

A LETTER FROM THE FOUNDER

Who Owns the On-Ramp to Intelligence?

The compute is landing. The road to it is not built. The fabric inside the data centers is being solved; access, the layer that decides who gets paid for intelligence, is still unclaimed. It should belong to the people who own the infrastructure.

Eli Karpilovski Founder and CEO, insidepacket

Consider a research team staring at a finished dataset and an empty calendar slot. The model is ready. The GPUs are available: a hall full of them sits four hundred miles away with capacity to spare this weekend. Between the two stands one of the oldest problems in networking: 400 terabytes, a procurement process built for three-year circuits, and a VPN that needs weeks to move what the science needs moved by Monday.

For decades, compute was scarce and the network was furniture. You bought it once, signed for three years and forgot it. AI broke that arrangement. Capacity now lives across hyperscale regions, GPU clouds, colocation halls and private clusters. Demand arrives in bursts, intense for hours or days, then gone.

*The compute is landing.
The road to it is not built.*

01

The era in three squares

The people leading this era describe it plainly. Sam Altman wants a factory that produces a gigawatt of new AI infrastructure every week; he titled the ambition “Abundant Intelligence.” The hyperscalers and frontier labs, xAI, Meta, Microsoft, Google and Oracle among them, are committing capital on the same scale. NVIDIA describes AI grids¹ spanning roughly 100,000 distributed network data centers, with more than 100 gigawatts of potential power over time. The infrastructure side reads it the same way: Zayo’s Bill Long describes a two-sided ecosystem⁷, intelligence as the supply, enterprises and consumers as the demand, and the operator’s job as removing the friction between them so infrastructure is reached at software speed.

Read those claims together and the era sorts into three squares. The compute square is being solved with more capital than any infrastructure cycle has seen. The fabric square, inside and between the data centers, is being solved by the best networking engineers alive. The third square is the one none of them is building: the commercial access layer, where an enterprise, or its software, discovers intelligence, buys it and receives a private path to it in minutes, and where the owner of the road gets paid.



Utilization says the third square is the constraint. Average GPU utilization sits near 5 percent², against the roughly 70 percent a 1,024-GPU H100 cluster needs to break even; the same cluster loses an estimated \$330,000 a month running at 55 percent³. The compute is bought and paid for. The road to it decides how many of its hours are ever sold.

~5%

Average GPU utilization across the installed base²

~70%

Utilization a 1,024-GPU H100 cluster needs to break even³

\$330K

Estimated monthly loss for the same cluster running at 55 percent³

02

The pattern we have seen before

Public clouds did not win enterprise workloads on processors alone. They won on reach and operability: private connectivity in major metros, a consistent interface and a common way to consume networking, storage and identity. When AWS launched Direct Connect, it looked like a useful feature. It became part of the moat.

Capital already agrees that value is moving to the layer above the assets. Lumen paid \$475 million for Alkira in May 2026⁴ and described the deal as a control plane for cloud connectivity. Whatever the final shape, intelligence will be created and consumed across many locations, owners and jurisdictions, and those resources must behave less like isolated assets and more like one programmable system.

\$475M

MAY 2026


Lumen's acquisition of Alkira, described by Lumen as a control plane for cloud connectivity: capital agreeing that value is moving to the layer above the assets.⁴

*When AWS launched Direct Connect, it looked like
a useful feature. It became part of the moat.*

03

Who owns the on-ramp

My answer is everyone who runs it. The on-ramp is not one company's product. It is a marketplace that lives on infrastructure people already own. Every GPU and neocloud provider is listed on it under its own brand, capacity and price. Operators run the road and bill both sides of it. MSPs and integrators wrap it in services and resell it. Providers sell more of the hours they have already bought. Enterprises trade networking projects for minutes. insidepacket builds the platform; the market keeps the margin. We sell no compute, which keeps the road neutral ground for everyone on it.



THE MARKETPLACE An on-ramp that lives on infrastructure people already own.

We sell no compute, which keeps the road neutral ground for everyone on it.

Providers that cannot expose their own and their partners' assets through a common, programmable model face the alternative: interchangeable capacity beneath someone else's customer relationship. The buildout alone settles nothing. The service layer that makes distributed infrastructure consumable decides who keeps the customer, and it should sit with the people who own the assets.

04

The buyer is changing species

The next enterprise customer of connectivity is software: an agent with a deadline, a budget and no patience for a quote. Each action such an agent takes on infrastructure is a routing decision, a security decision and, the part the industry has not priced yet, a purchase. Serving that buyer rests on three operating principles.

- | | | |
|----------|--|--|
| 1 | ONE SERVICE
DEFINITION LANGUAGE | A person, an application or an authorized agent describes the outcome across network, data center, cloud and compute without reproducing the operational logic beneath it. We use that model in production today and intend to publish an open proposal, because a service language becomes a standard through implementation and adoption, not declaration. |
| 2 | ANY SCALE, ANY
LOCATION, ANY TIME | Demand does not arrive in one fixed size, place or schedule, so provisioning and change run through software rather than a chain of manual orders. North America is live; Europe is live with Zayo Europe ahead of public availability. Multi-provider federation is the work in front of us, and “any location” grows as participating infrastructure grows, with each provider keeping control of its assets and economics. |
| 3 | FROM INTENT TO
GOVERNED ACTION | Within the production platform, authorized agents already use Model Context Protocol to discover approved capabilities and drive governed parts of the service lifecycle, with every action subject to available inventory, commercial rules, security policy and operational authority. This is not an agent let loose on a network. It is a network that answers agents, with the provider in control, and as of today it exists in production nowhere else. |

This is not an agent let loose on a network. It is a network that answers agents, with the provider in control.

05

Intelligence will be jurisdictional

Enterprises in Europe are already asking a question no hyperscaler answers cleanly: can my data reach the compute without leaving the country, on paths I can prove? A sovereign on-ramp, on-net, in-country and provable, is the condition under which the most regulated, most valuable workloads move at all. The access layer is where sovereignty stops being a policy paper and becomes an ordering option.

06

Proof before claims

I would rather show than claim. insidepacket's platform powers Zayo's DynamicLink, launched in September 2025⁵ and winner of the 2026 CRN Tech Innovator Award for enterprise networking⁶. DynamicLink reaches more than 90,000 enterprise locations⁶ and over 1,400 data centers across North America. The agent and MCP capabilities described above operate inside that production environment.

90,000+

Enterprise locations
DynamicLink can reach⁶

1,400+

Data centers across North
America

CRN 2026

Tech Innovator Award,
enterprise networking⁶

NA + EU

Live in North America; live with
Zayo Europe ahead of public
availability

Judge us on a dated calendar. In Q4 2026, DynamicLink moves from live with Zayo Europe to public availability, we publish agreed operating and commercial evidence from its production deployment, and we demonstrate a private enterprise-to-GPU service assembled from plain-language intent, with the tools, policies, approvals, actions and resulting service state visible, and a clear line between what is in production, what is being demonstrated and what remains roadmap. In H1 2027, we publish the open proposal for the service definition language and validate an initial multi-provider federation: one provider extending its sellable reach through another while each keeps control of its infrastructure, policies and economics.

THE DATED CALENDAR

Q4 2026

Public availability in Europe: DynamicLink moves from live with Zayo Europe to public availability.

Published evidence: agreed operating and commercial evidence from the production deployment.

Intent to service, demonstrated: a private enterprise-to-GPU service assembled from plain-language intent, with tools, policies, approvals, actions and service state visible.

H1 2027

The open proposal: the service definition language, published for implementation and adoption.

Federation validated: one provider extending its sellable reach through another, each keeping control of its infrastructure, policies and economics.

07

The world this builds

An enterprise agent asks for a path to the cheapest compliant GPU hour in reach. The marketplace answers with three options and a price. An operator's network carries the job and the operator bills both sides of it. A provider sells an hour that yesterday would have idled. Nobody left anyone's brand, and the whole exchange took less time than reading this paragraph. AI is the most urgent proving ground for this model, not its limit; the same layer extends to hybrid cloud, private data center access, on-demand interconnect and provider federation wherever infrastructure must become easier to productize, combine and sell.

The chips are deployed, the fiber is in the ground and the demand is real. If you own infrastructure and see a service your customers should be able to buy more easily, bring us the asset, the buyer and the commercial problem. We will show you the service it can become.

Eli Karpilovski

Founder and CEO, insidepacket

insidepacket.com

SOURCES

1. NVIDIA blog: AI grids for distributed inference across telecom networks. blogs.nvidia.com
2. Cast AI, GPU utilization analysis, April 2026.
3. ModulEdge, H100 cluster economics analysis, June 2026.

4. Lumen Technologies investor relations: Lumen completes Alkira acquisition, May 2026. ir.lumen.com
5. Zayo newsroom: Zayo redefines Network-as-a-Service with DynamicLink, September 2025. zayo.com
6. Zayo newsroom: Zayo's Naas platform wins CRN Tech Innovator Award, August 2026. zayo.com
7. Recon Analytics, The Week with Roger, episode 312: AI Infrastructure Connectivity with Zayo's Bill Long, September 2026. reconanalytics.com