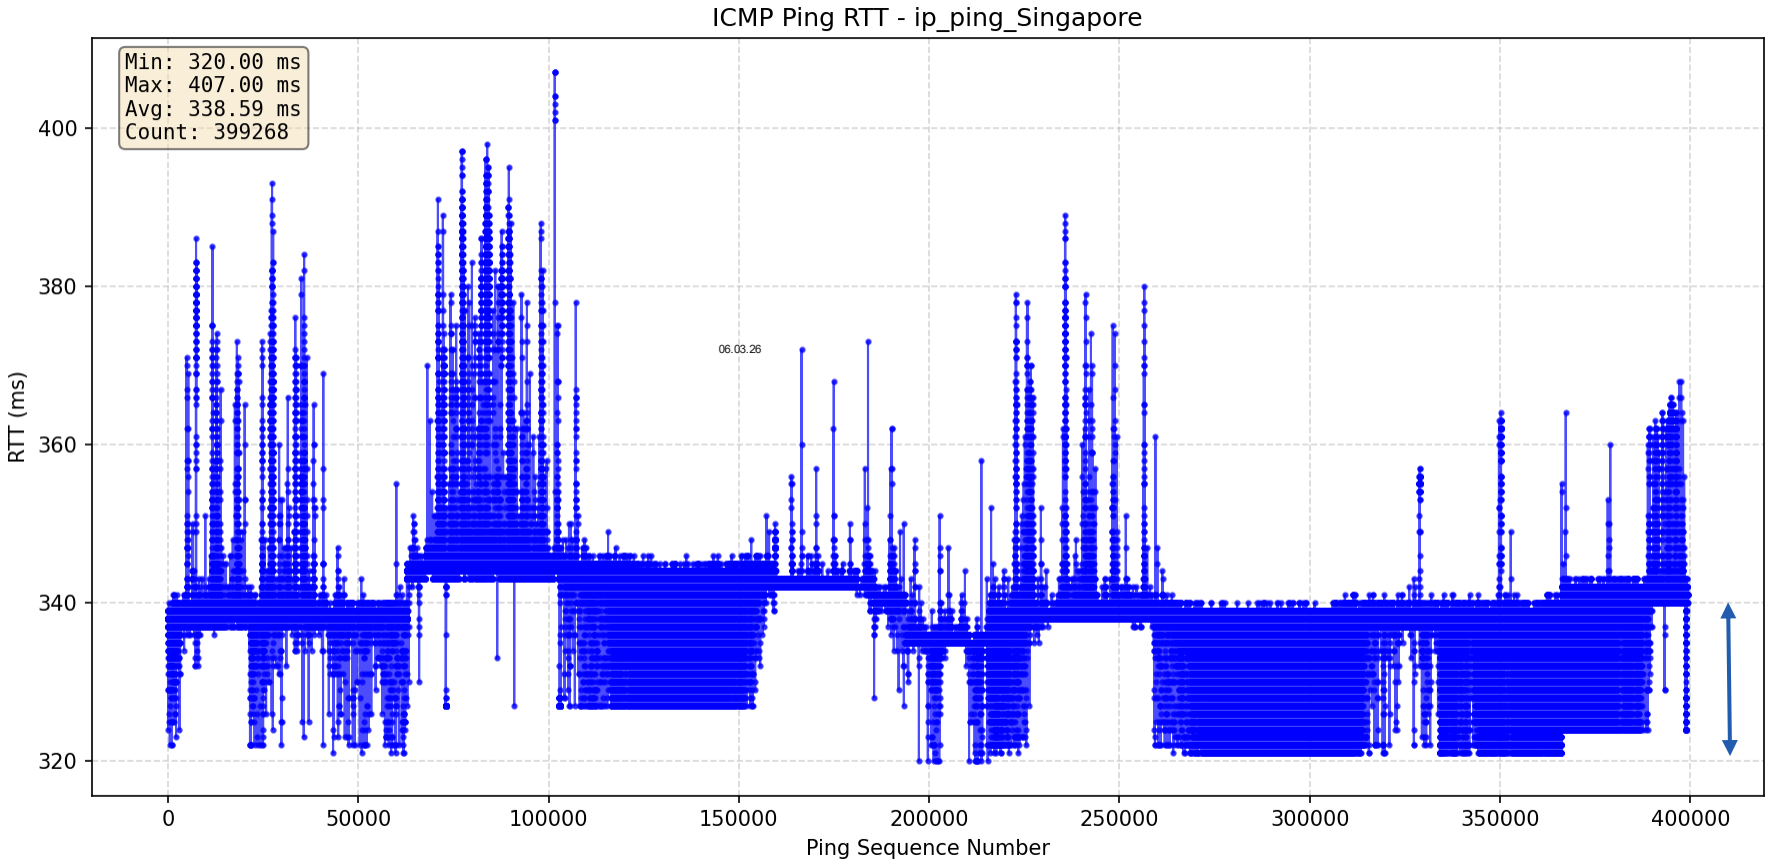
Network performance can be highly variable over time

Examp



Network performance can be highly variable over time

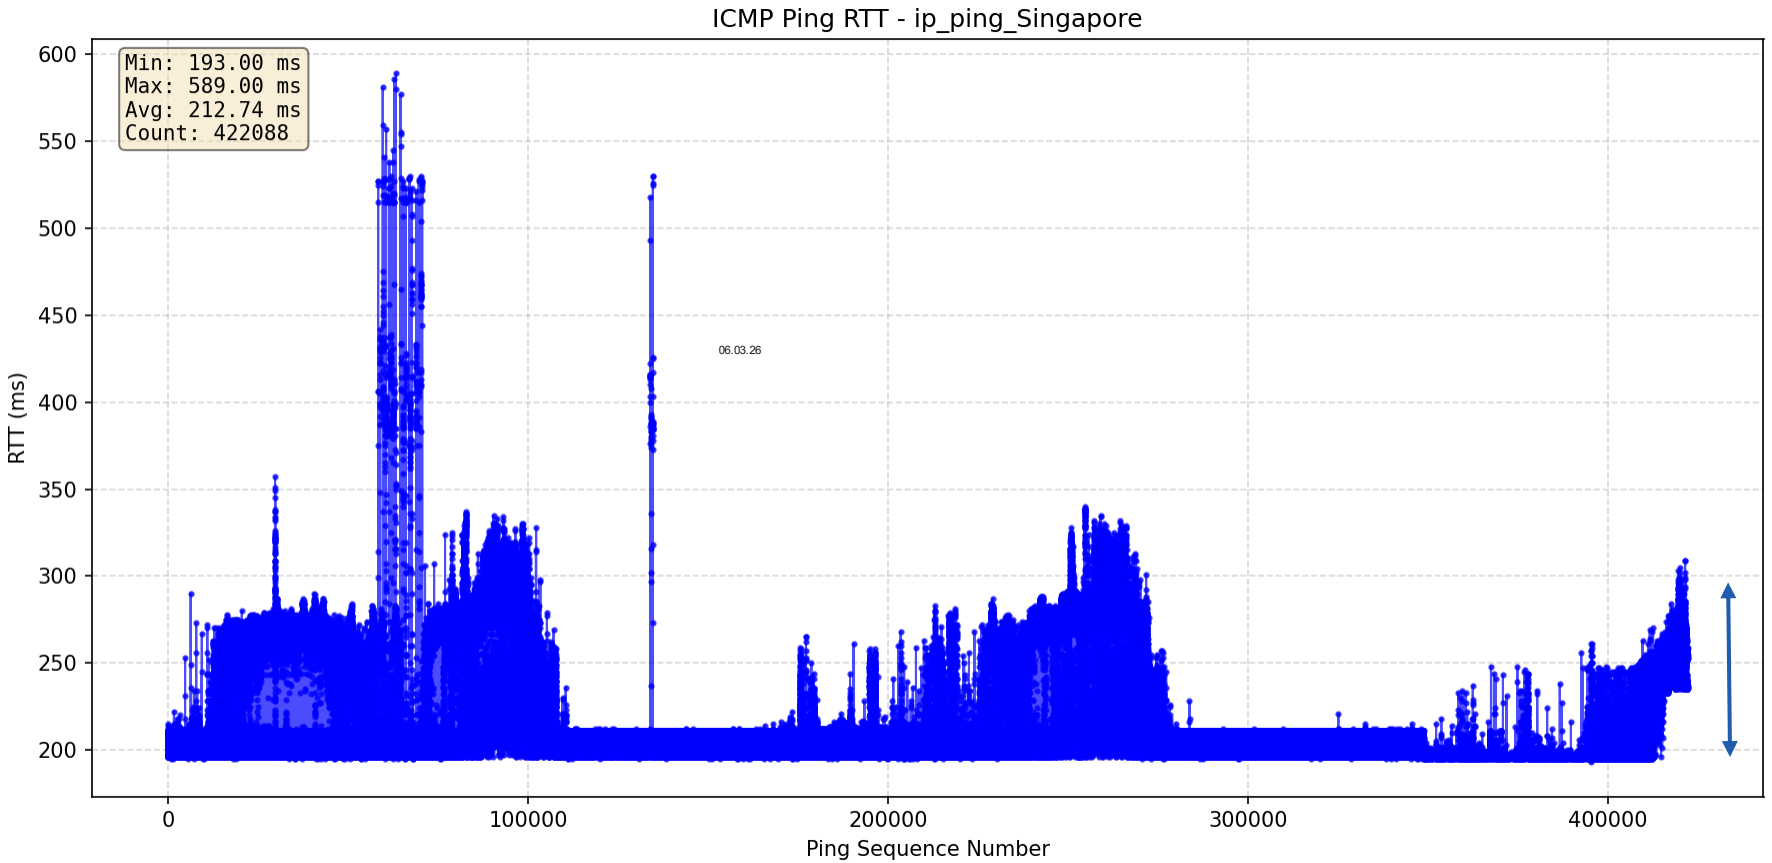
Example: OVH (Roubaix) – Vultr (Singapore)



6% latency increase

Network performance can be highly variable over time

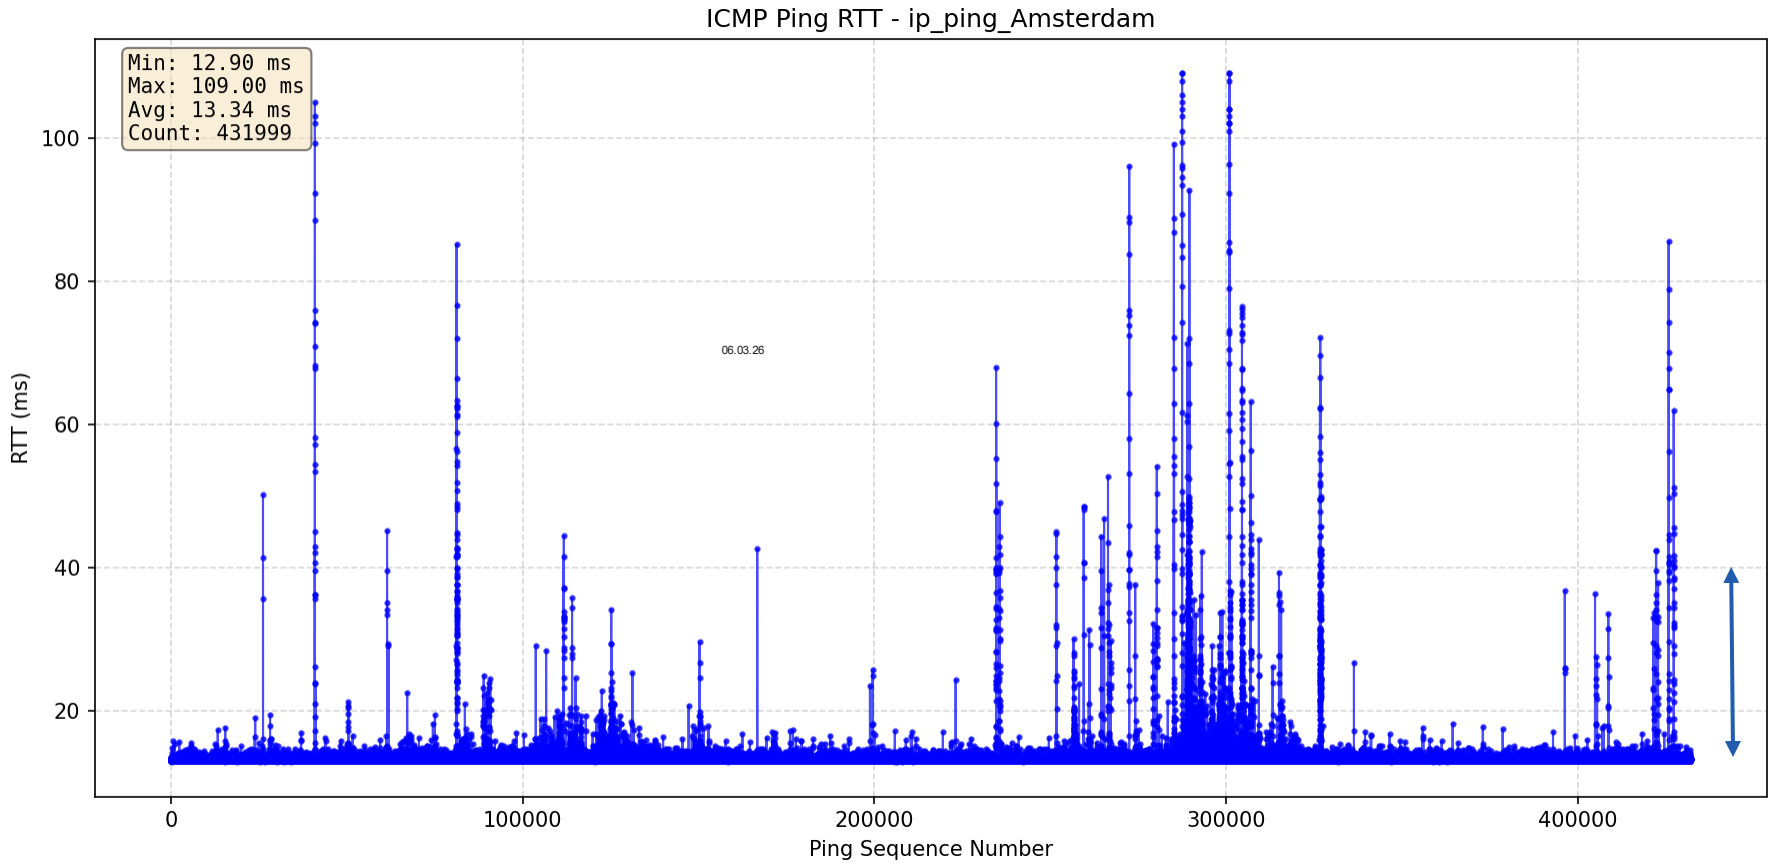
Example: WebNX (Salt Lake City) – Vultr (Singapore)



50% latency increase

Network performance can be highly variable over time

Example: Cloudscale (Zurich) – Worldstream (Amsterdam)



175% latency increase

Evidence from 129 days of RIPE Atlas 1.5B Measurements

Latency between 105 data centres, 31 May to 6 Oct 2026

Every Pair Among 105 Data Centres, for 129 Days

The data

105
data-center servers (RIPE Atlas anchors)

8
geographic regions

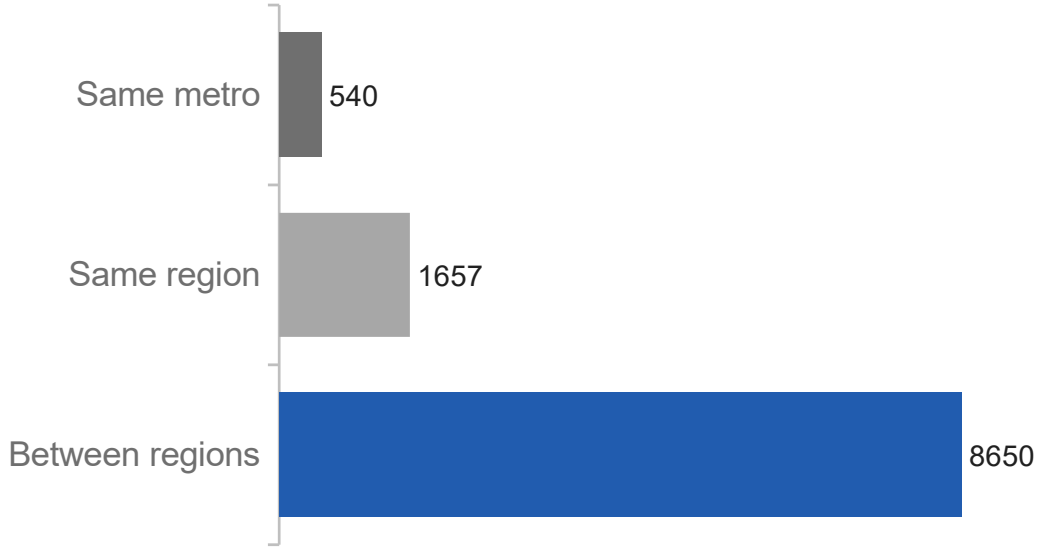
10,847
measured pairs, both ways (A pings B, B pings A)

1.48 bn
round-trip times recorded

How it was measured

- Every RIPE Atlas anchor pings every other anchor every 4 minutes, 3 packets each time.
- 31 May to 6 October 2026: 129 full days.
- Anchors mostly deployed in similar environments as validators.
- Pairs are grouped by distance: same metro, same region, between regions.

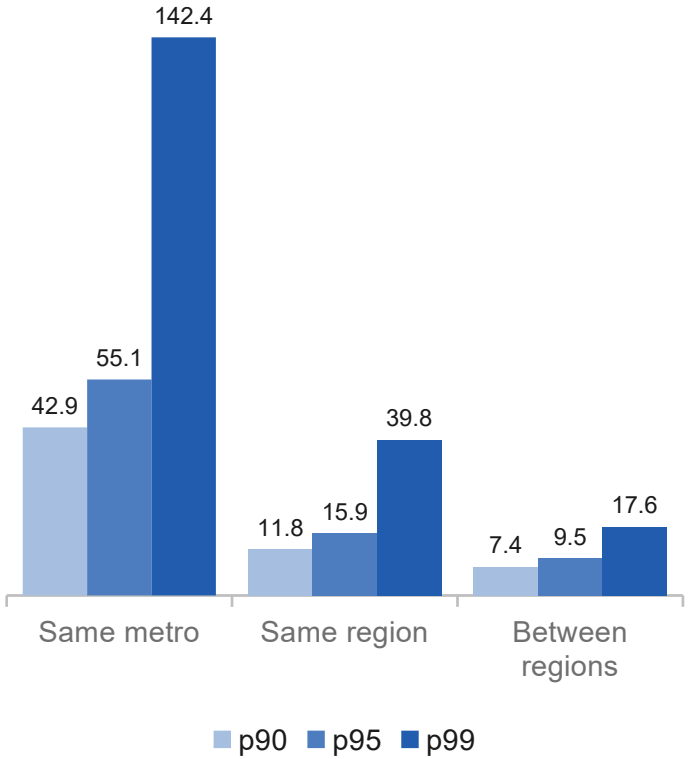
Measured pairs by distance



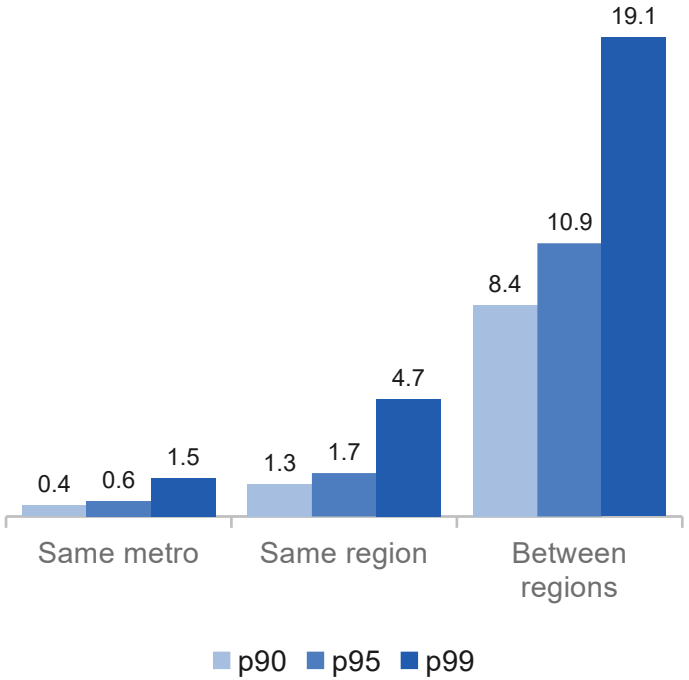
The “Unlucky” Cases Are Far Worse Than the “Lucky” Ones

Finding 1: spread

Latency above p10 over 129 days (%), median over pairs



Above p10 over 129 days (ms), median over pairs



17.6 %

p99 latency inflation compared to p10 for inter-region RTTs

73 ms

p99 above the median for the worst 10 % of pairs between regions (median pair: 16 ms).

Same-metro RTTs are about 1.1 ms, so 0.44 ms of jitter reads as 43 %.

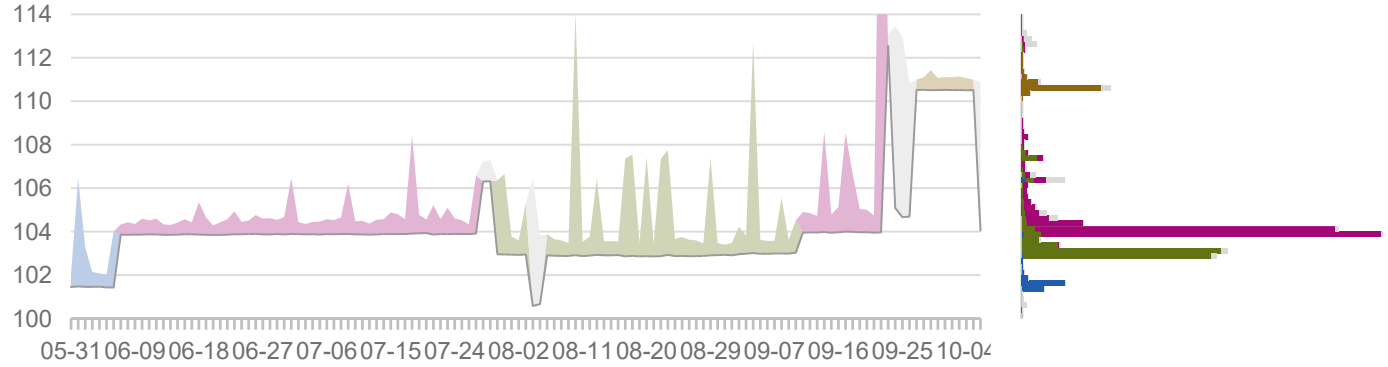
The Long-Term Distribution Is Built From Levels Held Over Time

Finding 2: baseline levels

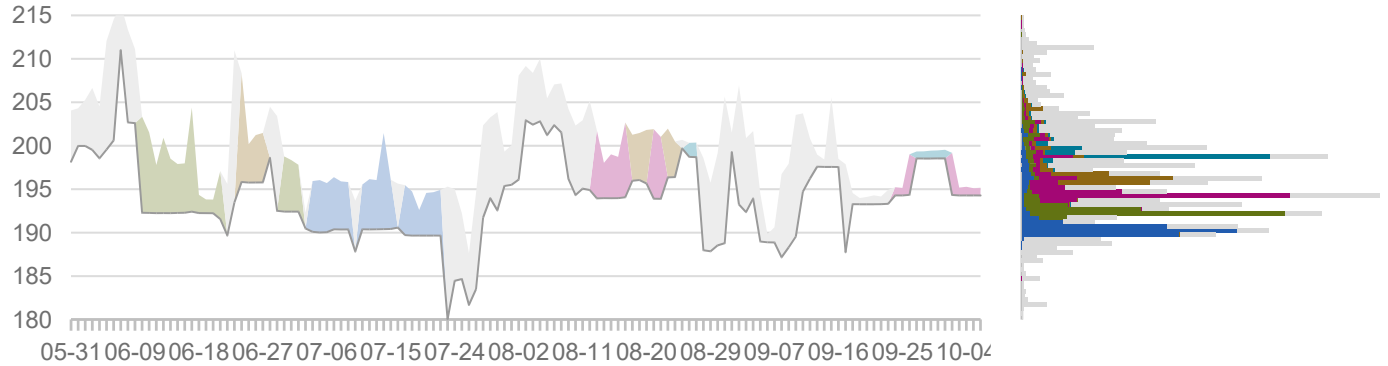
Daily p10 (line) and p10–p90 (shading); one colour per level, grey = other days

All packets pooled

Frankfurt → Atlanta: levels clearly apart: low p10 variance, levels clearly distinguishable



Dubai → New York: levels blend together: many levels only 1.6–2.7 ms apart, smaller than many daily spreads (about 5 ms)



Terminology: Level

A level is 2 or more days of consistent p10 latencies (color-coded in the left graphs)

Levels clearly apart: every gap between levels is larger than a day's spread (p90 − p10), so the pooled distribution has separate peaks.

Levels blend together: the gaps are smaller than a day's spread, so the peaks merge into one hump.

The public Internet cannot guarantee the performance critical applications need



Unpredictable bandwidth

Throughput on shared public paths changes with other traffic; no provider guarantees it end to end.



Latency, jitter and loss

Delay, jitter, loss and downtime raise cost and uncertainty for anything that runs in real time.

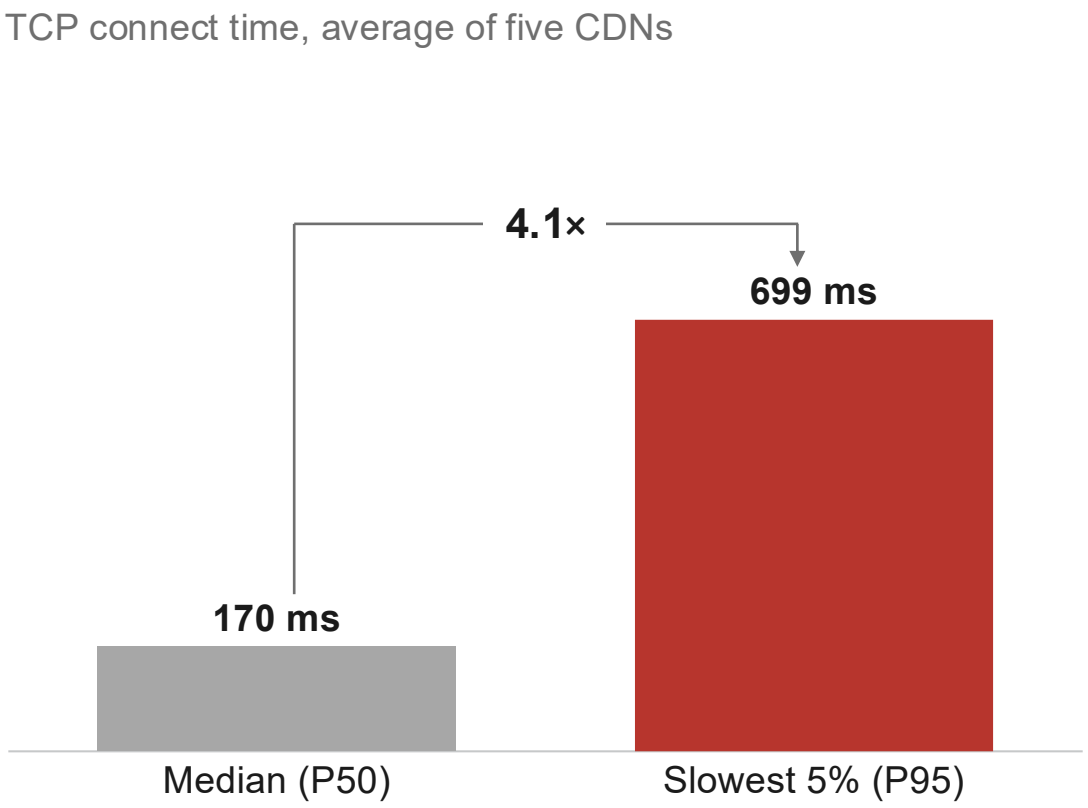


Heavy-tailed path quality

The worst 10% of a single path is very poor. With several WAN paths across entities, the chance that one is poor is high.

On the “Heavy Tail” Nature of Internet Connectivity

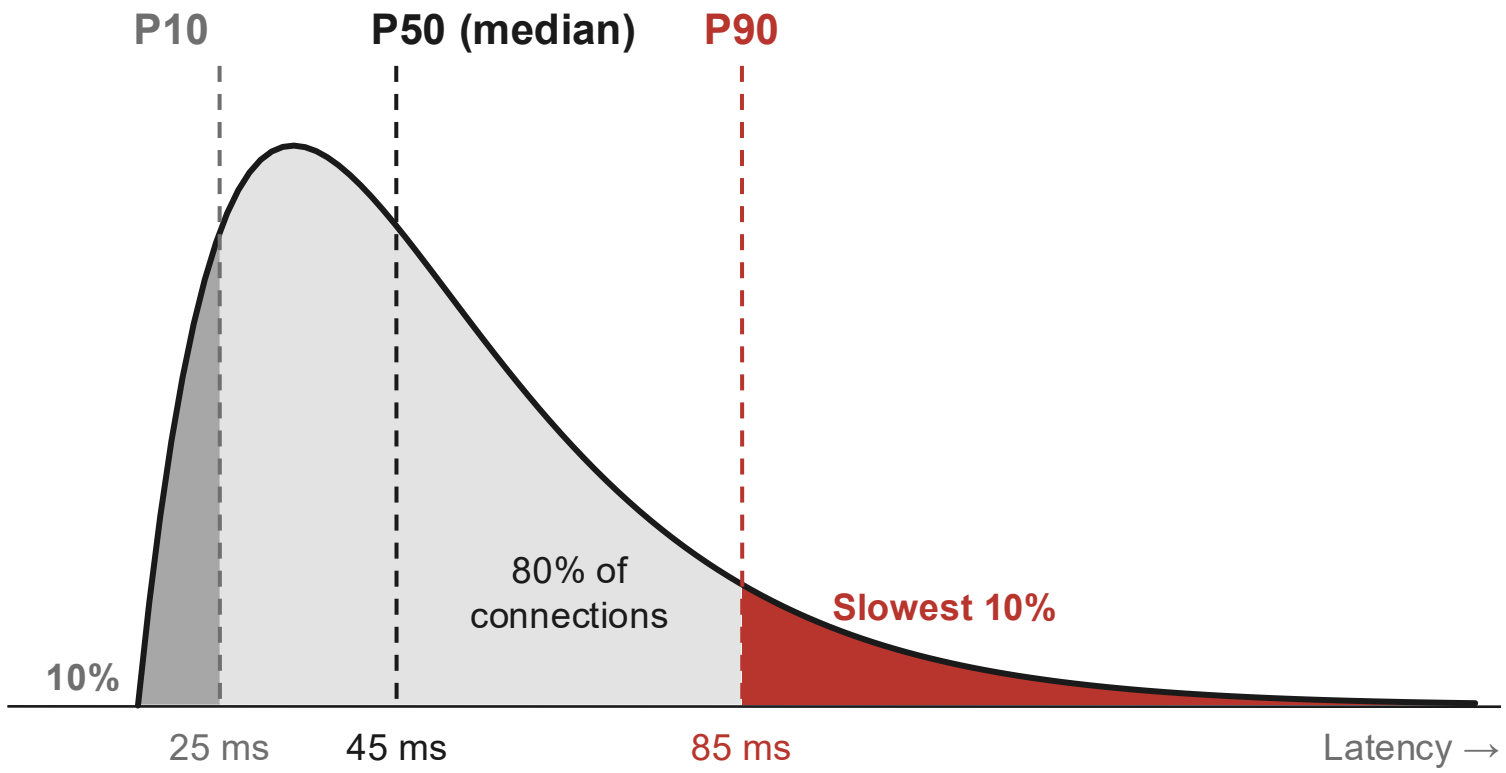
- Good performance on average, but the **worst or “unlucky” connections** see much higher latency and packet loss
- Median TCP connect time: **~170 ms** across five major CDNs
- Slowest 5% of connections: **~700 ms, 4× the median**
- The gap holds at every CDN: **3.9× to 4.5×**



P50 and P95 TCP connect times for Cloudflare, Amazon CloudFront, Google, Akamai and Fastly, Nov 2023–Feb 2024. Source: Cloudflare, “Network performance update: Security Week 2024”

Percentiles Show How Slow the Worst Connections are - Not Just the Typical One

Latency of individual connections (illustrative gamma distribution)



- **P10:** 10% of connections are “lucky” and faster than this point.
- **P50 (median):** half are faster, half slower.
- **P90:** 90% are faster; the slowest 10% are slower. The “unlucky” cases constitute the tail.
- **Heavy tail:** “unluckiest” cases are much slower

Latency-sensitive Applications in Many Industries Suffer from Best-effort Networks



AI and compute

Distributed AI training and inference

- Precise, reliable time synchronization



Finance and ledgers

Trading and transaction processing

- Blockchain consensus



Healthcare and emergency

Remote medicine and telesurgery

- Emergency response systems



Media and communication

4K video calling and conferencing

- 4K video streaming



Immersive and gaming

VR/AR with low-latency haptic feedback

- Video gaming



Critical infrastructure

Grids, pipelines, rail, defense
Industrial control, remote robots

- Guaranteed government services

How can SCION and Hummingbird Improve Connectivity Quality?

Path Selection and Multipath are Powerful to Improve Connectivity Quality



Path quality changes over time



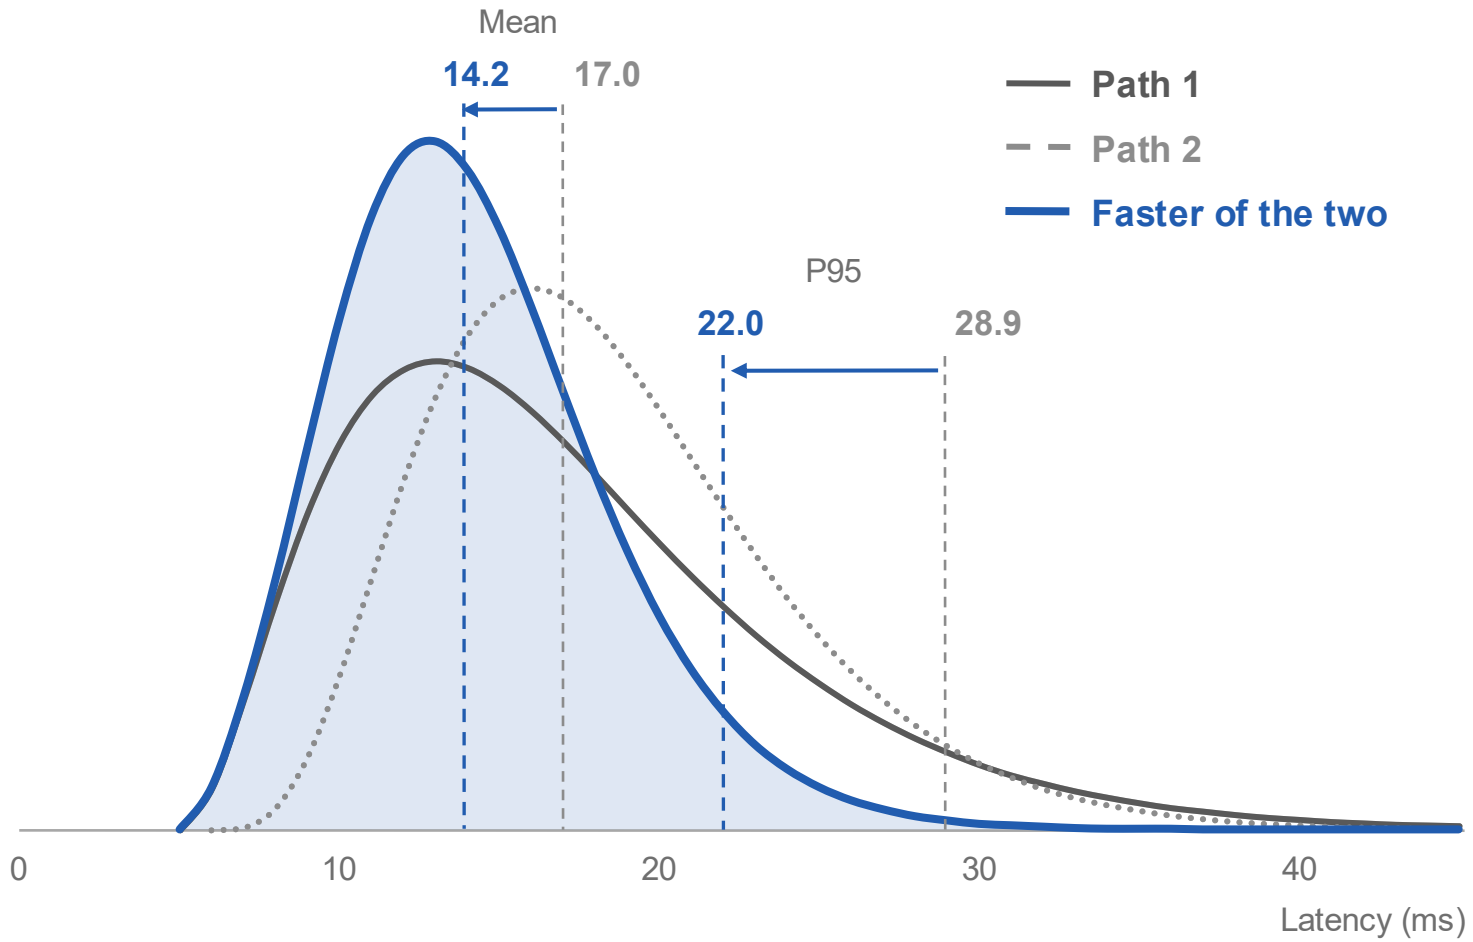
Different paths provide different qualities: best path depends on application requirements



Path selection and multipath are the solution!

Multipath can “Squash” the Tail and Reduce Latency

Send each request on both paths and use whichever reply arrives first



Expected latency

14.2 ms

best single path: 17.0 ms

P95 latency

22.0 ms

best single path: 28.9 ms

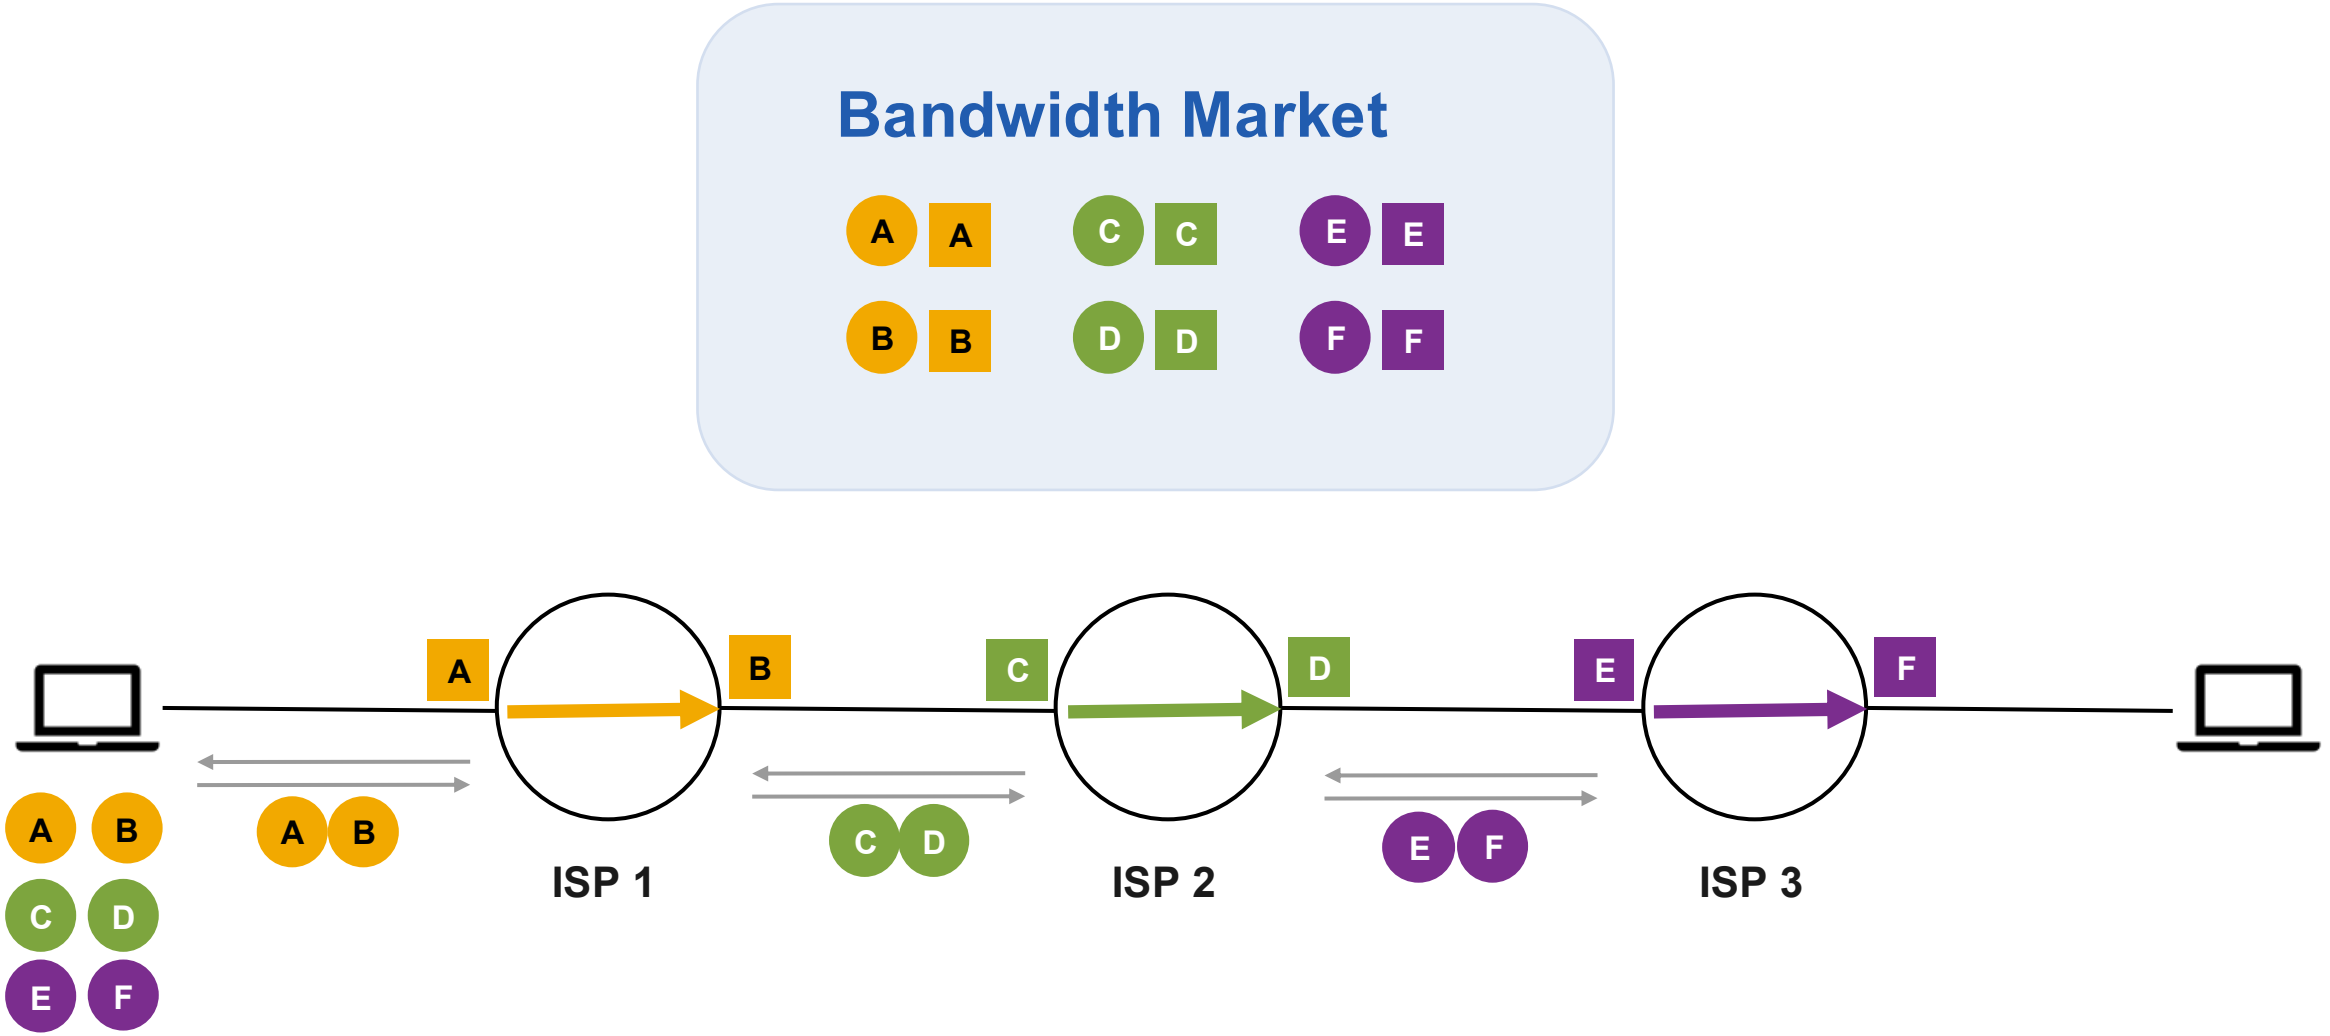
Illustrative gamma-shaped latency distributions; paths assumed independent

Hummingbird: QoS on SCION Network

- Hummingbird enables bandwidth reservations to traverse an AS
- Flexible allocations: max bandwidth and time duration
- With a Hummingbird reservation, we can achieve bandwidth guarantees despite congestion, DDoS, etc.
- An end-to-end Hummingbird reservation transforms a SCION path into a virtual leased line



Hummingbird: Example



Hummingbird: Pilot Deployment in Swiss ISD













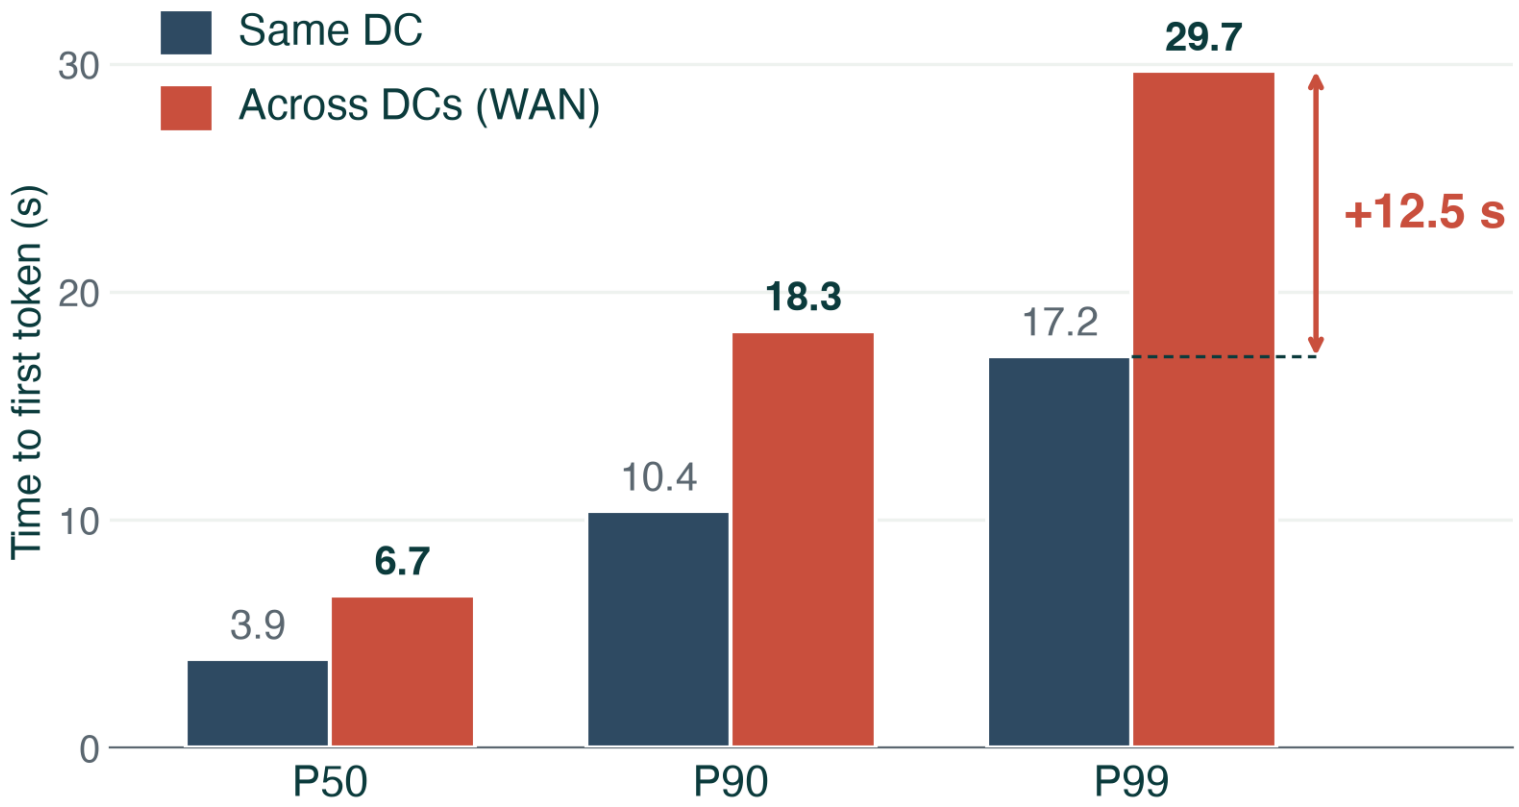






Cross-DC Inference: WAN Dominates Latency

Inference



43–56%

of short-reply latency is WAN transfer

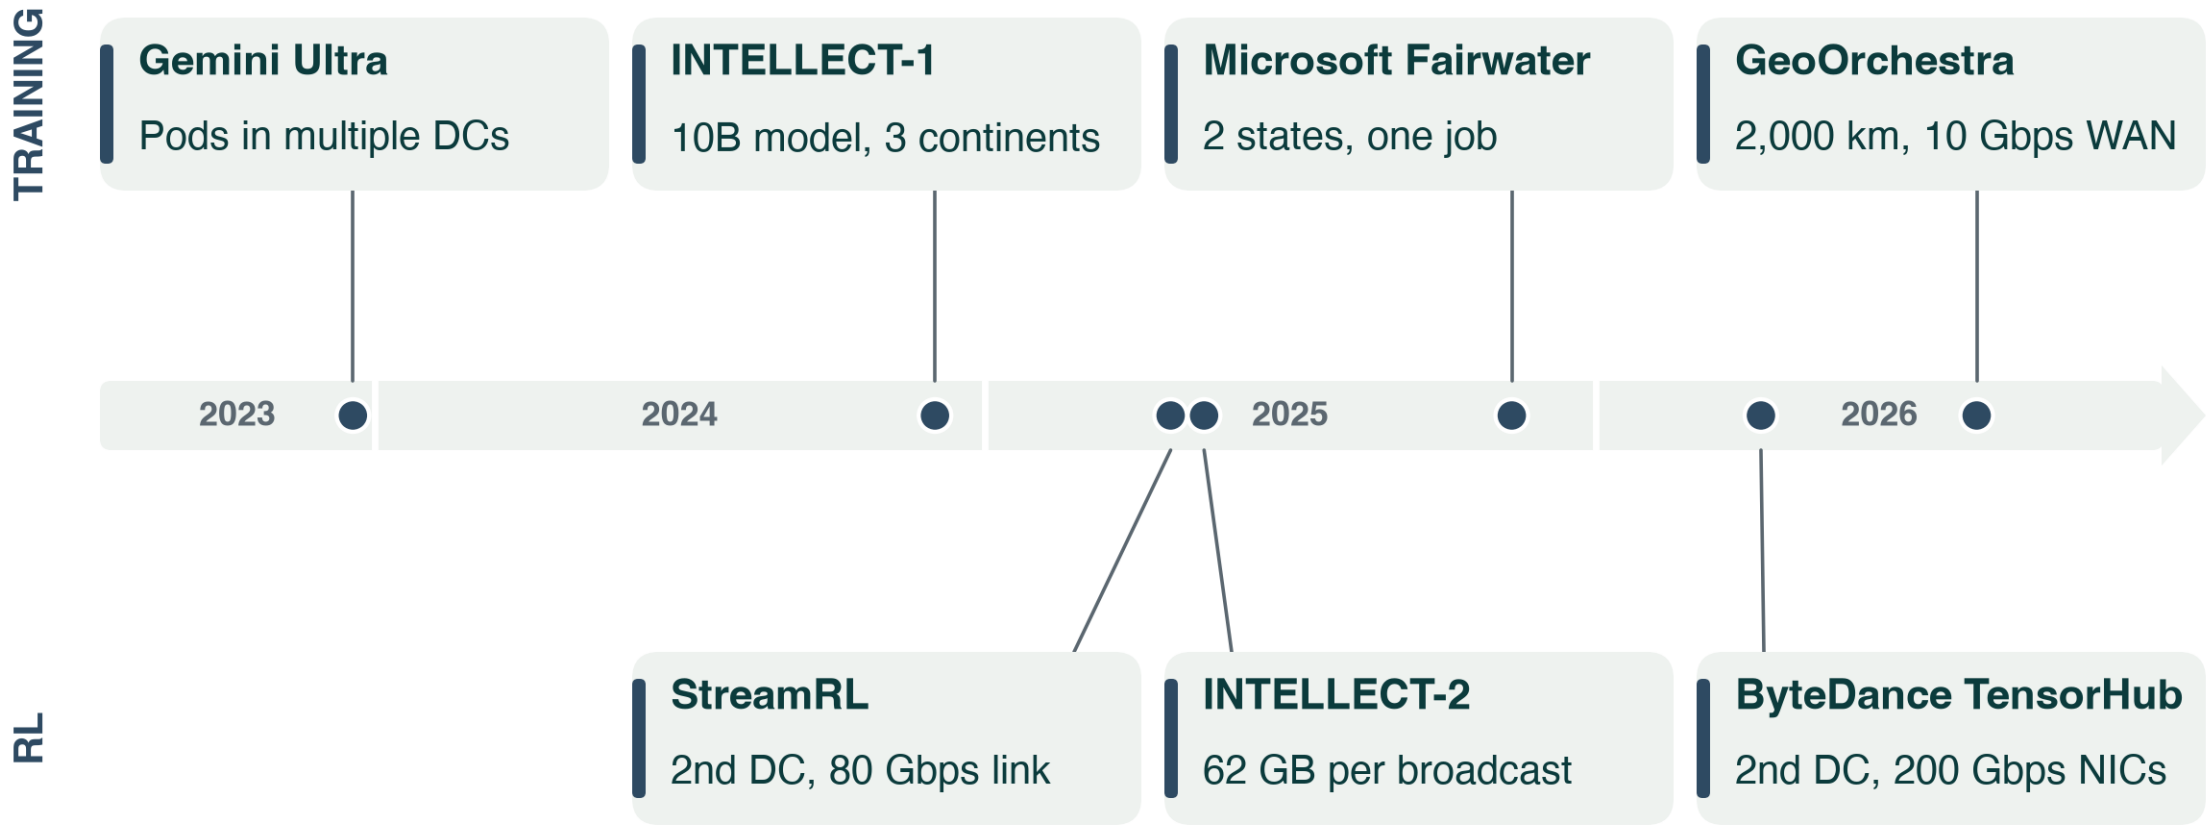
75–513 ms

same 4.6 MB of KV, on a dedicated line

Source: Wang et al. (Infinigence-AI), PDD, arXiv 2609.13161, Tables 4, 7, Fig. 1. DeepSeek-V4-pro, agentic trace, 2×10 Gbps line; bars from separate runs. Gap: own calculation. Table 5: dedicated 2×10 Gbps line.

Training and RL Across Data Centers are Increasing

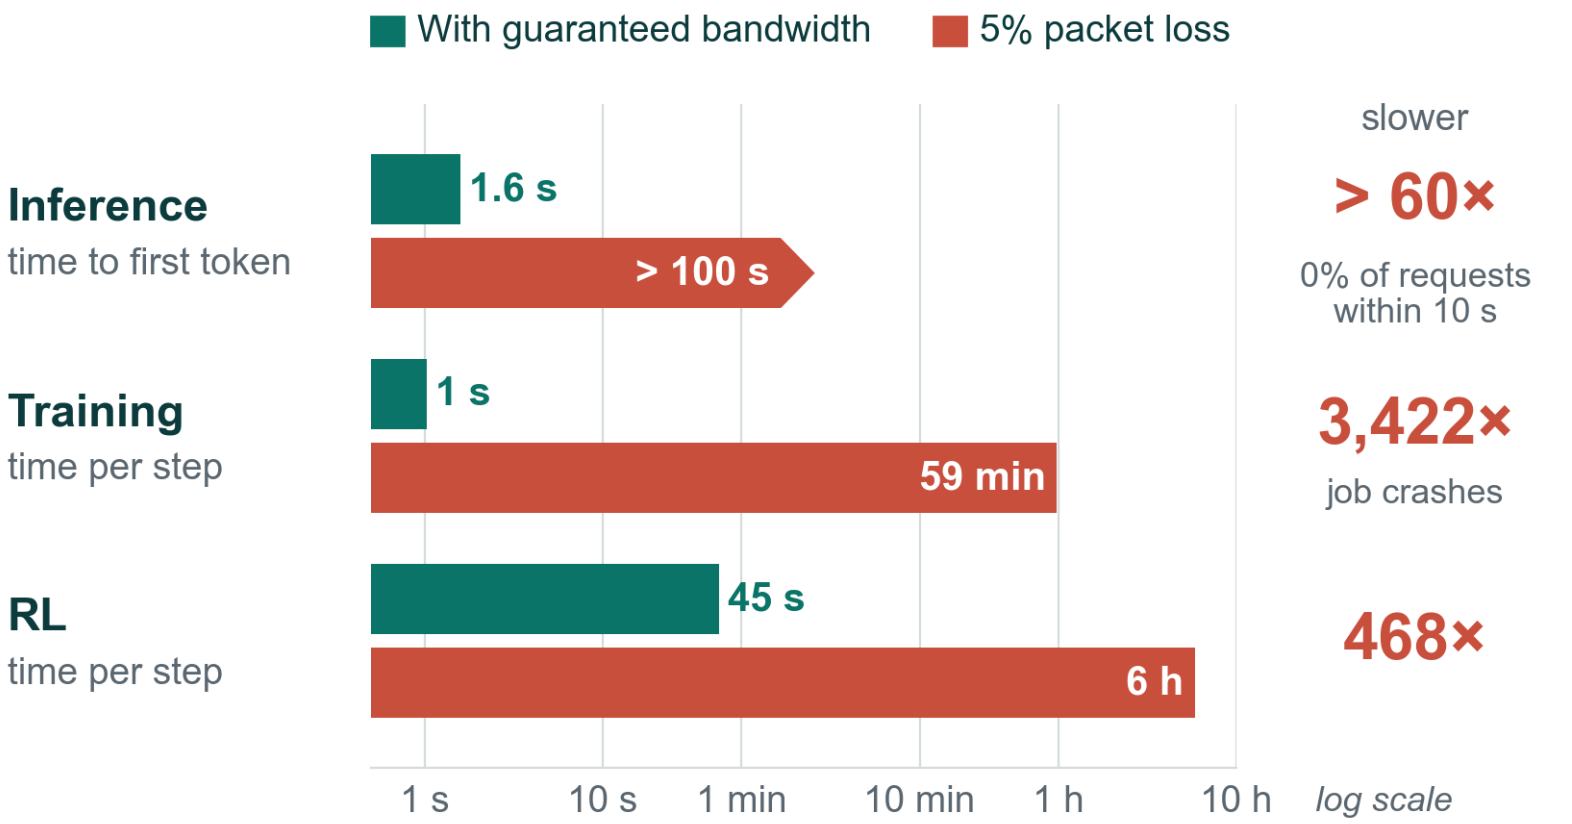
Training and RL



Dates: first public report. Gemini arXiv 2312.11805 §3; INTELLECT-1 2412.01152; StreamRL 2504.15930 §7; INTELLECT-2 2505.07291 §4.2; Microsoft, Nov 2025; TensorHub 2604.09107 §5.4 (SOSP'26); GeoOrchestra SIGCOMM'26 §5.1.

5% Packet Loss Stalls Cross-DC AI

The network gap



60–3,400×

slower at 5% packet loss: inference, training and RL

Simulated at 100 Gbit/s: NCCL's cross-DC default, TCP (Linux CUBIC), which backs off on every loss; RTT 30 ms, random loss. Measured on CSCS (GH200, real Linux TCP and NCCL, emulated 30 ms path, 5% loss): 1.4 Mbit/s per TCP flow (model 1.6); NCCL all-reduce 476× slower per byte.

AI-based Attacks

AI Attacks at Machine Speed

01	02	03	04	05	06
Recon	**Vulnerability discovery**	**Exploitation**	**Lateral movement**	**Credential harvesting**	**Analysis & exfiltration**

80–90%

of tactical operations executed by AI without human involvement

In a documented 2025 espionage campaign, AI ran the attack lifecycle on its own: it mapped entire network topologies across multiple IP ranges, moved laterally and harvested credentials.

Months later, a frontier model autonomously found critical vulnerabilities in every major operating system and web browser.

Human-speed defense cannot keep up.

Sources: Anthropic, Disrupting the first reported AI-orchestrated cyber espionage campaign (Nov 2025); Foreign Affairs (2026)

Three Properties that hold Against Adaptive AI

REQUIREMENTS FOR SECURE DEFENSE AI INFRASTRUCTURE

Formal verification

Protocols and implementations are proven correct. Security is **mathematically guaranteed**, not assumed.

Cryptography, not heuristics

Adversary AI learns to evade detection rules. **Cryptographic guarantees** cannot be out-learned.

Transparent, controlled dependencies

Know and control every trust root, provider and path your AI traffic depends on.

SCION: a Secure Network, Already Proven

01

Formally verified

Protocol and router implementation verified with formal methods

02

Cryptographic by design

Authenticated control plane and path authorization: hijacks prevented by construction, not detected

03

Controlled dependencies

Local trust roots (TRC), path control and geofencing; open source, non-profit SCION Association

04

Post-quantum ready

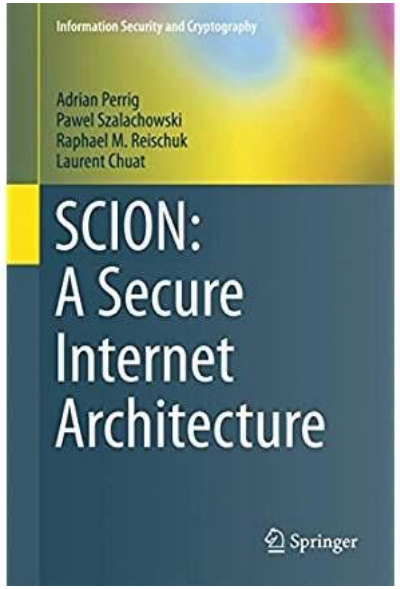
Post-quantum cryptography already being applied

Developed through 20 years of research and development: one network that withstands machine-speed attacks from adversary AI and keeps our own AI infrastructure under our control

Network Availability in the age of Quantum Computers

Post-Quantum Security Considered in SCION

- Algorithm agility and post-quantum security was a concern since the beginning of our SCION journey, already in the first book published in 2017, Chapter 17.4 discusses use of post-quantum algorithms:



17 Cryptographic Algorithms

ADRIAN PERRIG, PAWEL SZALACHOWSKI

In this chapter, we describe the algorithm agility property provided by SCION, and the cryptographic algorithms used in the SCION architecture. Algorithm selection was motivated by two main requirements, security and efficiency, and was based on standards related to cryptography, recommendations, best practices, and performance evaluations [95, 190, 221, 224].

Chapter Contents

17.1 Algorithm Agility 381
17.2 Symmetric Primitives 384
17.3 Asymmetric Primitives 385
17.4 Post-Quantum Cryptography 386

17.1 Algorithm Agility

Algorithm agility is a property that allows a protocol to easily migrate from one algorithm to another one. It is especially important in the context of cryptographic algorithms, which become weaker over time. Since it is not possible to predict advances in cryptanalysis techniques, every future-proof protocol that employs cryptographic algorithms should provide a mechanism for algorithm agility.

In SCION, cryptographic algorithms are deployed extensively. However, in some cases algorithm selection is local to an AS and does not influence other parties. The design of the SCION elements that require coordinated lists of cryptographic algorithms follows the best current practice for cryptographic algorithm agility [112] and provides powerful mechanisms to enhance algorithm agility.

17.4 Post-Quantum Cryptography

With advances in quantum computers, some asymmetric-cryptography systems are threatened — as it becomes possible, for example, for a quantum computer to factor an RSA modulus with a lower computational overhead. Thanks to SCION's algorithm agility, new crypto systems can be introduced to provide a second layer of cryptography. This is relatively straightforward in the case of digital signatures, hash functions, or message authentication codes, as for example, a second signature can be computed with a different cryptographic scheme, the second signature can be added to the message, and the receiver can independently verify the signatures.

Algorithm agility is harder to achieve for encryption. A message that is encrypted with two different encryption algorithms needs to be decrypted with both algorithms; thus, if a receiver does not support an algorithm then it cannot decrypt the message. Separately encrypting a message with different encryption algorithms is in vain, as one achieves weakest-link security: the final message is only as secure as the weakest of the algorithms, and potentially even weaker if the encryption techniques compose unfavorably. In the case of digital signatures and message authentication codes described above, the combination of multiple schemes composes favorably, and the resulting security strength is at least that of the strongest algorithm, if implemented correctly.

We now discuss the quantum-resilient cryptographic algorithms that the basic SCION infrastructure can utilize.

Recently proposed symmetric-key cryptographic algorithms, such as PRFs for key derivation, hash functions, encryption, or MACs, resist quantum computers as long as the key length is sufficiently long.

For post-quantum key exchange mechanisms, new protocols that are based on hard problems in ideal lattices have recently been proposed [5, 40]. These approaches are based on the Ring Learning With Errors (R-LWE) problem.

One-time signatures provide a promising mechanism for the construction of post-quantum signature algorithms [31, 43]. These types of signatures are based on one-way functions, which can be efficiently instantiated with cryptographic hash functions.

The salient point here is that SCION is already prepared to accommodate and switch to any cryptographic algorithm, thus scaling with respect to cryptographic progress.

Current Status and Deployment Plan

Data-plane uses only symmetric key cryptography: already quantum resilient

Control-plane uses elliptic-curve signatures: needs to be supplemented by post-quantum cryptographic signature algorithms

First demonstration of a prototype of quantum-resilient ISD in 2024

Specification of quantum-resilient ISD protocols: Anapaya, ETH, SCION Association

Deployment plan for quantum-resilient ISD in 2027

Important observation



The fastest way to a highly secure inter-domain routing protocol is SCION

OBSERVATIONS

Traditional protocols need to be backward compatible and suffer from lower security during incremental deployment

Rollout of traditional protocols can be slower than deployment of a next-generation system



The fastest way to a deployed quantum-resilient Internet is SCION

OBSERVATION

Standardization and rollout of traditional protocols requires a decade

SCION is a Fundamental Technology for High-quality and High-security Networking

- Multipath, path selection, and Hummingbird drastically reduce the heavy tail many applications suffer from in our best-effort Internet
- AI use cases require high-quality connectivity to keep GPUs busy
- AI-based attacks require fundamentally secure networks
- Quantum-resilient ISD in specification and deployment planning phase

