

III FORUM

Forum Edge AI: Inference Infrastructure Opportunity

October 2026

Disclaimer

Unless the context otherwise provides, “we,” “us,” “our,” the “Company,” “Forum,” “Forum Markets,” and like terms refer to Forum Markets, Incorporated and its subsidiaries.

Disclaimers and Other Important Information


We have not authorized any other person to provide you with any information other than that contained in this presentation or in any other information prepared by or on behalf of us or to which we may have referred you. We take no responsibility for, and can provide no assurance as to the reliability of, any other information that others may give you. You should assume that, unless otherwise noted, the information appearing in this presentation is accurate only as of the date of this presentation. Our business, financial condition, results of operations and future prospects may have changed since those dates.

This presentation contains statements that constitute forward-looking statements within the meaning of the Private Securities Litigation Reform Act of 1995 and other applicable securities laws. All statements other than statements of historical fact are forward-looking statements, including, but not limited to, statements regarding the Company’s future financial position, business strategy, budgets, projected costs, and plans and objectives of management for future operations. These statements refer to many things, including future performance, and all other statements that are not historical facts, or that are intended to be forward looking statements, should be read as forward-looking statements. There are risks associated with the contemplated transactions, including regulatory and legal uncertainty, risks of loss associated with the industry, line of business, trade, customers, partners, custodians, and vendors of the Company, and other risks. Historical facts are presented without intent to persuade. These statements can be recognized by the use of words such as “believe,” “expect,” “anticipate,” “potential,” “create,” “intend,” “could,” “should,” “would,” “may,” “plan,” “seek,” “will,” “look,” “future,” “assume,” “continue,” or the negative of such terms or other variations thereof, or words of similar substance or meaning. Such forward-looking statements are not guarantees of future performance and involve risks and uncertainties, and actual results may differ from those in the forward-looking statements as a result of various factors and assumptions, that could cause actual results to differ materially from those contained in any forward-looking statement and which are inherently subject to significant uncertainties and contingencies that are or may be difficult or impossible to predict and are or may be beyond our control. The Company and its affiliates, shareholders, controlling persons, directors, officers, employees, agents, advisors and representatives assume no obligation to and do not undertake to update such forward-looking statements to reflect future events or circumstances.

All trademarks, service marks, and trade names of any party of their respective affiliates used herein are trademarks, service marks, or registered trade names of such party or its respective affiliate, respectively, as noted herein. Any other product, company names, or logos mentioned herein agree the trademark and/or intellectual property of their respective owners, and their use is not alone intended to, and does not alone imply, a relationship with any party, or an endorsement or sponsorship by or of any party. Solely for convenience, the trademarks, service marks and trade names referred to in this presentation may appear without the ®, TM or SM symbols, but such references are not intended to indicate, in any way, that any party of the applicable rights owner will not assert, to the fullest extent under applicable law, their rights or the right of the applicable owner or licensor to these trademarks, service marks and trade names.

Actual results may vary greatly from any assumptions or models built in reliance on this presentation. Results may vary due to market conditions, unforeseen circumstances or changes, competition, and results are subject to a multitude of risks, uncertainties, and changes. Those include but are not limited to, market conditions, the regulatory landscape, and other risks of loss.

Highlights




FORUM
EDGE AI

- ✓ **Preeminent management team:** Decades of collective experience across infrastructure, aviation, energy, and AI real assets paired with Edge Node AI's technical expertise to drive modular HPC deployments at powered sites from day one.
- ✓ **Modular deployment – Forum has fully funded Phase 0:** Forum's 51% JV interest funds an initial ~11 MW deployment, across tier 1 TowerCo site with spoke-and-hub model with two initial hubs
 - ✓ **Differentiated power footprint unlocked at tower sites and existing, owned hubs:** Dallas energized today; North Carolina energized in Q1 2027; excess allocated power capacity at Tower sites accessed via MSAs rather than multi-year interconnection queue, first online in Q2 2027
 - ✓ **Customer-led deployment model:** GPUs ordered against paired offtake agreements with B300 and GB300 costs and supply established; Forum Edge AI meets demand via proximity and rapid, modular deployment capability. Focus on edge sites meets future demand for low-latency inference.
- ✓ **Financeable, discrete project architecture:** Separate SPVs with paired offtake/collateral/equipment financing; no cross-collateralization; Forum controls the JV and consolidates without carrying 100% of project equity.
- ✓ **Scalable platform:** Phase 0 fully funded today; Phases 1 and 2+ are options earned through measured power, achieved utilization, ready offtake, and equipment financing. Optioned Tier 1 tower site footprint supports componentized roll-out across multiple phases.


Phase 0 Scope

Footprint and Economics At Fully-Funded Phase 0 Only




Illustrative Earnings


Consolidated at Forum




Where the Leverage Comes From



Already in Place



Then It Scales, Beyond It



Forum Markets Leadership

Executive Leadership



**McAndrew Rudisill,
Chairman and CEO**

- 25 years experience private and public equity & credit markets
- Founder Pelagic Capital Advisors
- Former Chief Investment Officer, Capital Vacations
- Former Chief Investment Officer, Bridger Aerospace



**John Saunders,
CFO**

- Senior finance executive across aerospace, defense & digital assets
- Former CFO, then SVP Finance, Bridger Aerospace
- Former CFO, Ascent Vision Technologies; 2020 sale to CACI



**Robert Spake,
General Counsel**

- Capital markets, fintech & regulatory compliance background
- Former Head of Litigation, Republic; CCO, Republic Capital
- Previously Polsinelli and Weil Gotshal; Delaware Chancery clerk



**John Kristoff,
SVP Corp Comm & IR**

- 30+ years in fintech marketing, communications & IR
- Former Chief Marketing & Communications Officer, Diebold Nixdorf
- Former VP, Investor Relations, EXL Service

Board of Directors



Jason New

- Vice Chairman, Investment Banking, Lazard
- Former CEO, Onex Credit
- Former Senior Managing Director, Blackstone
- Co-founder, NovaWulf Digital Management



Angela Dalton

- Founder, Signum Growth Capital
- Cofounder Evercore Equities
- Former Managing Director, UBS
- Former Managing Director, Guggenheim Partners



Michael Edwards

- Former Head of US Business at Arrowgrass Capital Partners
- Former Deputy CIO Weiss Advisors
- Former PM, D.E. Shaw
- Former M&A Banker Credit Suisse



Andrew Suckling

- Founder and Partner, Verulam LLC, Metals Trading House
- Non-executive Chairman, Cadence Minerals
- Former Partner, Ospraie Management




Ryan Smith


- President and CEO, Big Sky Industrials Inc.
- Former CFO, Emerald Oil
- Former VP Investment Banking, Canaccord Genuity

Extending Digital Infrastructure to Compute Financing

Forum's Platform Approach



Four On Platform RWA Classes Today



Naturally Meets AI Token Production


The compute footprint is moving; Forum has identified an under-served and under-capitalized opportunity

Centralized hyperscale → **Distributed metro edge**
 Greenfield megaprojects → **Underserved existing footprint**
 Capital-intensive builds → **Capital-efficient deployment**

Unlocking AI tokens at the edge in partnership with Edge Node

Forum Finances and Structures Every Link

Energy is the binding constraint on new AI capacity



Forum finances the equipment and places the cash flow

Origination, structuring & distribution under one roof

And Powers Multiple Value Propositions

Direct Project Ownership

51% of Forum Edge AI

The cash flows and distributions of every project entity, consolidated into Forum

Additional Forum Business Lines

Equipment Financing

GPUs are serial-numbered, standardized collateral with an established secondary market. The platform that already places real-world asset cash flows can place these


Credit Support and Structuring

Project SPVs, advance rates, guarantees, and intercreditor terms - the same toolkit, pointed at a different collateral type

AI Tokenization

GPU hours / compute tokens are emerging as a distinct, tradable asset class


Forum Edge AI Platform



Alignment through 9.95% cross-ownership. Forum will own 9.95% of Edge Node AI. Edge Node will receive 1.46m shares in Forum upon achievement of certain performance hurdles.

Unlocking Powered Compute Adjacent to Demand

Delivering distributed capacity in already energized footprints: quickly and without mega project risks




Faster and Lower Risk Across Every Vector

	Hyperscale Greenfield	Forum Edge AI — Tower and metro sites
Latency	80–200 ms round trip, outside real-time inference budgets	~5 ms round trip to the end user
Power lead time	A 4–7 year queue for a new grid interconnect	Existing permitted power, already at the meter
Permitting	2–5 years for new permits, with incremental political risk	Administrative approvals at an existing site, measured in weeks
Capital per site	\$500M–\$2B+, locked up for the length of the build	~\$40–50M, fast to deploy and financeable against the hardware
Capital at risk	Goes first into land and stays illiquid until built out	Goes into hardware that can be redeployed or sold
Capacity	Must be contracted in advance to justify the spend	Under-utilization is monetizable; surplus capacity sells at spot
Resilience	A single point of failure - one outage takes the campus down	Sites are independent; an outage at one does not reach the network
Counterparty	A utility interconnection queue and a municipality or state	An existing tower owner under an MSA

Risk-Controlled Scale into a Large, Durable Edge Inference Market

Inference is where the market is going; we can already deliver proximity with granular, controlled expansion


AI Inference and Edge Compute Markets



Future proofed by position, not just by forecast

Training concentrates; serving disperses. The market expects a migration on the order of ~70% of demand by 2030. We already own the hub and site destinations

Where That Future Demand Lands



We Scale Into It One Cohort at a Time

Phase 0 Dallas, NC and the first tower site

~11 MW

Underwritten on power that is live, cost that is priced and demand paired with offtake.

Phase 1 NC to 9 MW; the tower program up to fifty sites

63 MW

Released only after the pilot cohort has produced measured power, cost and utilisation.

Phase 2+ Two more regional hubs; towers to one hundred

120 MW


Held as an option. Its value rises with every metro entered; none of it is pre-funded or pre-obligated.

The Metro Module - One Hub, Up to Fifty Tower Sites

Tower sites carry 91% of the GPUs and 91% of the capital in a metro module

One Hub, Up to Fifty Tower Sites - Discretely Financed

Up to 50 tower sites
21,600 GPUs | 48.6 MW



Where the Value Sits

Up to fifty tower sites carry 91% of the GPUs and 91% of the capital in a metro module. The hub is the anchor and the aggregation point

Entities Are Split

Real estate and equipment carry different lenders, advance rates and tenors. Separating them keeps each financeable on its own terms

The Entities Behind It

Forum Edge AI JV

Forum Markets 51% | Edge Node 49%

Regional Hub

One per metro - separately funded and separately collateralized

Hub LandCo

Owns or leases the metro campus. Leases space and power to Hub EquipCo

Hub EquipCo

Owns the GPU fleet at the hub. Aggregation, orchestration and burst load

Tower EquipCo

Owns the GPU fleet at all tower sites - where the capital and the capacity sit. Pays each tower owner a monthly recurring charge under applicable MSA

Tower owner sits outside the JV; Land, structure, and power under the MSA

Hub Offtakers

Enterprise and AI-native customers. Contracted GPU hours


Tower Offtakers

Latency bound workloads inside the metro

Offtake Sits at the EquipCos

The contract has to be written by the entity that owns the compute. Each cohort separately bankable

Discrete Pods Scale by Tower Site, then by Metro Hub



One Tower Site, In Detail




- 6 GB300 racks, 432 GPUs
- 1–2 MW at the meter
- Liquid-cooled, six-rack build
- ~12 weeks from site order to live

Tower + equipment shelter

	Dallas	High Point, NC
Units	5 pods	37 → 55 racks
GPUs	1,280	2,664 → 3,960
Capacity	4 MW	6 → 9 MW
Energized	2 MW today	Q1 2027

The hub carries orchestration and burst load. The tower sites carry the capacity and the revenue

Capacity Compounds by Phase⁽¹⁾



Modeled:


Per B300 Pod 32 Nodes, 256 GPUs, 0.55 MW of IT Load	Per Tower 6 GB300 Racks, 432 GPUs, ~1.0 MW	Per Regional Hub 5 MW Hub + Up to 50 Towers	Contracted Rate >\$4.00 (\$/GPU-hour)	Uptime High 90% of GPU Hours	Contracted Share⁽²⁾ 100% of Phase 0 Hours
--	---	--	---	--	--

FORUM Note: Represented figures are per underwriting assumptions & model inputs.
1. Each block = 1 MW of contracted capacity per underwriting assumptions & model inputs.
2. Additional earnings may be available through virtualization.

Contracted Customers and Segmentation Strategy

Balances offtaker support for equipment finance with attractive spot given speed-of-delivery

Representative Offtakers



Multi-Tenant Virtualization
Virtualizing GPUs; Multiple tenants share each chip and sold through their marketplace

Stage **MSA⁽¹⁾**



Inference and Fine-tuning Platform
Install optimization software on client GPUs to run inference, fine-tuning, and model training

Stage **Finalizing MSA**

Open-Model Inference Provider

Open-model Inference Serving
Serving open-model inference at production latency on reserved metro capacity

Stage **Negotiating MSA**

GPU Spot Marketplace

Spot Marketplace Distribution
Marketplace distribution of uncontracted GPU-hours, absorbing the spot leg

Stage **Negotiating LOI**

Segments of the AI Stack

Segment	Edge Value Add	Contract Shape	Stage
Enterprise production AI	Data residency, fixed cost	Multiyear offtake	Documented
Exchanges	Discrete pricing, consumption	Token parcels	In discussion
Model providers	Burst and overflow capacity	Capacity reservation	TBD
AI-native software	Latency-bound product	GPU-as-a-service	TBD
Autonomous and robotics	Sub-5 ms latency in metro	Reserved capacity, metro	TBD


Workload Requirements

- Latency** Autonomous, robotics, AR and live vision can require metro-local response.
- Data Control** Private / regulated workloads benefit from isolated, location-specific compute.
- Capacity Timing** Customers can contract ready capacity rather than wait for greenfield interconnection.
- Burst & Overflow** Regional hubs absorb non-latency-bound load while tower sites serve the edge.

Edge Demand Specifications

Workload	Latency Budget	Location Requirement
Autonomous vehicle perception	Under 10 ms	Metro-local inference
Industrial robotics	Under 5 ms	On-site or tower-adjacent
Drone detect-and-avoid, 15 m/s+	Under 25 ms end to end	Metro-local inference
AR and spatial computing	20 ms motion-to-photon	Metro-local inference

Market \$ per GPU-hour



Tower Power

Tier 1 Partners

Owner	# of Sites	% Market Share	Sites Screened for Day 1	Available MW at 1-2 MW / site	Phase 1 MW Deployable ⁽¹⁾	Stage
	~17,000	~4%	~486	486-972 MW	217-434 MW	MSA signed
	>2,500	~1%	~192	192-384 MW	67-134 MW	Finalizing MSA
	>15,000	~4%	~429	429-858 MW	194-388 MW	To be disclosed
	[]	[]	[]	[]	[]	In evaluation
Total	>34,500	~8%	~1,107	1,107-2,214 MW	478-956 MW	—

First Contracted Site: ~1.4 MW

Measured available power at the executed site

Pass-Through

Power cost passed through at the tower rate

~717 MW

Immediately available for Phase 1

~1,661 MW

Optioned across the screened portfolio

The Excess Power Unlock



Installed Power is Already there

Each generation of radio equipment draws materially less power than the equipment it replaces, and at many sites the utility service sized for the earlier load is still in place

Contracted vs. Planned

After Phase 0, tower access sits under MSAs within a screened pipeline. Capital is deployed after a sales order is executed.

Screened Tower Portfolios

Available MW is screened sites at 1-2 MW each. Contracted and energized MW are a model output at the Phase 0 close. Phase 1 must clear 11 tests where three-phase service is a binding screen.

Why Neoclouds Do Not Do This

- Access is contracted, not purchased
- Forum Edge AI is a first mover to contracted access
- Footprint below minimum scale
- Operators built for 100MW hauls cannot cover overheads in 1-2MW bites


Why TowerCos Do Not Do This

- Different business model
- Most TowerCos are REITs or built for long-dated leases against capital-light assets. Partnering with network builders is their business and deploying capital quickly keeps tower partners happy

GPU Supply Secured

Installed cost per GPU, annual revenue per GPU, and residual value against the loan balance

Equipment Cost Bridge: Percent of All-In Cost per GPU



What One GPU Costs and Earns

All-in Cost \$88-102k Per GPU of equipment, GB300 to B300	Gross Revenue per Year ~\$38k Per GPU at modeled rates which are below contracted	As a Share of cost 38-44% Annual revenue against installed cost
---	---	---


Residual Cover of the Loan at Year 3	1.9x ~\$35k against ~\$18k
---	--------------------------------------

Supply Agreements


Vendor A: Integrated Pod Supply Quotation received for three 32-node B300 pods covering servers, support, fabric, storage, rack, PDU and installation 768 GPUs quoted at ~\$100k per GPU	Vendor B: GB300 Rack Supply Rack cost and supply established for North Carolina and the tower site: ~\$6.4m per 72-GPU rack, ~\$88k per GPU all-in
--	--

FORUM Note: cost lines are the vendor quotation as issued (GDT QT-000038431, 1 September 2026, valid to 15 September 2026) plus two Cisco firewalls per 32-node pod at \$250,000 each, divided by the GPU count; the quotation does not price the accelerator separately from the server, so no figure below the node has been allocated. Liquid cooling, freight, duty and sales and use tax are not on the quotation. Residual is straight line from cost to the modeled 20% residual over the 48-month contract and is not a market valuation. Loan at a 65% advance, six months interest-only, then level payments to month 48 at 10%. Photographs are illustrative.

What Goes In at Every Site




Residual Value Against the Loan Balance



Tower Footprints Motivate Discretely Built, National Network

Strategic start in North Carolina and Texas with outsized expansion opportunity across continental united states



30

Largest US Metros
150m people, ~45% of the US

2

Hub Markets⁽¹⁾
Dallas 4 MW · NC 6 MW

Up to 50

Sites Per Metro
Up to 50 towers +
one 5 MW hub

50-75 miles

Service Radius
1-2 ms round trip on metro
fiber

Thirty Targeted Metros

- The 30 largest MSAs hold ~150m people; the three screened tower portfolios hold their deepest tower portfolios in the same metros
- **Premium demand - autonomous fleets, industrial robotics, and AR are primarily in metro areas**

Hub and Spokes


- **The hub** handles orchestration, burst, and training-adjacent load on utility power; **the spokes** deliver inference at the tower on power already metered
- Each metro is its own cohort - separately funded and collateralized

Fifty Mile Radius

- Latency: **~1-2 ms round trip on metro fiber keeps every premium workload in budget.**
- **One field team per hub** reaches spokes inside the radius which makes **30 metros operationally manageable**

Immediate Capacity, Rapid Expansion

Phase 0 establishes first tower deployment; site additions and new hubs drive capacity beyond 100MW



Phasing Unlocks

Phase	Incremental Expected MW	Cum. Exp. MW Energized	Platform
Phase 0	11	~11	Dallas 4 MW, NC 6 MW, 1 Tower
Phase 1	51	63	NC to 9 MW, Towers to 50
Phase 2	58	120	2 Regional Hubs; Towers to 100
Total	-	120	4 Data Centers, 100 Towers

2 MW Energized Today at Dallas⁽¹⁾


Utility service available	Yes
Power behind a Forum-controlled-meter	Yes
Data hall complete	Yes
Cooling installed	Yes
Racks installed	Yes
GPUs installed and drawing load	Q1 27

FORUM 1. The site is energized and metered. The first two pods energize in January 2027, which is when load and revenue begin.

Unit Economics

Discretely scalable project deployments – equity underwriting perspective

Percent of Net Revenue⁽¹⁾



GB300 NVL72 unit economics: illustrative and intentionally conservative vs Phase 0 contracts. B300 pods (Dallas) carry ~15% higher equipment cost per GPU (~\$102k vs ~\$88k).

Model Inputs

REVENUE DRIVERS	TOWER	HUB
Contracted rate (\$ / GPU-hour)	>\$4.00	>\$4.00
Contracted share of hours	100%	100%
Uptime: hours available to bill ⁽⁴⁾	98.5%	98.5%
SLA credit allowance	1.0%	1.0%

Returns

CONFIGURATION AND RETURNS	TOWER	HUB
GB300 racks	6	37
GPUs	432	2,664
Load at the meter	~1.0 MW	~6.0 MW
Total capital at energization	\$40M	\$254M
Facility and deposits, share of capital	96%	93%
Project contribution margin	80%	81%
DSCR⁽²⁾, contract average	1.29x	1.31x
Equity invested⁽³⁾	\$1.6M	\$17.4M

Debt / Equity Mix: Phase 0

PHASE 0 CAPITAL STACK	\$M	% OF TOTAL
Equipment facility (65% of \$403M equipment)	\$262	59%
Customer deposits	\$138	31%
JV equity (Forum 51% / Edge Node 49%)	\$42	10%
Total capital	\$442	100%

Market Competitive Context

Strategic positioning to become the market leader in Edge and Inference compute provisioning

Company	Model	Market Cap / Valuation ⁽¹⁾	Scale	Edge / Inference Focus	Tower Partnership	Public
 FORUM Forum Edge AI Phase 0	Tower-anchored Edge GPUs		>10 MW in 1H27 Access to ~1,000 sites (1–2 MW per site)	✓	✓	✓
 NEBIUS	Neocloud	\$63.7 billion	10–100 MW+	✗	✗	✓
 CoreWeave	Hyperscale GPU Cloud	\$48.9 billion	100 MW+	✗	✗	✓
 boostrun	Neocloud	\$1.3 billion	20 MW+	✗	✗	✓
 Crusoe	Hyperscale GPU Cloud	~\$30.0 billion	Variable	Partial	✗	✗
 Lambda	Neocloud	~\$5.9 billion	10–50 MW	✗	✗	✗
 VULTR	Multi-cloud GPU	~\$3.5 billion	10–50 MW	✗	✗	✗
 Vapor	Carrier-neutral edge	Undisclosed	1–2 MW	✓	Carrier sites	✗