



January 20, 2026

Green Zone Capital 2025-2030 Investment Thesis & Strategy White Paper

I. Executive Summary

Green Zone Capital is an emerging Atlanta-based investment management firm led by Ahijah Ireland, focused on long-only U.S. public equities positioned in the AI-era physical infrastructure stack, the “forced-buy” components that data centers and hyperscalers must procure to scale AI in the real world. The core insight is simple: AI does not scale digitally, it scales physically. As AI transitions from experimentation to industrial deployment, the binding constraints have migrated to power availability, grid interconnect, long-lead electrical equipment, thermal management, commissioning capacity, and now, context infrastructure.

Our edge is in our repeatable proprietary **Bottleneck-to-Ticker (BTT)** screening framework which translates real-world constraints such as lead times, backlog acceleration, commissioning delays, and qualification moats, into identifiable and investable opportunities before the broader market consensus reprices value. GZC precisely compounds ownership in companies that are undervalued “scarcity specialists” with durable pricing power and backlog visibility, supported by our valuation bands, technical analysis, and drawdown controls. Green Zone Capital identifies a generational investment opportunity in "Forced Capex" spending that hyperscalers (e.g., Google) must undertake to bypass critical bottlenecks. By positioning in these enablers that hyperscalers are forced to buy, we aim to capture a scarcity premium that traditional mega-cap narratives might overlook.

Strategic Update: “Context Infrastructure” is emerging as the next bottleneck in inference, expanding forced spend beyond power and into memory, storage, and networking. This extends our thesis by adding another constrained layer to our strategy.

II. Main Thesis: Physical Constraints → Forced Spend → Re-Rating

GZC’s strategic framework is designed to avoid crowded consensus and target deep value-chain enablers that customers are forced to buy and organized around a repeatable causal chain:

1. **Identify Constraint:** What physically gates deployment, uptime, revenue realization, or time-to-market?
2. **Validate Forced Spend:** Is capex non-discretionary due to reliability, regulation, qualifications, or opportunity cost?
3. **Map Value Chain:** OEMs → subcomponent suppliers → distributors → services.
4. **Identify Mispriced Beneficiaries:** Under-covered firms with pricing power, backlog visibility, and capacity leverage.
5. **Apply Risk Controls:** Valuation bands, technical timing, position sizing, drawdowns, correlation caps.

III. Macro Analysis Foundation: Physical Constraint System

AI adoption is often framed as a software computing story but it is increasingly a infrastructure problem requiring:

- Massive reliable electric power and grid interconnection capacity
- Long-lead electrical mechanical equipment and thermal management at unprecedented densities
- High-performance memory, storage, and data-movement systems

As workloads scale, power becomes the gating factor; capacity must be secured before revenue can be realized, turning growth capex into mandatory infrastructure spend. At the system level, this creates a multi-year regime where:

- Lead times extend and backlogs accumulate
- Pricing power improves and visibility increases
- Scarcity becomes structural rather than cyclical

Market narratives historically concentrated on the "compute layer" (GPUs); however, the fundamental bottleneck for the next decade is the physical infrastructure required to sustain it.

- **Energy Constraint:** Data center power consumption is projected to scale from 415 TWh (2024) to ~945 TWh (2030).
- **Grid Crisis:** ~2,300 GW trapped in U.S. interconnection queues doubled median wait times in commercial operations.
- **Peak Load Growth:** Forecasts grew within the past 3 years (24 GW to 166 GW); ~55% linked to data center demand.



IV. Methodology: GZC “Bottleneck-to-Ticker” (BTT) Framework Criticality Scoring Index

GZC utilizes a hybrid methodology combining Perez Technological Surge Theory with Goldratt’s Theory of Constraints (ToC) to identify the specific physical components throttling system throughput. To quantify substitutability risk and identify importance of individual “forced buy” components, GZC applies a Criticality Index proprietary scoring system using 5 weighted factors:

Factor	Description	Weight
Substitutability	Availability of Tier-2 alternatives or "second-sourcing" viability for hyperscalers.	35%
Lead-Time Delta	Variance between current delivery times and historical 5-year mean (Z-Score).	25%
Capex Intensity	Total cost of components as % of total data center facility capital expenditure.	15%
Regulatory Moat	Barriers to entry due to utility standards, NERC/FERC compliance, or permitting.	15%
IP / Qualification	Time required for a hyperscaler to "certify" a new vendor (typically 12–24 months).	10%

“Forced Buy” Mechanism Concept

A “Forced Buy” tends to emerge when a component scores in the 80th percentile or higher on the Criticality Index, especially where qualification cycles lock incumbents in for extended periods. In a forced-buy environment, the primary beneficiaries are often not the most visible assemblers, but the niche suppliers controlling mission-critical components. For example, when a hyperscaler builds a \$10 billion campus, the \$300 million spent on electrical switchgear is non-discretionary; a delay in these components renders the entire multi-billion dollar investment useless. When hyperscalers face a choice between delaying a multi-billion-dollar buildout or paying up for constrained equipment, they pay up.

Bottleneck-to-Ticker (BTT) Implementation Engine Screening Rulebook:

GZC operationalizes the strategy using a **Screening + Confirmation + Valuation + Execution** pipeline.

Step 1: Market Cap Bands & Liquidity Minimums (ADV \$)

- Core Sleeve: \$2B–\$80B & \geq \$20M/day
- Alpha Sleeve: \$250M–\$10B & \geq \$5M/day
- Micro Sleeve: \$75M–\$750M & \geq \$1.5M/day

Step 2: Revenue Exposure Score (0-5)

Advance names with exposure or strong bottleneck evidence with partial exposure.

Step 3: Bottleneck Confirmation (Requires 2 of 5)

1. Explicit lead-time inflation
2. Backlog acceleration
3. Capacity expansion tied to demand
4. Pricing power (stable/expanding margins)
5. Commissioning/Installation constraint

Step 4: Valuation Bands

Define valuation “bands” by company type and only buy when valuation sits in the lower portion of its historical band (or when backlog inflects and price has not repriced).



V. Physical Stack

Phase 1 (2025-): Power, Grid, and AI Physical Stack: Power Delivery & Grid Interconnect

Power & grid constraints are the primary gating factor for AI deployment. The strategy concentrates where forced spend clusters:

- Data-center power chains (switchgear, busway, UPS, rack power)
- Grid modernization/interconnect (transformers, substations, HV gear)
- Transmission buildout and reliability upgrades

We believe this is structurally investable because physical infrastructure scarcity behaves differently than normal cycles:

- Lead-time inflation signals pricing power before revenue acceleration
- Backlogs convert to revenue with high operating leverage
- Capacity expansions lag demand, extending scarcity windows
- Substitutes are limited due to certification, reliability, and safety constraints

The primary bottleneck of the mid-2020s is the grid. The initial bottleneck of 2024-2025 was raw compute power. The bottleneck of 2026 is Power Density and Thermal Management.

- Transformers: Lead times for high-voltage transformers have surged from 6 months in 2021 to 36+ months in 2026.
- Switchgear & PDUs: These components determine time-to-revenue for hyperscalers.
- Nuclear/SMRs: GZC views Small Modular Reactors as the only viable carbon-free baseload for AI, positioning uranium suppliers as critical upstream beneficiaries.

Component Rating & Bottleneck Ranking

- Switchgear / Breakers / PDUs: 9.3 / 10 (Bottleneck: multi-year lead times and factory slot scarcity).
- High-Voltage Transformers: 9.1 / 10 (Bottleneck: extreme grid interconnection delays; 36+ month lead times).
- UPS Systems & Power Conditioning: 8.5 / 10 (Bottleneck: supply chain constraints on power electronics).
- Liquid Cooling (CDUs/Manifolds): 8.2 / 10 (Bottleneck: rapid phase-shift from air to liquid architectures).

Phase 2 (2026-): Context Infrastructure Bottleneck

Context is the new forced-spend layer as AI shifts to multi-turn agentic systems, the bottleneck moves to memory, storage, and networking. What changed: Industry disclosures (notably NVIDIA CES 2026) confirm that context is emerging as a major bottleneck in AI inference. As models expand context windows and agentic/multi-turn workflows proliferate, performance becomes constrained by the system's ability to store, retrieve, and move context efficiently.

- Memory tiering: HBM → DRAM → SSD
- AI-optimized storage platforms
- Low-latency networking fabrics (RDMA Ethernet)
- Data-movement offload & orchestration (DPUs / SmartNICs)

Why it matters: GPUs are most efficient when continuously fed data. When context spills across memory tiers:

- GPU utilization declines
- Tokens-per-second falls
- Power efficiency degrades
- Capex efficiency suffers

This forces system-level redesign, not just incremental compute additions.

This follows the logic: physical constraint → forced spend → visibility before consensus → undervalued suppliers benefit first.



VI. Portfolio Construction Strategy & Risk Management Controls

GZC employs institutional-grade discipline to manage volatility and capture alpha.

Allocation Strategy Framework

- **Core Strategy (70-80%):** Concentrated exposure to physical infrastructure bottlenecks (Power, Cooling, Grid).
- **Grid Modernization Sleeve (15-25%):** "Picks & Shovels" like transformers and substation automation.
- **Logistics/Industrial Overlay (0-15%):** Timing and conviction-weighting based on real-time industrial signals.

Execution & Risk Controls

- **Drawdown Trigger:** If the portfolio drops -9% from peak, gross risk is automatically reduced by 15-25%.
- **Exit Strategy:** Triggered by backlog normalization or valuation exceeding the top 20% of its historical band.
- **Execution Model:** Utilize Almgren-Chriss model to balance impact against volatility, ensures optimal liquidations.
- **Diversification:** Monitor correlation to avoid excessive NAV tied to one shared factor/catalyst.

Portfolio Guardrails & Risk Management

To manage volatility inherent in industrial or small-cap equities, GZC adheres to strict allocation caps and liquidity requirements:

- **Market Cap Bands:** Core Sleeve (\$2B-\$80B), Alpha Sleeve (\$250M-\$10B), Micro Sleeve (\$75M-\$750M).
- **Allocation Caps:** Power Chain (35-45% Max), Grid / Interconnect (30-40% Max), Cooling Thermal (20-30% Max)
- **Position Sizing:** Tier 1 Core (4-8%), Tier 2 Alpha (2-5%), Tier 3 Micro (0.5-2%)

Trim Exit Logic

- Trim on sharp moves without fundamental revision, valuation entering extreme band, or bottleneck easing signals.
- Exit on backlog rollover, pricing deterioration, demand assumption invalidation, or governance/dilution spiral risk.

VII. GZC G-Force Valuation Framework

To properly value companies operating within structurally constrained infrastructure, GZC applies a constraint-adjusted earnings framework that augments traditional earnings-based valuation with a disciplined supply-friction premium. This approach captures the economic reality that prolonged lead times, backlog persistence, and limited substitution materially increase pricing power and earnings durability beyond standard cycle assumptions.

Valuation Formula

GZC Fair Value Calculation: $GZC\ Fair\ Value = Forward\ EPS \times Industry\ Multiple \times Supply\ Constraint\ Multiplier$

Supply Constraint Multiplier: $Supply\ Constraint\ Multiplier = 1 + \min(Lead\ Time\ Premium, Maximum\ Constraint\ Premium)$

Lead Time Calculation: $Lead\ Time\ Premium = (Current\ Lead\ Time - Normalized\ Lead\ Time) \div Normalized\ Lead\ Time$

Rationale and Discipline

This framework ensures that valuation premiums derived from supply inelasticity remain bounded, repeatable, and cycle-aware. By explicitly capping the constraint premium, GZC avoids extrapolating temporary bottlenecks into perpetuity while still recognizing that extended lead times materially enhance pricing power, backlog visibility, and return on invested capital.

This model is applied selectively to sectors where: Demand exceeds installable capacity on a multi-year basis, substitution risk is limited or nonexistent, and capital expansion timelines materially lag demand growth.

The G-Force Valuation Framework serves as a core input to position sizing, trim discipline, and risk-adjusted return forecasting across GZC's infrastructure-focused portfolios.



VIII. Implementation Standards

Because parts of this strategy can include under-covered names, execution discipline matters. GZC's approach explicitly integrates institutional execution, liquidity constraints, and treats risk management as structural, not reactive. GZC also emphasizes an institutional communication posture aligned with best practices around material information handling and public disclosures (e.g., Reg FD-oriented discipline referenced in the research framing).

Research Signal Evidence

GZC focuses on **leading indicators** that typically precede consensus re-rating:

- **Lead-time Commentary** and procurement bottlenecks
- **Backlog / RPO trends**, book-to-bill where applicable
- **Margin Resilience** (pricing power vs input volatility)
- **Commissioning Capacity** and services attach rates
- **Grid Data**: interconnection backlog, permitting, substation upgrades, regional load hotspots

Core arbitrage constraints become obvious only after they appear in earnings reports, the objective is to position before re-ratings.

Scenario Analysis Range

GZC frames outcomes based on (1) constraint persistence and (2) hyperscaler capex trajectory. We model three forward-looking scenarios for the "Physical Buildout" theme to guide portfolio positioning and risk controls:

1. Bull Case: AI agent adoption accelerates and hyperscaler capex expands above trend. Grid relief remains gradual, extending key bottlenecks into the early-2030s. **Result:** Scarcity premiums persist across power delivery, commissioning capacity, thermal management, and context infrastructure. **Target:** >35% gross annualized.

2. Base Case: Hyperscaler capex grows at a steady run-rate ($\approx 15\text{-}25\%$ YoY), with power scarcity and context infrastructure remaining binding constraints. **Result:** Forced-spend visibility and durable pricing power drive sustained re-rating for constrained suppliers. **Target:** 29-33% gross annualized.

3. Bear Case: Permitting delays, deployment friction, or policy headwinds slow the pace of buildout and shift spending toward efficiency/optimization. **Result:** Returns driven primarily by quality, cash-flow resilience, valuation discipline, and drawdown control. **Target:** 12-15% gross annualized.

Note: Return targets are forward-looking objectives (gross of fees/expenses) and are not guarantees; actual results may differ materially.

Why GZC Is Structurally Differentiated

GZC is not: A generic AI fund, a momentum tech strategy, or a passive infrastructure allocator. GZC is a scarcity-focused research institution that allocates capital using: Real-world bottleneck intelligence, early identification of forced spend, valuation execution discipline, and willingness to concentrate where constraints are undeniable.

IX. Conclusion

The next decade of AI-driven growth will not be won by those who only chase computing power. It will be won by those who understand where AI physically breaks and who owns the constrained solutions when it does. GZC is designed to be a specialist.

The binding constraints on the AI revolution have shifted from silicon to physics. By identifying the companies that control the "Forced Buy" mechanisms in grid access, thermal management, and memory substrates, GZC captures alpha that is structurally protected from competition. Our "Bottleneck-to-Ticker" framework provides the repeatable methodology required to navigate this transition and capture the multiple expansion that follows the realization of physical scarcity.