# How Computer Networks Route Data

Routing is like sending a letter through post offices — each office decides where to forward it next.

Data doesn't travel in a straight line; the internet is a mesh of independently managed networks.

Key concepts: packets carry addresses, routers check tables, paths adapt to conditions.

Learning goals: why packets take multiple hops, how routers choose next steps, and what affects path selection.

# Packets, Addresses, and the First Hop

- Data is split into small chunks called packets before transmission.
- Each packet carries a source and destination IP address — like a letter's From and To.
- Your device doesn't know the full path; it sends packets to a default gateway.
- A request from a laptop begins by heading to the local router.

# The Hop-by-Hop Journey Across Networks

- Each router-to-router step is called a 'hop'.
- Routers act as traffic directors, choosing the next hop.
- IP addresses stay the same, but MAC addresses change at every hop.
- The TTL field prevents infinite loops by limiting maximum hops.

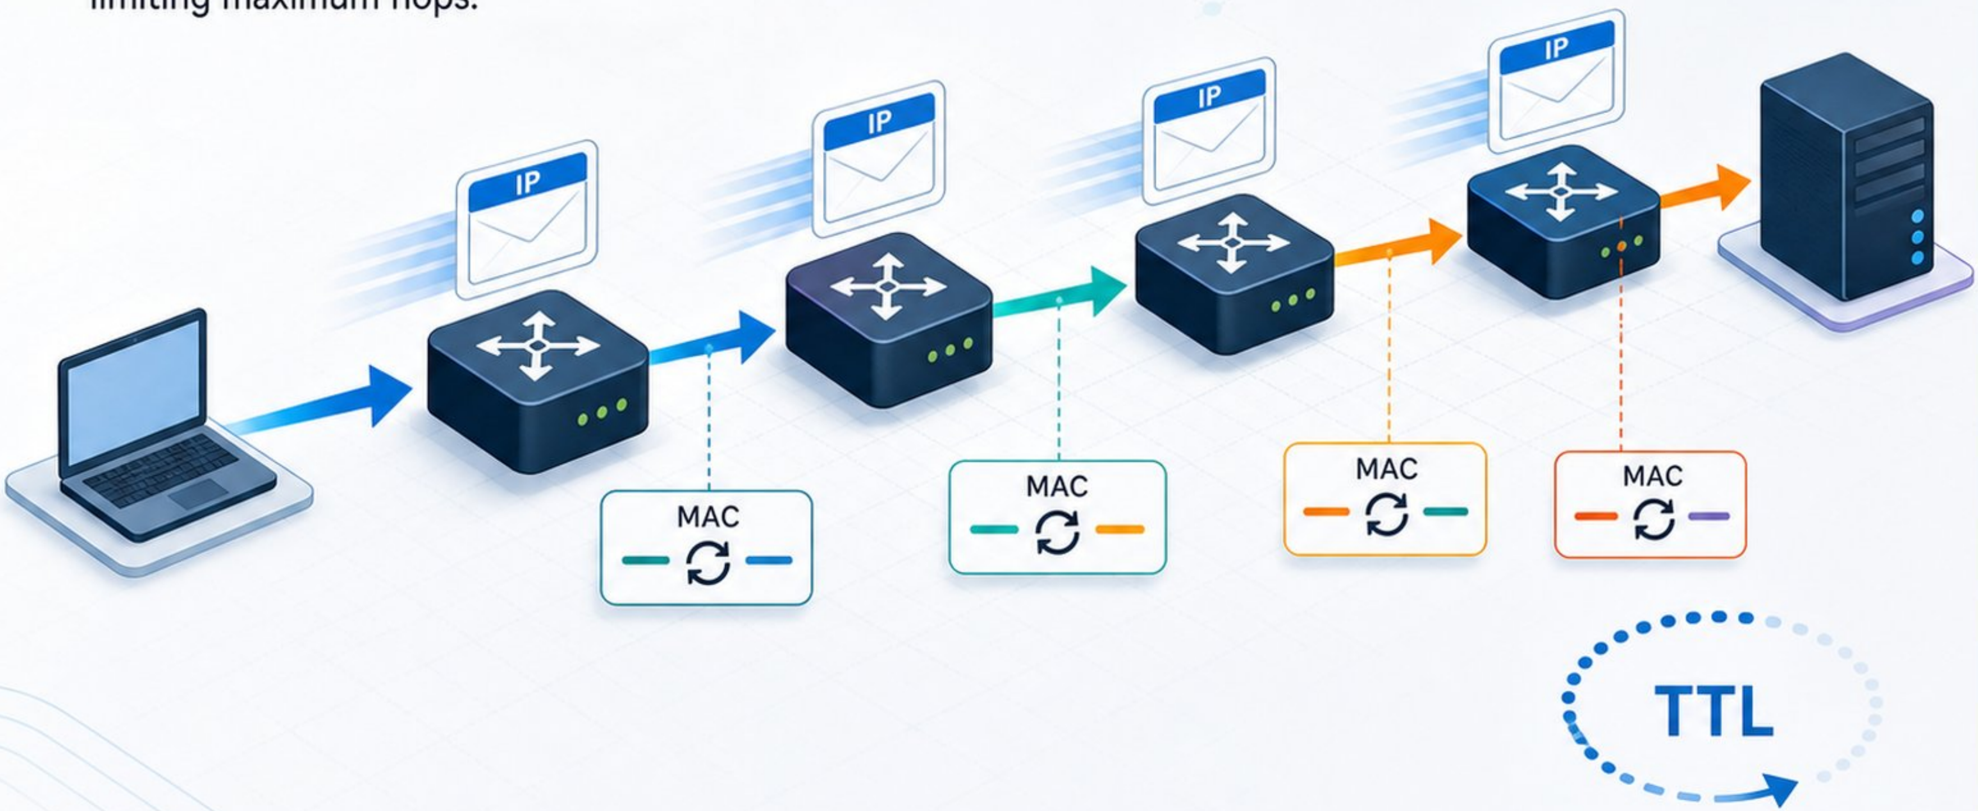

# Inside a Router: The Decision Table

- Each router maintains a routing table listing destination networks and the next hop to reach them.
- Routers only know the next hop and trust it will continue forwarding the packet correctly.
- Longest prefix match: when multiple routes fit, the most specific prefix (e.g., /24 over /16) wins.
- A routing table works like a transit map — you just need to know the next station.

# Why Packets Go Through Multiple Networks

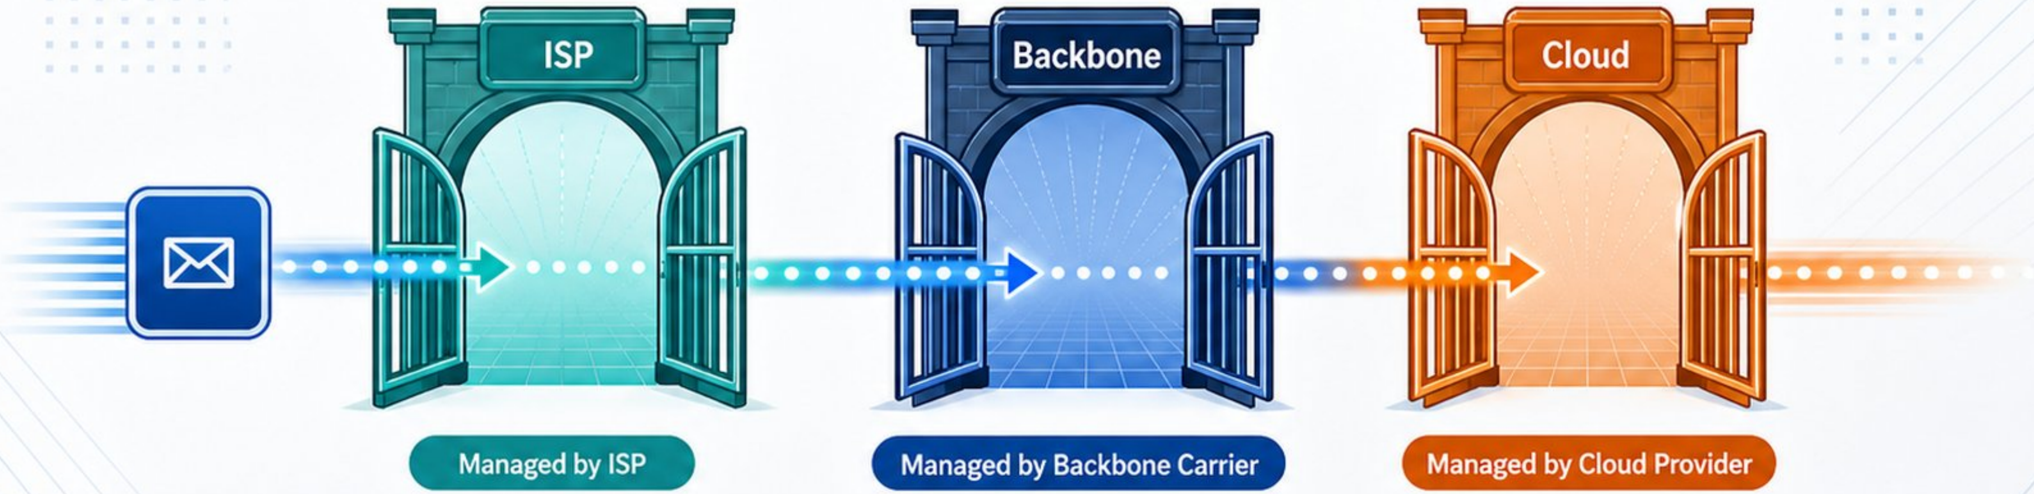


- The internet is a network of networks — no single company owns the entire path.
- Different organizations run different parts: your ISP, backbone carriers, cloud providers.
- **Peering:** direct, settlement-free traffic exchange between networks for mutual benefit.
- **Transit:** paid service providing access to the rest of the internet through an upstream provider.
- Like international mail crossing postal systems, packets traverse independently managed networks.

# How Routers Learn Paths: Routing Protocols

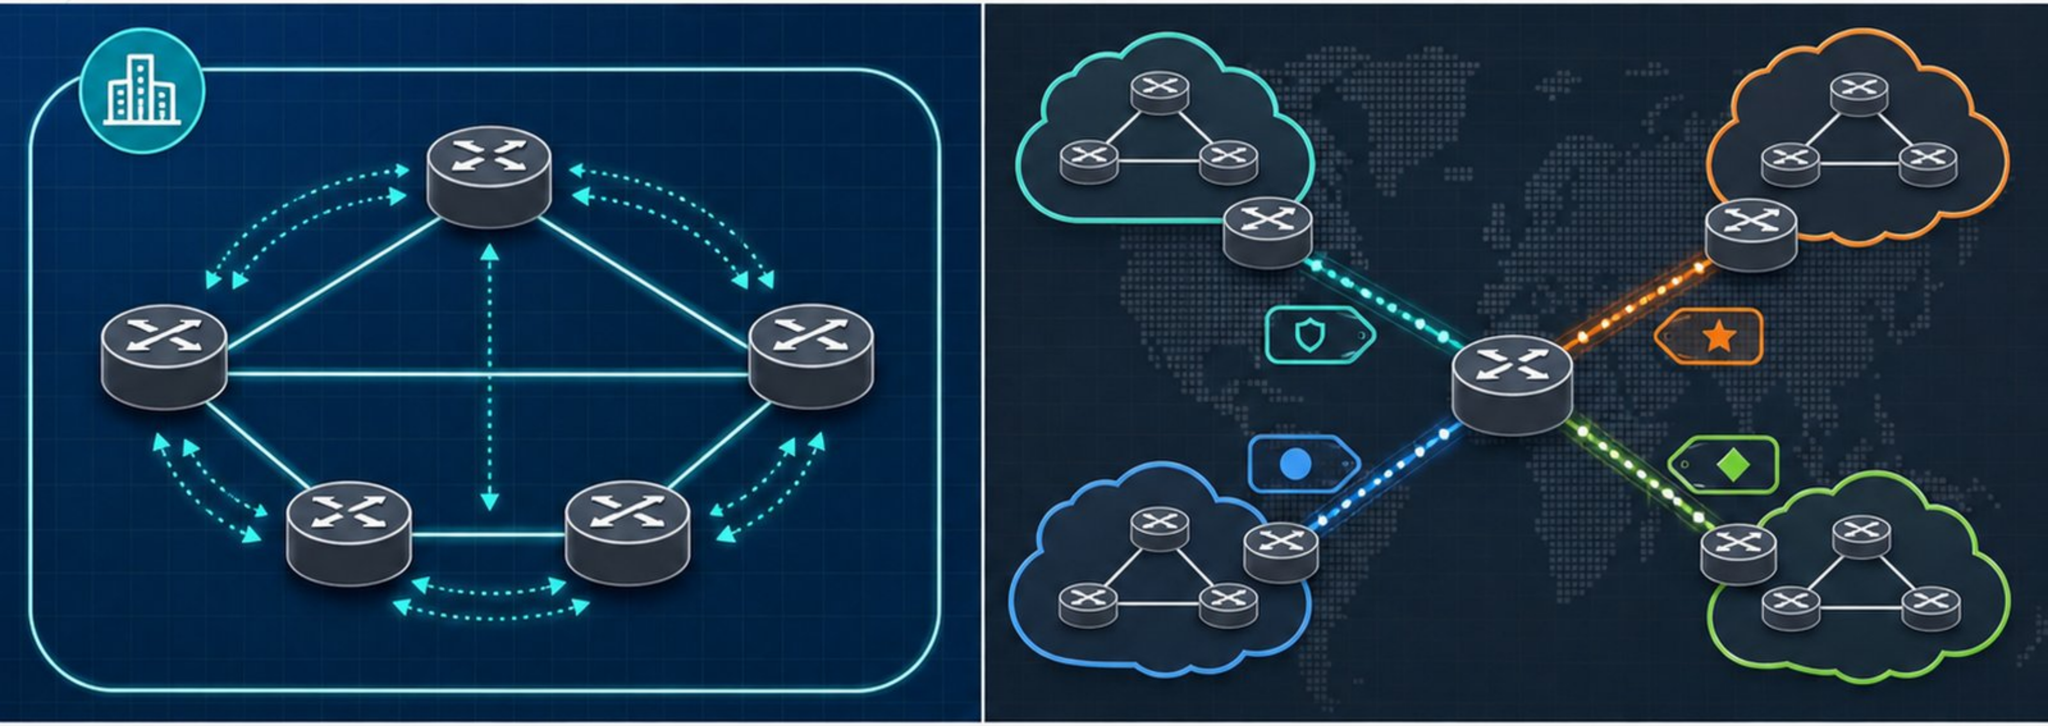

- Routers automatically share reachability data with neighbors—manual config can't scale
- OSPF (Interior Gateway Protocol) handles routing inside one organization;
  BGP (Exterior) links separate ASes worldwide
- OSPF picks fastest internal path by link speed; BGP selects inter-domain paths via policy attributes
- BGP is the "glue" of the Internet, enabling ISPs to exchange reachable IP prefixes

# BGP and OSPF: A Layered Partnership

- OSPF runs inside a network; fast, bandwidth-aware, and handles internal reachability.
- BGP runs between networks; policy-driven and scales to the entire internet.
- Most networks use both: OSPF inside, BGP at the edges.
- OSPF provides the underlying reachability that keeps BGP sessions alive.
- They solve different problems – they are complementary, not competing alternatives.

# What Can Change the Chosen Path

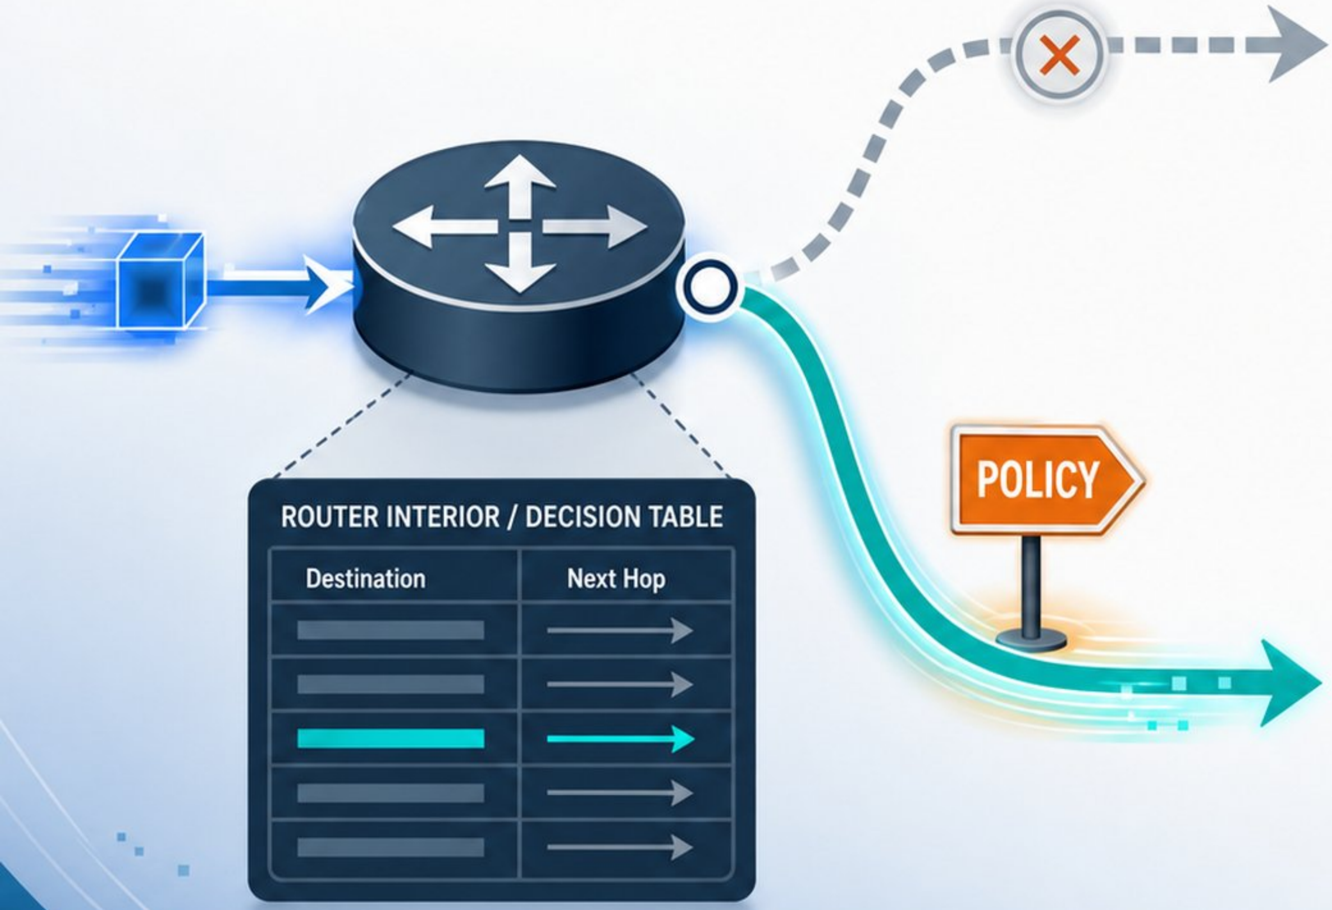


- **Link failures**: routers detect changes and find alternatives in seconds
- **Congestion**: traffic engineering shifts data to less busy paths
- **Business policies**: cost and peering often override shortest technical path
- **Example**: two nearby cities may route through a distant hub due to peering economics

# Real-World Routing Anomalies

- 200 km Saipan-Guam packets re-routed 12,000 km via LA when local peering failed during typhoon
- Cape Town-Sydney traffic took a London detour due to peering economics and cable topology
- Detours aren't bugs: routing prioritizes reachability and economics over geographic efficiency
- The internet was designed to be reachable — not necessarily geographically optimal

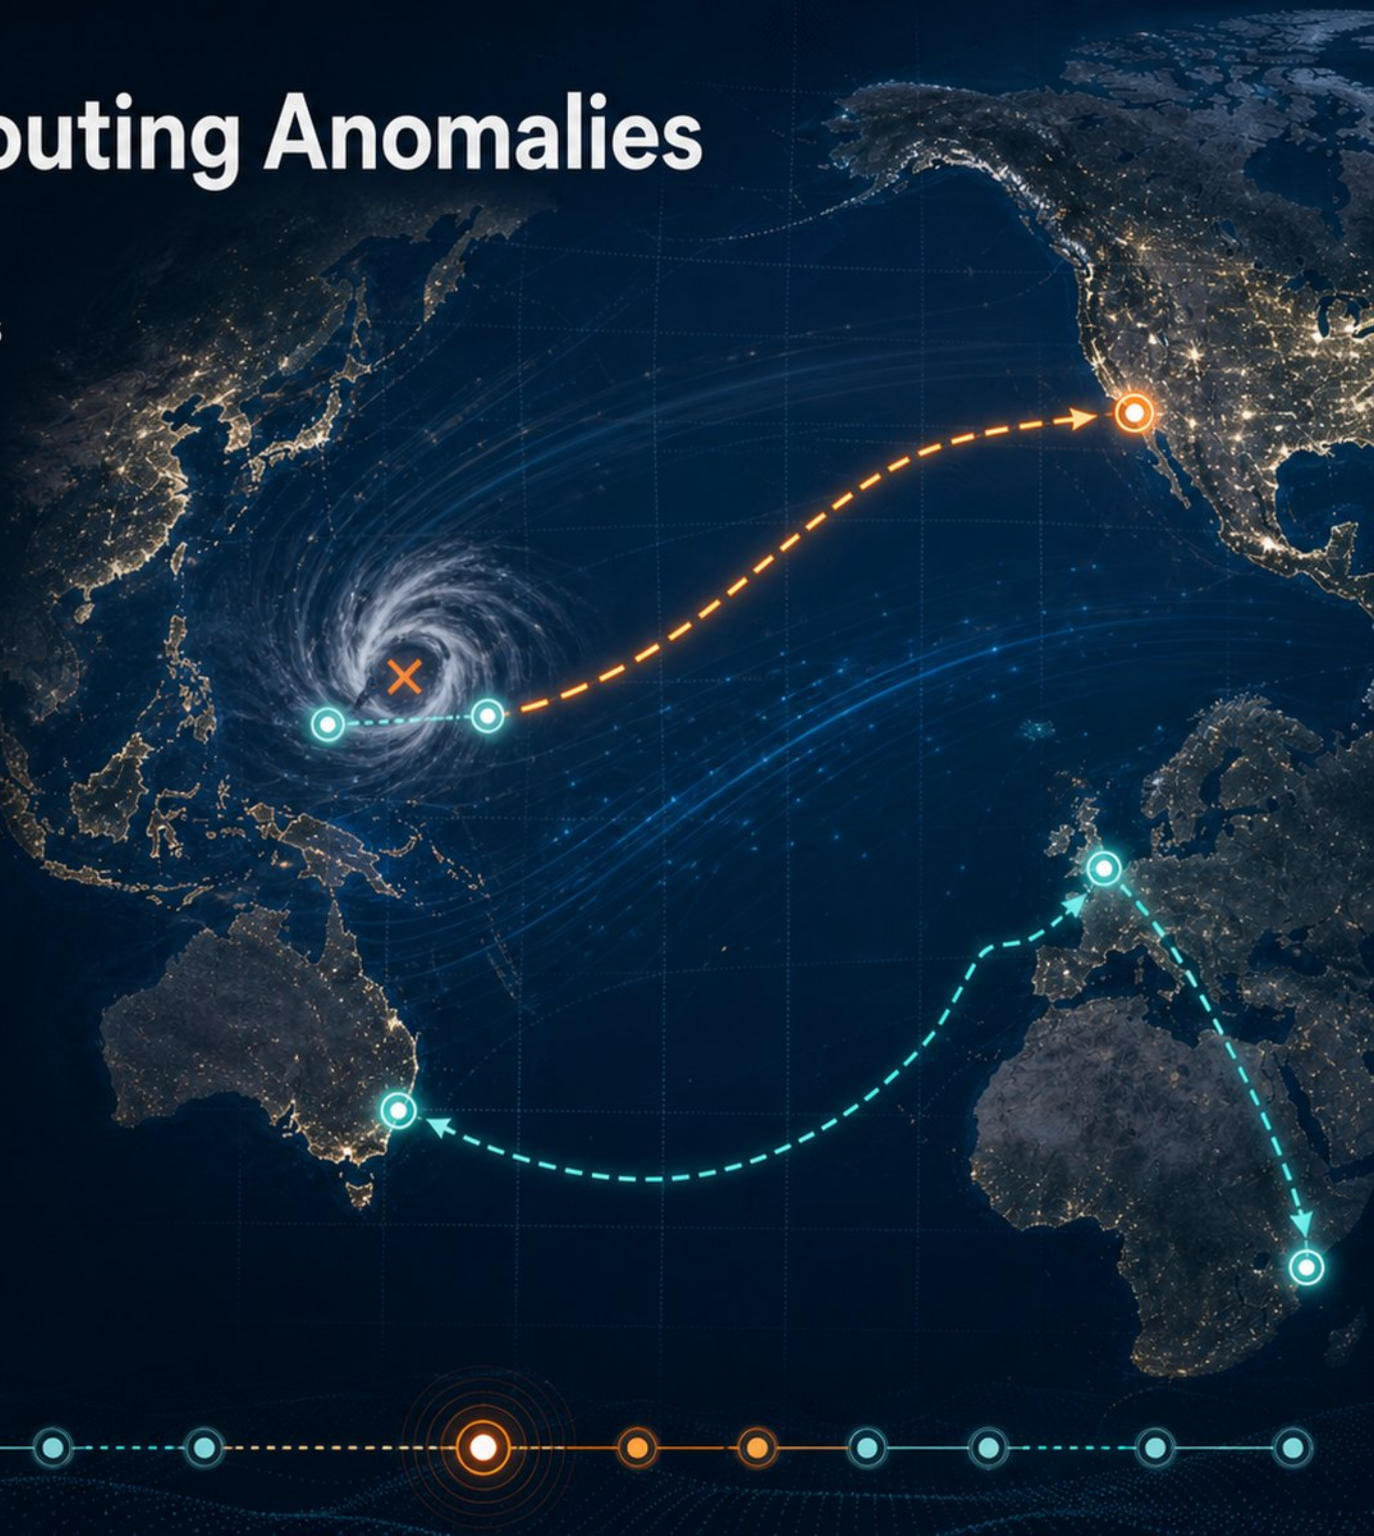

# Traceroute: Seeing the Path for Yourself

- Traceroute maps each router hop between your device and a destination
- Each line shows hop number, IP/hostname, and three RTT measurements in ms
- Asterisks (*) usually mean a router ignores probe packets, not a failure
- Focus on where latency jumps and stays high; a single slow middle hop is harmless

```
traceroute to example-destination (198.51.100.1), 30 hops max
TTL   HOP      ADDRESS / HOSTNAME          RTT (ms)
1     192.0.2.1                            1.1    1.2    1.0
2     198.51.100.2                         2.3    2.1    2.2
3     203.0.113.5                          3.6    3.4    3.5
4     *   *   *                            *      *      *
5     203.0.113.9                          45.7   46.2   45.9
6     198.51.100.25                        44.8   45.0   45.1
7     198.51.100.50                        6.2    6.0    6.1
8     198.51.100.1                         6.3    6.2    6.2
```

**TTL (Time To Live)**

Each probe starts with a TTL. Routers decrement it; when it reaches zero, the router responds, revealing that hop.

**Asterisks (*)**

Asterisks mean the router didn't return a response to our probes, not a failure.

**Latency Spikes**

Look for where latency jumps and stays high. Those hops may be the source of slowdown.

# Putting It All Together: An End-to-End Trace

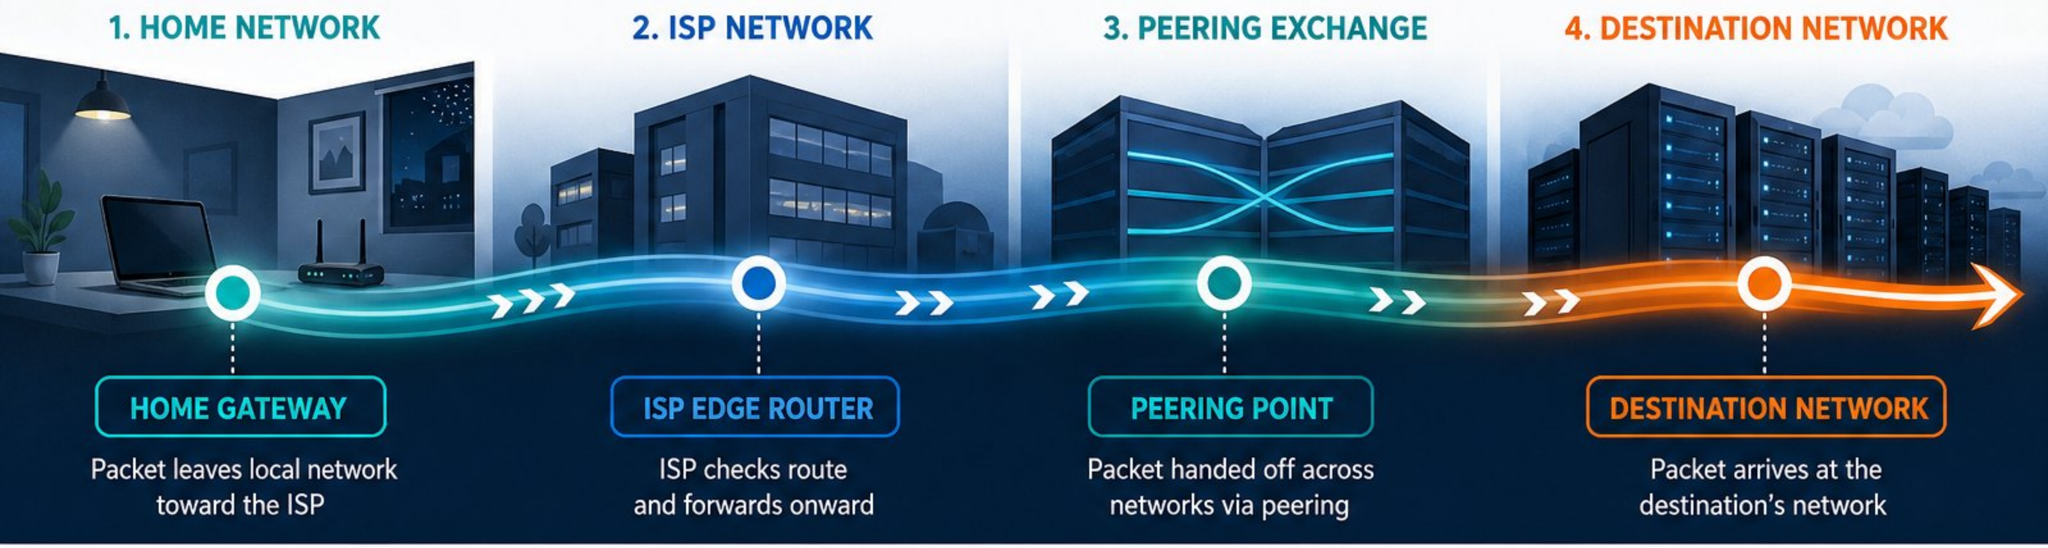


- Laptop sends packet to default gateway when the destination is outside the local network
- Home router checks its routing table and forwards the packet to the ISP
- ISP uses BGP to send the packet across a backbone to the destination's network
- Each router rewrites the MAC address, leaves the IP alone, and reduces TTL by one
- Traceroute maps this hop-by-hop journey using incrementing TTL values

## HOP-BY-HOP JOURNEY

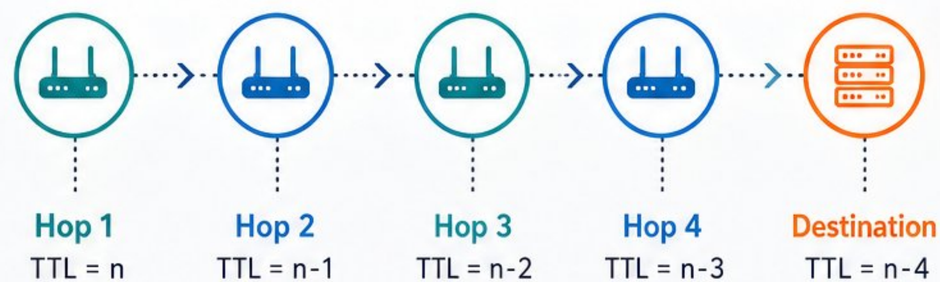


MAC address is rewritten | IP address stays the same | TTL is reduced by one

# Common Myths and Misunderstandings

- Packets don’t always follow the shortest physical path — routing prioritizes policy and reachability.
- No single organization carries your packet end-to-end; the Internet is a network of networks.
- Routers only know the next hop, not the entire end-to-end path.
- Peering exchanges only partner traffic; you still need transit for global access.

# What This Means for Everyday Internet Use

- Your latency and reliability depend on routing decisions by many organizations — not just your ISP.
- Traceroute reveals where delays or packet loss occur along your path when diagnosing issues.
- Routing awareness explains why CDNs place servers closer to users for faster websites.
- Peering and transit economics cause regional performance differences for the same service.

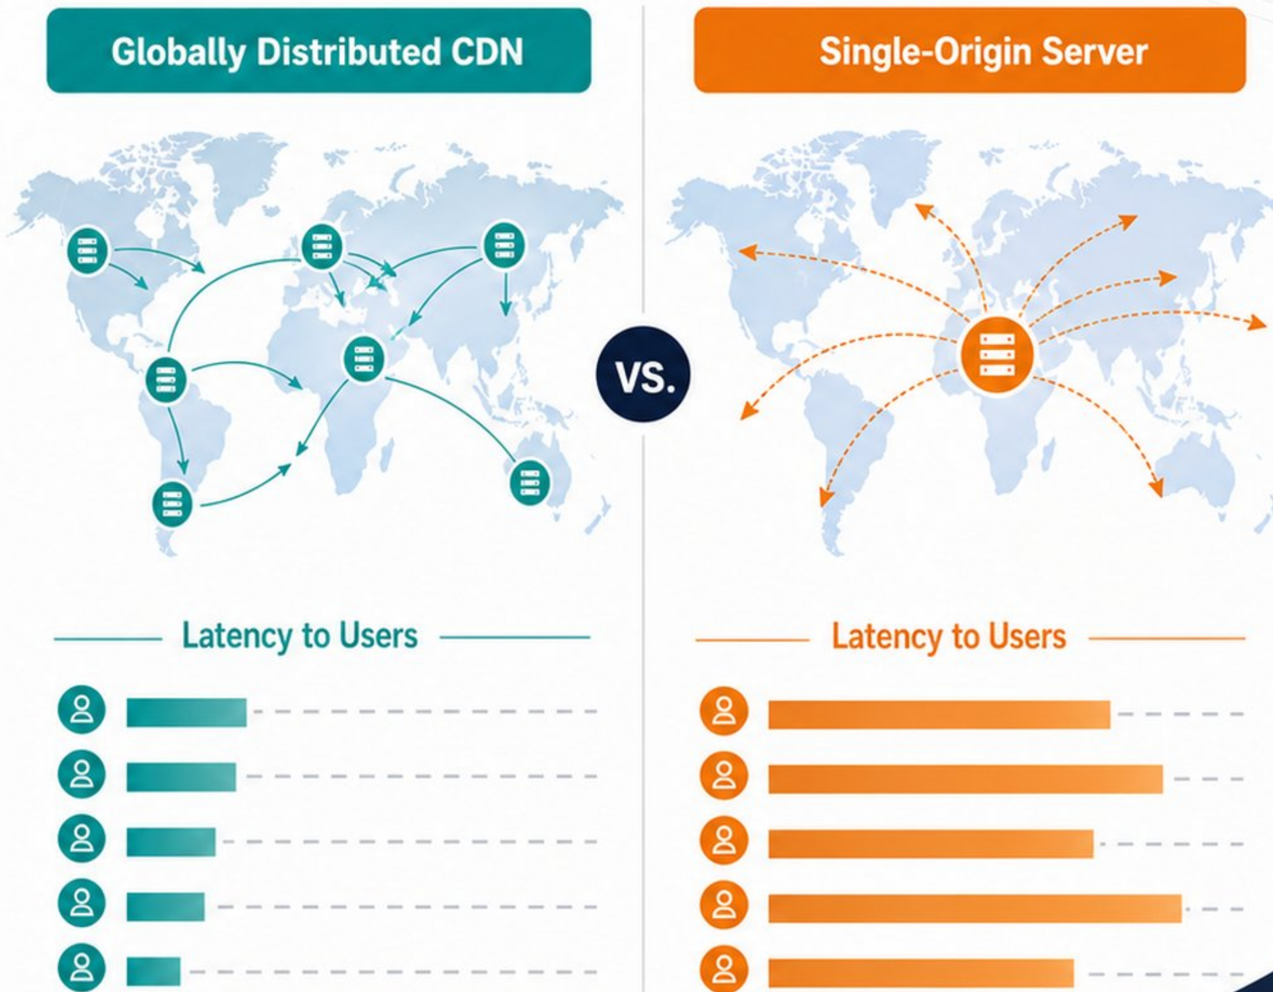

# Key Takeaways and Next Steps

- Packets carry addresses; routers use tables to pick the next hop
- OSPF and BGP automate path discovery inside and between networks
- Paths change with link failures, congestion, and business policies
- The internet favors global reachability, not geometric efficiency
- Try traceroute to see real paths; explore how CDNs optimize performance