Flexential

Hybrid colocation and connectivity provider riding the enterprise and edge AI infrastructure wave

Explainer · June 1, 2026

ABSTRACT

Flexential is a US data center operator that sells colocation, cloud, connectivity, data protection, and professional services through its FlexAnywhere platform, a footprint of roughly 40-plus facilities across 18-19 highly connected markets stitched together on a 100 Gbps-plus private network backbone. What distinguishes it is positioning: rather than chasing hyperscale campuses in Northern Virginia, Flexential concentrates on Tier 2 and edge markets such as Denver, Atlanta, Portland, and Tampa, where it offers high-density colocation with custom cooling for enterprise and increasingly AI workloads. The company is privately held, formed in 2017 from the merger of Peak10 and ViaWest under GI Partners, and is funding growth through asset-backed securitizations rather than equity raises. The implication: Flexential is a real-estate-and-power business with a connectivity overlay, whose timeliness rides directly on AI compute demand spilling out of primary markets into secondary ones, and whose financial structure ties it to data-center capex cyclicality.

Keywords: colocation; edge data centers; high-density; AI infrastructure; FlexAnywhere; GI Partners; asset-backed securities; Tier 2 markets

1. Snapshot

Flexential (flexential.com) is a US provider of colocation, cloud, connectivity, data protection, and professional services, headquartered in the Charlotte/Denver orbit (press releases dateline both Charlotte, NC and Denver, CO). The company was formed in 2017 when investment firm GI Partners merged Peak10 and ViaWest. It remains privately held under GI Partners, and in August last year, GI pumped further investment into Flexential and secured new external investment for its data center developments across the US. The company is led by Ryan Mallory, who has been referenced as both President/COO and CEO across sources, so his exact current title is not cleanly confirmed. Financially, a 2026 securitization disclosure noted the pledged data centers generate approximately $663.3 million in annualized revenue and $353.2 million of annualized adjusted net operating income. Not publicly known: total company revenue, full headcount, customer count, and consolidated profitability.

2. Thesis: Why This Company, Why Now

The bet is that AI compute demand overflows the saturated primary data center hubs and lands in the secondary and edge markets where Flexential already owns capacity. The expansion of digital services and platforms is driving demand for advanced data center capabilities, but long-standing primary markets are unable to keep pace, prompting hyperscalers and enterprises to leverage second-tier markets to support their infrastructure and critical workloads. Flexential's CEO frames the AI linkage directly, calling US demand "tremendous" and positioning edge facilities to serve enterprise AI and "agentic AI modeling" workloads. The AI-demand connection is not a side use case here; it is the reason the company is timely, because high-density GPU deployments are precisely what spill into the markets Flexential serves. On scale, the reachable market is narrower than the headline AI capex frenzy: Flexential plays in retail-to-wholesale colocation for enterprises and mid-size cloud/AI tenants, not the gigawatt hyperscale builds. In its own survey, 79% of enterprises are mapping out data center capacity more than a year in advance, which supports the demand-pull thesis but also signals a tightening, competitive supply race.

3. The Core Idea in Plain English

Flexential rents secure, powered, cooled space and the network pipes to reach it, then wraps services around that. Think of it as a landlord-plus-utility for computers: instead of an enterprise building and operating its own data center, it leases racks or megawatts from Flexential and plugs into a private backbone connecting all the sites. The old-world version was a company owning a server room or a single colo cage in a major metro. The new-world version Flexential sells is a distributed platform spanning many secondary markets, engineered for high-density, power-hungry hardware, so a customer can place AI or latency-sensitive workloads closer to where data and users actually are rather than concentrating everything in one expensive primary hub.

4. The Technical Space

Colocation is fundamentally a power, cooling, space, and connectivity problem, and "good" is measured on a few axes that matter to buyers. First is power density and thermal handling: legacy colo was built for roughly 5-10 kW per rack, while AI training and inference hardware pushes 40-100 kW or beyond, demanding liquid or advanced air cooling. Second is reliability and resiliency, expressed through redundancy design and uptime commitments, plus the security and compliance certifications regulated buyers require. Third is connectivity richness: low-latency interconnection between sites, to clouds, and to network exchanges. Fourth is cost and proximity, the tradeoff between cheaper, available secondary-market power and the latency or ecosystem density of primary hubs.

The market splits between hyperscale wholesale (massive single-tenant campuses) and retail/edge colocation (multi-tenant, smaller footprints in more locations). Hyperscale data centers, which support massive cloud-scale workloads and platforms, accounted for 44% of total global capacity as of Q1 2025. The enterprise-and-edge segment Flexential occupies competes less on raw scale and more on density-per-rack, geographic spread, and the ability to serve hybrid deployments that mix colo, cloud, and connectivity.

5. How Their Technology Works (and What's Proprietary)

Flexential's "technology" is mostly an operated physical platform plus a connectivity and services layer, not novel silicon or software IP. The substance breaks into three components.

› **The physical footprint and density engineering.** The FlexAnywhere platform anchors services in roughly 40 data centers across 18 highly connected markets on a scalable 100Gbps-plus private network backbone. The company markets high-density colocation capabilities featuring innovative cooling designs and geographical proximity, from single rack to multi-megawatt deployments. The specific cooling designs and supported GPU densities are not disclosed in available material, so the "AI-ready" claim should be treated as a marketing assertion pending verified rack-density and liquid-cooling specs.

› **The private network backbone.** The interconnect linking sites is the closest thing to a differentiated asset, enabling multi-site, hybrid deployments and inter-data-center traffic without crossing the public internet.

› **The services and platform wrap.** Colocation, cloud, connectivity, data protection, and professional services are bundled under one platform and contract, and Flexential states its solutions are strategically engineered to meet the most stringent challenges in security, compliance, and resiliency. The specific certifications (SOC 2, ISO 27001, PCI, HIPAA) are not enumerated in available sources and warrant verification.

On defensibility: almost none of this is technically proprietary. Cooling designs, backbones, and service bundles are all replicable by well-capitalized competitors. What is hard to copy is the physical capacity itself, the owned real estate and secured power in the right markets, which is a capital and siting advantage more than a technology one.

6. Business and Go-to-Market

The model is recurring colocation revenue (space, power, cross-connects) plus higher-margin cloud, connectivity, and professional services, sold predominantly through a sales-led and channel-partner motion. Flexential runs an active partner program, references regional partner managers, and recently achieved Pinnacle Partner status in the Broadcom Advantage Partner Program for North America, and was named to CRN's 2025 Cloud 100 list. Pricing is not publicly disclosed but follows standard colo norms (per-kW or per-rack plus interconnect fees), scaling from single-rack retail to multi-megawatt wholesale.

On traction, the securitized portfolio's approximately $663.3 million in annualized revenue and $353.2 million of annualized adjusted net operating income implies strong facility-level margins, though that figure covers only the pledged assets, not the whole company. Demand concentration is a watch item: KBRA said 70 percent of the operators' annual revenue was concentrated in Portland, Denver, Atlanta, Dallas, and Nashville. The capital strategy is notable: rather than diluting equity, Flexential funds expansion via debt, including a move to secure $1.4 billion in asset-backed securities backed by a 28-site portfolio spanning 14 markets across 13 states, after an $800 million ABS fundraise the prior year. It is also shifting from leasing to owning facilities, e.g. recent Hillsboro and Atlanta acquisitions, to enable long-term investment for AI customers.

7. Competitive Landscape and Moats

Flexential sits in the crowded enterprise/edge colocation tier, below the hyperscale wholesale giants and above small regional operators. The comp set:

- **CyrusOne, CoreSite, and QTS** as adjacent US colocation and interconnection players.
- **Equinix and Digital Realty** as the scaled incumbents whose interconnection and global footprint dwarf Flexential's.
- **Cologix** as a direct edge/secondary-market rival.

The closest direct rival is Cologix, another North American edge-and-interconnection colocation operator focused on secondary markets and network density. Where Flexential wins: a broader high-density colocation push and an explicit AI/hybrid services bundle spanning more US markets with owned real estate. Where it loses: Cologix's interconnection density and network-neutral carrier ecosystem in certain hub markets can exceed Flexential's connectivity gravity, and the other adjacents (Equinix on global interconnection scale, the hyperscale-adjacent wholesalers on raw capacity) outclass Flexential on their respective dimensions.

On moats, two are real, one is asserted.

- **Owned capacity and power in the right markets.** Real and strengthening. The shift from leasing to owning, plus secured power in supply-constrained secondary markets, is a durable capital and siting advantage as power becomes the binding constraint on AI buildout.
- **Switching costs and the connected backbone.** Real but moderate. Once workloads, cross-connects, and hybrid architectures are deployed across the private backbone, migration is painful, but it is a normal colo lock-in, not unique.
- **"AI-ready" technical differentiation.** Largely asserted. The high-density and cooling claims are not yet substantiated with disclosed specs, and platform risk is real: hyperscalers and larger wholesalers can build dedicated AI capacity faster and cheaper at scale.

8. Risks and Open Questions

The picture turns on a handful of items, spanning technical, commercial, and capital-structure risk.

› **AI-readiness substance.** Can Flexential actually support 40-100 kW-plus racks with liquid cooling at scale today, or is "agentic AI" positioning ahead of the physical plant? Specific rack densities, cooling architectures, and AI reference designs are unverified.

› **AI-capex cyclicality.** Demand is being driven by an AI infrastructure boom; a slowdown in enterprise AI spending or a hyperscaler pullback would directly hit absorption in exactly the secondary markets Flexential is expanding into. This is a first-order risk, not a footnote.

› **Leverage and refinancing.** With growth funded through serial ABS issuance against the data center assets, rising rates or weaker NOI could pressure the capital structure. What is total company leverage and debt maturity profile?

› **Revenue concentration.** With roughly 70% of pledged-portfolio revenue in five metros, how exposed is the company to power, permitting, or demand shocks in Portland, Denver, Atlanta, Dallas, and Nashville?

› **Footprint clarity.** Sources cite 40, 41, 42, and 44 facilities across 18-21 markets; the exact current count is unconfirmed.

9. Bottom Line

Flexential is a credible, mid-tier colocation operator whose value sits in owned power and capacity across secondary markets that AI demand is now spilling into, not in any proprietary technology. It works if the edge-AI thesis holds and the company can deliver genuine high-density capacity before larger players saturate those same markets; it stumbles if AI capex cools or the ABS-funded buildout outruns absorption. Watch whether Flexential converts its "AI-ready" marketing into verified high-density, liquid-cooled deployments with named AI tenants, the single proof point that would separate substance from positioning.

10. For the Nerds

The defining technical bet is thermal. Retrofitting a colo fleet originally engineered for sub-10 kW air-cooled racks to support 50-100 kW AI nodes is not a marketing toggle; it requires rear-door heat exchangers, direct-to-chip liquid loops, or immersion, plus the facility water, power distribution, and floor-loading to match. The open question is how much of Flexential's existing footprint is genuinely liquid-ready versus newbuild-only (e.g., the 22.5 MW and 27 MW Denver and Hillsboro projects), because retrofit economics differ sharply from greenfield.

The second deep tension is the edge-versus-latency claim. "Agentic AI" inference at the edge is plausible, but training is centralized and bandwidth-hungry, so the backbone's east-west capacity and peering between secondary sites determines whether distributed inference is actually performant or just geographically dispersed colo. A 100 Gbps-plus backbone is respectable for enterprise hybrid traffic but modest against GPU-cluster east-west demands, which is the architectural detail that decides whether the AI story is real or repackaged.