

Data Center Operators

## Backbone for AI:

Make AI-ready connectivity easier to buy—and easier to trust




### AI puts pressure on DCO connectivity in two ways:

- **Training** drives large, irregular surges in data movement between sites and clouds.
- **Inference** becomes a persistent load as usage grows—performance issues show up as variability (latency spikes, timeouts), not always as clean outages.

When the backbone isn't designed for this, tenants feel it as slow deployments, inconsistent performance, and escalating operational friction.

### What the architecture shows



400G waves connecting Datacenters and Cloud On-ramp  
Dark Fiber connecting smaller Datacenters in the Houston metro

Data centers sit in the middle of a larger AI path: enterprise sites, AI compute environments, cloud on-ramps, and end users.

The handoff points—on-ramps, gateways/PoPs, intercity routes—are where latency variability and correlated failures show up first.

## Decisions that matter early

### 1 Connectivity tiers tenants can choose

- What does “unprotected vs protected vs diverse” mean for routes in your facilities—specifically?
- Can tenants match resiliency tiers to workload criticality (not vague promises)?

### 2 Diversity that’s defined, not implied

- Do “two paths” share the same corridor, conduit, or metro chokepoint?
- Are diverse paths truly geographically separate where it matters?

### 3 Capacity planning built for step-function growth

- How quickly can you add capacity when AI demand jumps—then jumps again?
- Which routes are long-haul intercity vs metro, and what’s your scaling plan for each?

### 4 Interconnect as part of the tenant experience

- Are on-ramps positioned and engineered as performance paths—not just cross-connect logistics?
- Can you explain failure behavior at the on-ramp/gateway layer in plain terms?

### ✓ Quick checklist (use in planning and new builds)

- ☐ Do tenants have clear resiliency tiers they can select with confidence?
- ☐ Can you explain diversity without jargon (geography, not diagrams)?
- ☐ If a path fails, does performance degrade gracefully—or collapse?
- ☐ Can you add capacity without redesigning the backbone?
- ☐ Quick test: if one metro construction event could impact multiple “redundant” paths, your diversity story isn’t real yet.

## What to do next

Map your highest-demand AI routes (metro and intercity), define resiliency tiers in operational terms (what degrades first), and pressure-test capacity expansion against “doubles—and then doubles again.”

## Download the AI Backbone Playbook



## About Lumen

Lumen is guided by our belief that humanity is at its best when technology advances the way we live and work. We deliver a fast, secure platform for applications and data to help businesses, government, and communities deliver amazing experiences.

## Get in touch

[lumen.com](https://lumen.com) | 1-877-453-8353