KASPR KASPR Datahaus

Methodology factsheet · data to 16 July 2026

How KASPR measures the internet

KASPR Datahaus measures internet connectivity and quality from outside the networks, every hour, across most of the world. This factsheet sets out what is measured, how it becomes data for districts, regions and countries, and what that data can and cannot tell you.

3 billion+
measurements a day

400 million+
IPv4 endpoints probed

26,324 districts
in 2,921 regions and 136 countries

Since Feb 2019
hour by hour

How KASPR measures

KASPR sends ICMP echo (ping) queries to geolocated IPv4 addresses from cloud infrastructure worldwide — over 3 billion measurements a day across more than 400 million endpoints. Each response records a round-trip time; results are aggregated to district, region and country by hour. Where address space is dense a random sample of 100,000 addresses is measured; where it is sparse, every address is. The design targets a consistent baseline across thousands of small boundaries over time, rather than a globally representative sample of all addresses.

From addresses to places

Each responding address is placed in its district (ADM2), region (ADM1) and country by IP geolocation. Responses are then aggregated by area and hour: how many addresses answered, how quickly and how steadily. Every figure is an aggregate over many addresses and providers in an area, never a single household or connection.

The core metrics

Metric	What it measures
Connectivity	The number of unique IPv4 addresses that answer in a period. Outages and shutdowns show as falls in this count.
Connectivity index	Connectivity set against the area's own recent baseline, which reads 100. Below 100, fewer addresses are answering than usual.
Latency	The mean round-trip time of the area's responses. Lower is better.
Stability (jitter)	How much the round-trip time varies. Lower is steadier; high jitter breaks video calls and live streams even where latency looks fine.
Relative quality	The gap to the fastest district in the same country, in absolute terms and as a percentage.

What it's built on

7+ years
of KASPR's own hourly measurements

Global monitoring
We continuously measure over 400 million endpoints across 130+ countries.

6.7 trillion records
Machine learning models forecast typical hourly patterns from them.

Consistent
remote measurement technology and granular resolution.

KASPR Datahaus measurement has been covered by Reuters, The Economist and The Guardian, and is distributed via Snowflake, Carto and AWS. A Monash University spin-out.

Anomalies and disruptions

Connectivity and latency anomalies come from a statistical model over a 90-day trailing window. Each series is normalised so its long-run normal reads 100; a fixed-effects model on hour-of-day and day-of-week gives the expected value, and 99.9% prediction intervals ($\alpha = 0.001$) set the normalcy band. A reading outside the band is flagged as an anomaly. Flagged readings are joined into episodes with a start, an end and a depth, and classed as a spike (a single day), a sustained drop, an outage or a collapse.

The typical week

KASPR's models, trained on its hourly history of 6.7 trillion records, give every area a typical week: one national weekly shape per country, with each district setting its own level, how hard it swings and how spread out its measurements are. For every local hour this gives the expected latency and stability and the chance of a bad hour, judged against a bar shared by every place so that places stay comparable. KASPR Schedule and the dashboard's timing views read it.

Connectivity Score and grades

Each district gets a score from 0 to 100 over the past 90 days, from five components: speed 30%, stability 25%, evening 20%, reliability 20% and trend 5%. Each component is a percentile among scored districts worldwide. A district is scored when it has at least 60 usable days in the last 90. Location grades A to E are fifths of the scored districts in the same country; a country with fewer than 10 scored districts is graded against the world. A district without enough measurement of its own takes its region's score and is marked as estimated.

Products and cadence

Product	Geography	Cadence
Regional Intelligence	District, region, country	Hourly, daily, weekly, monthly and yearly, from February 2019
Situational Awareness	District, region, country	Hourly, at 1, 3 or 6-hour aggregation
Telco / ISP	Operator by district	Daily, matched to stock tickers
Typical week	District, region, country	One profile per local hour, renewed with each release

Delivered through the KASPR Pulse dashboard, the KASPR Data MCP server for AI assistants, Snowflake Marketplace, CARTO Spatial Data Catalog and AWS Marketplace.

What the data is not

- IPv4 only. Addresses that do not answer ICMP are not measured, and geolocation is third-party, refreshed monthly, so its accuracy varies across places and over time.
- Measured from outside the networks. Part of a district's latency is its distance from the measurement, so compare places within a country, or one place over time.
- An aggregate of the measured addresses in an area, not what any one person's connection does.
- Counts of responding addresses, not subscribers or market share.

Data is descriptive and provided as is. It approximates infrastructure and user experience through remote measurement.
