METISPOWER

DATA CENTER SERIES

Behind the Meter

POWER GENERATION FOOTPRINT

Assessing Key Factors Impacting the Data Center Operations

The Footprint Factor: Navigating Space Constraints in Behind-the-Meter Power Generation

Imagine a sprawling data center humming with activity, or a critical manufacturing facility operating around the clock. These operations share a common, non-negotiable need: Increasingly, organizations are turning to Behind-the-Meter (BTM) power generation—building their own on-site power plants—to reduce energy costs, gain control over power quality, enhance reliability against grid failures, and bypass lengthy grid interconnection queues. It's a move towards energy independence. But as facilities embrace BTM solutions, they often encounter a hidden challenge lurking beneath the surface of MW capacity and efficiency ratings: the physical space required. The footprint of the power generation system is not just a detail; it's a fundamental constraint that can significantly impact project feasibility, cost, and future flexibility. This paper delves into the critical nature of footprint in BTM deployments, exploring how design choices, specifically the size of modular power blocks and the level of redundancy, shape the spatial demands of these vital systems.

1. The Premium on Space: Why Footprint Matters

In today's world, space often comes at a premium. Whether it's expensive urban real estate, a tightly packed industrial site, or a data center floor meticulously planned for servers, finding room for a multi-megawatt power plant can be a significant hurdle. The footprint dictates:

- Capital Costs: Larger systems require more land or building space, directly inflating initial investment.
- Site Integration: Fitting generation equipment into existing facility layouts can become a complex, costly jigsaw puzzle.
- Permitting & Compliance: Zoning laws and local regulations frequently impose restrictions based on equipment size and density.
- Future Growth: A poorly optimized initial footprint can box in future expansion possibilities.

Minimizing the footprint, therefore, isn't just about tidiness; it's about maximizing resource efficiency and ensuring the long-term viability of the BTM investment.

2. Modeling the Building Blocks: A Single Unit Perspective

Modern BTM power plants are often constructed using modular "power blocks." To understand the total system footprint, we first need a model for the footprint of a single block relative to its capacity. Our analysis employs a constrained linear model, visualized in Figure 1 (Estimated Single Block Footprint Trend).

This model, derived from analyzing sample industry data and applying a specific design constraint, assumes a baseline footprint (anchored at 500 sq ft/MW in our model) and then projects a linear increase in space requirements as the block's capacity grows (Footprint ≈ m * Power (MW) + c). While real-world relationships might exhibit more complexity, this constrained linear trend provides a starting point. It represents our foundational assumption: bigger individual blocks need more space, following a defined, predictable slope after an initial minimum size. This single-block model is the basis upon which we build our understanding of the entire system.

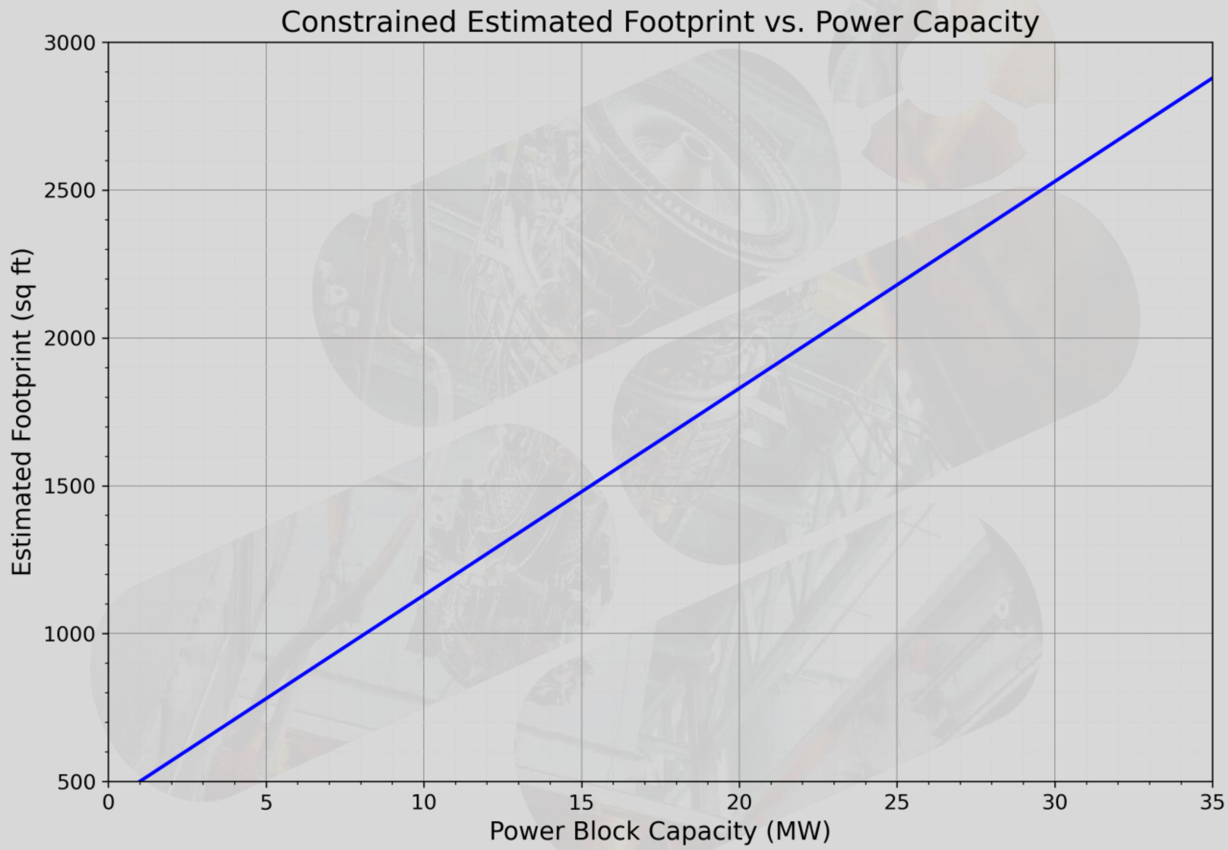

Figure 1- Estimated Footprint vs. Power Generation Capacity

3. The Foundational Principle: Impact of Block Size without Redundancy

Now, let's consider the bigger picture: deploying enough blocks to meet a facility load of exactly 100 MW. In Figure 2, the solid black line (N Config) represents this baseline scenario, plotting the total system footprint against the capacity of the individual power blocks chosen (from 5 MW to 35 MW).

At first glance, the overall downward trend of the black line confirms an intuitive principle. However, a closer look reveals two distinct behaviors:

- The Overall Trend (The Advantage of Larger Blocks): As we move from left to right, using fewer, larger blocks (e.g., 20, 25, or 30 MW) dramatically reduces the total required footprint compared to using a large number of smaller blocks (e.g., 5 MW). While individual units are bigger, the sharp reduction in the quantity of blocks, foundations, and auxiliary equipment needed more than compensates for their size.
- The "Sawtooth" Effect: The jagged, sawtooth pattern is caused by step-changes in the number of units (N) required to meet the 100 MW load. For instance, the footprint decreases steadily as block size increases from 26 MW to 33 MW (where N=4). But at 34 MW, only three units are needed (N=3), causing the total footprint to "reset" at a new, lower starting point for that configuration.

The clear message from this baseline analysis is that, when not accounting for redundancy, using the largest possible power blocks to minimize the number of units results in the smallest total footprint. This straightforward finding provides a crucial foundation for the more complex analysis of a resilient system.

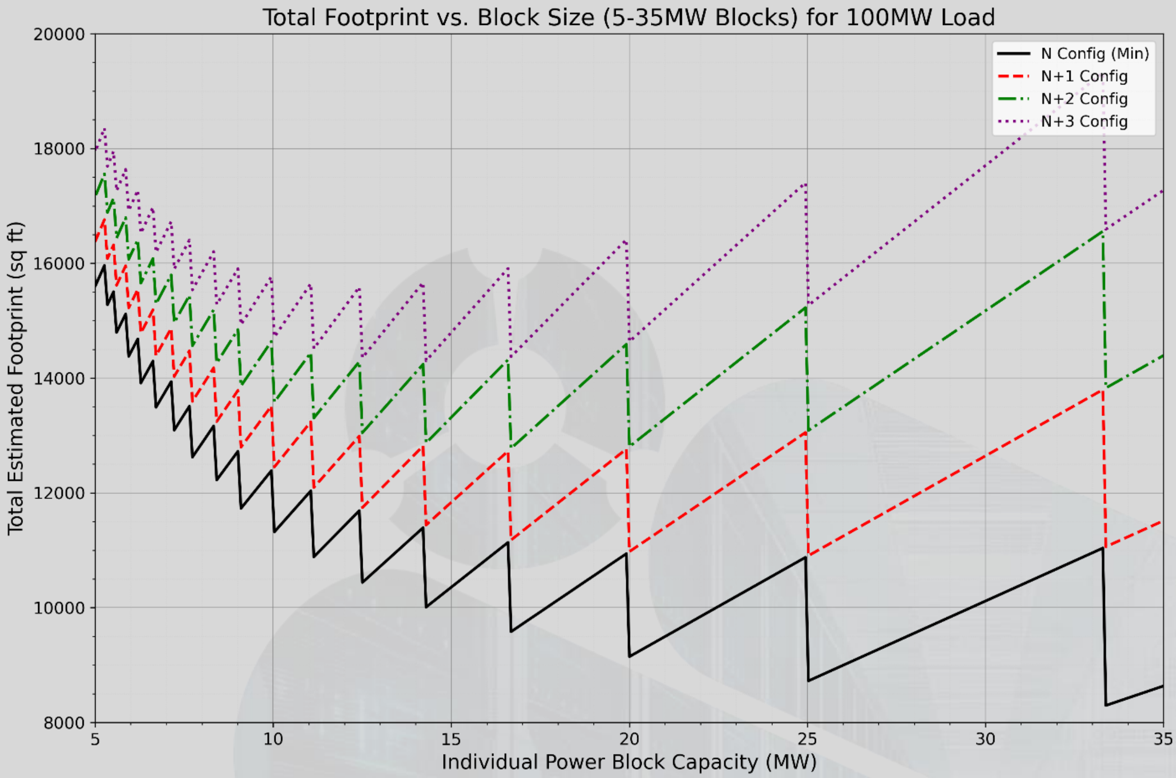

Figure 2 - Total Estimated BTM Footprint vs. Block Size for 100MW Load

4. The Price of Reliability: Footprint and N+X Redundancy

Meeting a target load is only the first step; ensuring uninterrupted power requires designing for reliability through redundancy. The N+X methodology is the industry standard for this, ensuring the facility remains fully powered even if 'X' blocks are offline. For a comprehensive discussion on the critical role of reliability frameworks, **readers are encouraged to consult the METIS Power white paper, "Behind the Meter POWER GENERATION RELIABILITY."**

While it's logical that adding redundant blocks increases the required space, Figure 2 reveals a more surprising and nuanced relationship between block size and the true cost of reliability.

- The Obvious Cost: The red (N+1), green (N+2), and purple (N+3) lines vividly illustrate the direct footprint penalty of redundancy. Each successive line sits higher than the last, confirming that every layer of resilience demands a larger physical installation.
- The Surprising Trade-Off: The key insight lies in how redundancy interacts with block size. The intuitive conclusion from our first analysis—that bigger is always better—is now challenged. For example, in an N+2 configuration, choosing very large blocks (e.g., 30 MW) means the two redundant units add a massive footprint penalty. In contrast, a moderately sized block might result in a lower overall redundant footprint, even if its baseline N configuration was not the absolute minimum. The optimal block size for an N system is not necessarily the optimal size for an N+2 or N+3 system.

This visual evidence proves that designing for high reliability is not as simple as adding extra units. Planners must weigh the operational value of redundancy against a more complex reality: the choice of individual block size creates a critical trade-off that can either minimize or multiply the physical and financial cost of resilience.

5. Beyond the Graphs: A Holistic View

While this study identifies the optimum size of power blocks based on idealized footprint estimation, it is essential to acknowledge that real-world project decisions involve a multitude of factors. In practice, the final selection of a power generation solution is influenced by a comprehensive analysis of economic viability, environmental regulations, integration with hybrid power blocks, site-specific load profiles, deployment timelines, seasonal demand variability, local permitting (especially air permitting), and social acceptance. These practical considerations can often lead to selections that differ from a purely theoretical optimum based on a single variable.

To aid in this comprehensive evaluation, METIS Power has published a series of in-depth white papers. We encourage readers to explore these resources for a deeper analysis of key decision-making areas:

- For Capital Cost Analysis: Behind the Meter Power Generation Capital Cost for Data Centers
- For System Reliability: Behind the Meter Power Generation Reliability for Data Centers
- For Environmental & Emissions Performance: Behind the Meter Power Generation Emissions for Data Centers

By consulting these companion documents, stakeholders can build a more complete picture, which can help align their final power generation strategy with key project variables.

6. Conclusion: Designing for Power, Planning for Space

The journey to energy independence and resilience via Behind-the-Meter generation is paved with critical design choices. As facilities demand more power, the physical space available becomes an increasingly valuable commodity. Our analysis demonstrates that the relationship between power block size, redundancy, and total system footprint is nuanced and crucial.

Figure 1 established our baseline model for a single block's footprint. Building on this, Figure 2 illustrates two fundamental dynamics. First, for a non-redundant system, it reveals a clear trend: using fewer, larger power blocks significantly reduces the total footprint. Second, it quantifies the spatial cost of reliability, showing that each level of N+X redundancy progressively increases the required space.

Taken together, these factors lead to a compelling, counter-intuitive conclusion: the optimal block size for a simple, non-redundant system is often not the most space-efficient choice once the significant footprint of a high-resiliency design is factored in.

Successfully deploying BTM generation requires a strategic balancing act. Planners must navigate the trade-offs between the desired level of reliability, the operational efficiencies of different configurations, capital costs, and the very real constraints of the physical footprint. By understanding the dynamics illustrated here, stakeholders can make more informed decisions, optimizing their BTM power systems not just for megawatts, but for the valuable square footage they occupy.

METIS POWER

Powering What's Next ...

12006000220