San Van Flexscale Advanced Data Center Solution

Enable Rapid Construction Capability of Data Centers for the AI Era

San Van HPC

The AI Infrastructure Challenge

Today's enterprise AI initiatives face a critical bottleneck: **traditional data centers can't keep pace with exploding compute demand.** Hyperscalers and enterprise leaders alike are struggling with:

Time-to-Market Pressure

12-24 month buildout timelines for traditional data centers severely limit AI deployment schedules

Power & Cooling Constraints

Legacy infrastructure simply cannot handle the intensive power and cooling requirements of today's AI accelerators

Capital Inefficiency

Fixed, monolithic designs lead to overcapacity in some areas while bottlenecking others

Sustainability Concerns

Traditional data centers struggle to meet modern efficiency requirements, creating conflict with corporate ESG goals

The need for a fundamentally different approach has never been more urgent as AI workloads continue to grow exponentially.

Introducing: San Van Flexscale Advanced Data Center Solution

SVFADS is a revolutionary approach to data center infrastructure—**entirely modular and fully prefabricated** to meet the demanding requirements of modern AI workloads:

- Complete turnkey AI data center delivered in **under 3 months**
- Supports both air-cooled and liquid-cooled deployments up to **190kW per rack**
- Industry-leading **PUE < 1.1** for exceptional energy efficiency
- Uptime **Tier 3+** reliability standards
- Integrated support facilities including monitoring halls, maintenance offices, conference rooms, and staff living spaces

Our modular approach transforms every aspect of the data center development process—from design to deployment—with a factory-prefabricated system that dramatically reduces time-to-market while enhancing quality, efficiency, and flexibility.

> "The San Van Flexscale solution represents a paradigm shift in data center construction, enabling enterprises to deploy AI infrastructure at the speed their business demands."

Core Advantages: Why Flexscale Outperforms Traditional Approaches

Rapid Deployment

While traditional data centers require 12-24 months to build, Flexscale modules can be assembled on-site within **just 3 months**, reducing time-to-market by up to 70%.

Superior Energy Efficiency

Factory-optimized systems deliver industry-leading **PUE < 1.1**, substantially reducing both operating costs and carbon footprint compared to typical enterprise data centers (PUE 1.5-2.0).

Extreme Power Density

Purpose-built for AI workloads, supporting up to **190kW per rack**—sufficient for the most demanding GPU clusters and AI accelerators on the market today and tomorrow.

Flexible Scalability

Modular architecture enables **precise capacity matching** to workload requirements, with the ability to expand incrementally without disrupting existing operations.

Complete Functionality

Includes not just computing space but **all supporting facilities**—from power and cooling to offices, monitoring centers, and even staff living quarters—in a unified, integrated design.

These advantages combine to create a solution that fundamentally transforms the economics and capabilities of enterprise AI infrastructure.

The Flexscale Process: From Factory to Operation

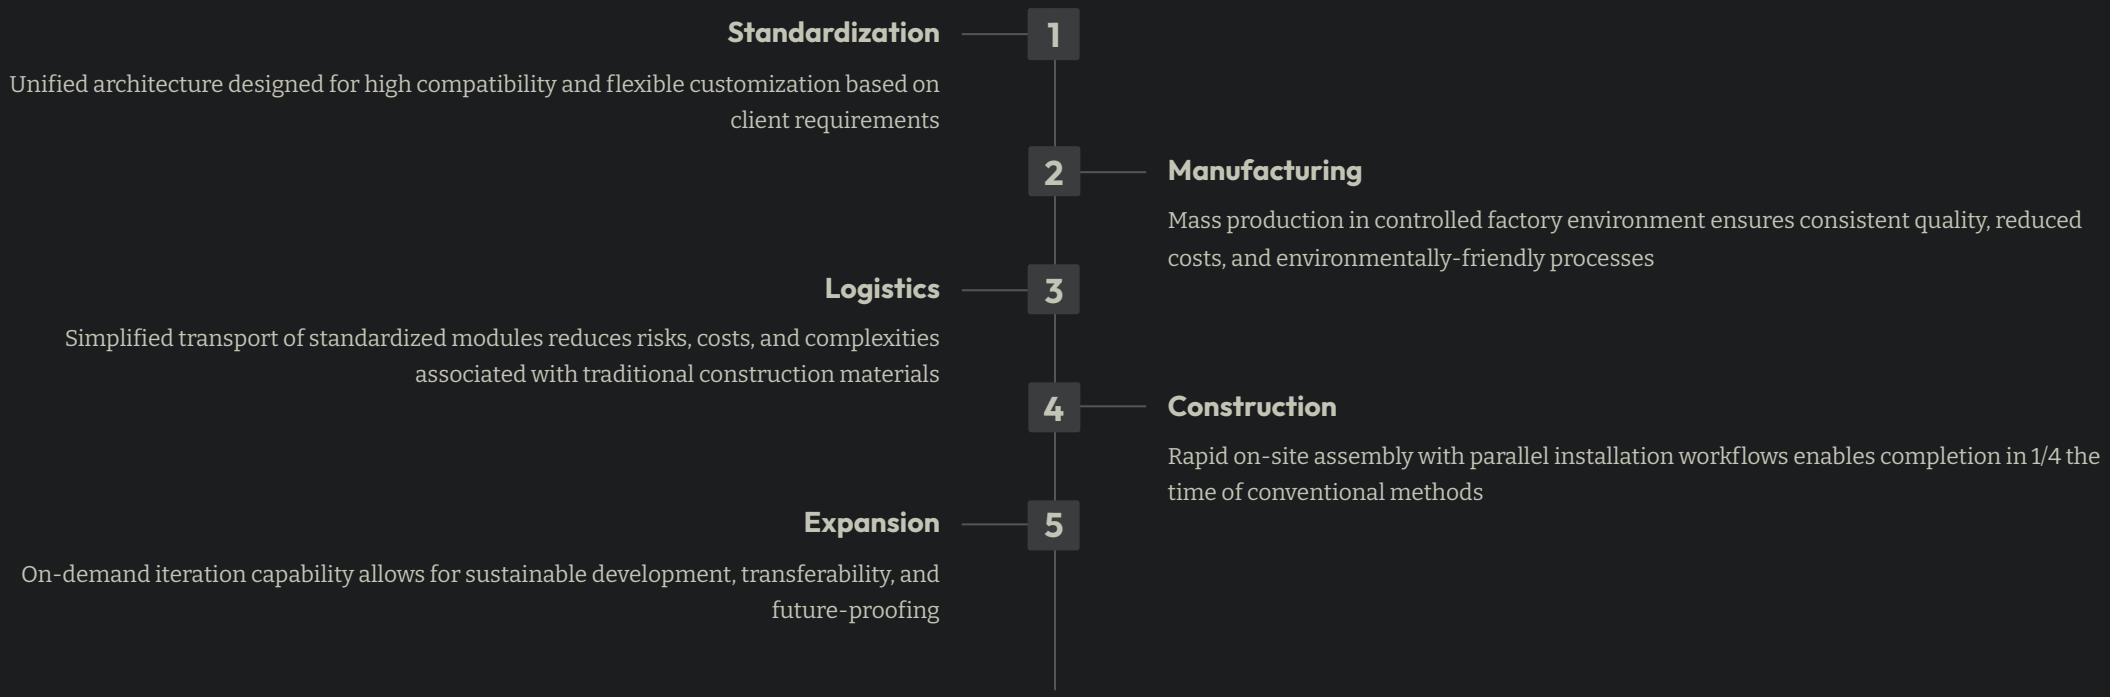

Key Differentiators:

- **Quality Control:** Factory testing of all components before shipping minimizes on-site issues
- **Parallel Workflows:** Site preparation occurs simultaneously with module manufacturing
- **Reduced Labor Requirements:** Assembly rather than construction reduces on-site workforce needs by up to 60%
- **Weather Independence:** Factory production continues regardless of site conditions

Our factory-based manufacturing process ensures consistent quality while dramatically accelerating deployment timelines.

Modular Architecture: Building Blocks of the AI Data Center

Power Unit

Complete power infrastructure including:

- Generators
- UPS systems
- Dry transformers
- High/low voltage cabinets
- Power distribution

Support Unit

Operational facilities including:

- Meeting/office spaces
- Reception/exhibition areas
- Restrooms/showers
- Dormitories
- Kitchens
- Activity spaces

HVAC Unit

Advanced cooling systems featuring:

- Pump-driven refrigeration
- Backplate air conditioning
- Primary circulation pumps
- Cooling Distribution Units
- Water treatment systems
- Cooling towers

Rack Unit

Computing infrastructure including:

- High-density server cabinets
- Bridge systems
- Communication systems
- In-row CDUs
- In-row PDUs
- Optimized airflow management

These modular components can be scaled, coordinated, and combined to form optimal computing power modules tailored to specific workload requirements. The standardized interfaces between modules ensure seamless integration while maintaining the flexibility to adapt to changing needs.

Financial Comparison: Flexscale vs. Traditional Buildouts

Aspect	San Van Flexscale	Traditional AI HPC
Cost per MW	**$7-10 million**	$12-20 million
Construction Timeline	**3 months**	12-24 months
CapEx	**Up to 50% lower** due to prefabrication and modularity	Higher due to on-site labor, custom engineering, and delays
OpEx	**Lower energy costs** (PUE <1.1), reduced maintenance	Higher ongoing costs from inefficiencies
ROI Timeline	**Faster payback** (6-12 months post-deployment)	Slower (18-36 months)

Total Cost of Ownership Analysis

Clients can achieve **30-50% total cost savings on per megawatt buildout alone** by opting for Flexscale, with the most significant advantages in:

- Regions with high labor costs or regulatory complexities
- Deployments requiring rapid time-to-market for competitive advantage
- Organizations with fluctuating or unpredictable capacity needs
- Enterprises with strong sustainability commitments

For a customized financial model based on your specific requirements, our team can develop detailed projections that account for regional variations, scale considerations, and your organization's unique priorities.

Deployment Process: From Concept to Operation in 90 Days

Initial Consultation

Week 1

- Site assessment
- Power requirements analysis
- Cooling specifications
- Capacity planning
- Initial budget framework

Design Customization

Weeks 2-4

- Module selection
- Integration planning
- Site preparation details
- Regulatory compliance
- Final specifications

Factory Prefabrication

Months 1-2

- Component manufacturing
- System assembly
- Quality control testing
- Performance validation
- Shipping preparation

On-Site Assembly

Month 3

- Module transportation
- Foundation installation
- Module placement
- System connections
- External finishing

Testing & Commissioning

Week 12

- Full system checks
- Performance validation
- Hyperscaler integration
- Security protocols
- Handover procedures

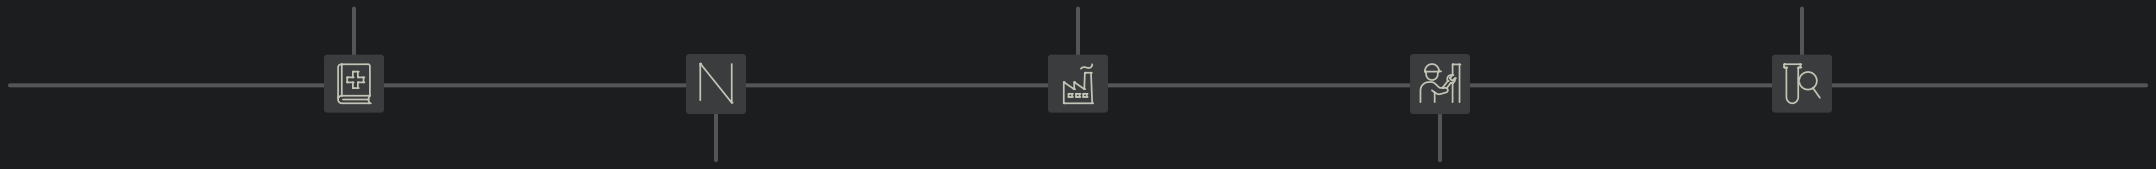

This streamlined process minimizes common risks faced in traditional data center construction, including supply chain delays, weather disruptions, labor shortages, and coordination challenges between multiple contractors. The result is predictable timelines, consistent quality, and significantly reduced project risk.

Sustainability and Innovation Features

Environmental Leadership

San Van Flexscale emphasizes green AI infrastructure through multiple innovations:

Industry-Leading Efficiency

PUE <1.1 reduces carbon footprint by 20-30% compared to traditional data centers, translating to thousands of metric tons of CO_2 savings annually per facility

Sustainable Materials

Factory processes utilize recyclable components and precision manufacturing to minimize waste, with up to 95% of materials eligible for eventual recycling

Water Conservation

Advanced cooling systems reduce water usage by up to 80% compared to traditional cooling towers through closed-loop designs and efficiency measures

Next-Generation Technologies

AI-Driven Management

Integrated environmental control systems use machine learning to optimize operations in real-time, predicting maintenance needs before failures occur

Geopolitical Adaptability

Flexscale's design ensures compliance with diverse regulatory environments, including U.S. export controls and regional data sovereignty requirements

Renewable Integration

Optional solar integration packages and green energy procurement support align with hyperscaler sustainability mandates and corporate net-zero goals

Technical Deep Dive: Power and Cooling Innovations

Advanced Power Infrastructure

Flexscale's power distribution system is engineered specifically for the unique demands of AI workloads:

- **N+1 Redundancy:** All critical power components maintain N+1 redundancy while minimizing excess capacity costs
- **Dynamic Load Balancing:** AI-driven power management adjusts in real-time to workload fluctuations
- **Scalable Power Delivery:** From 50kW to 190kW per rack without infrastructure modifications
- **Voltage Optimization:** Factory-calibrated systems reduce conversion losses by up to 15%
- **Predictive Analytics:** ML-based systems forecast power needs and preemptively address potential issues

> "The Flexscale power architecture eliminated the overprovisioning we typically faced with traditional data centers, reducing our capital expenditure by 32% while actually improving reliability."

Next-Generation Cooling Solutions

Supporting extreme compute density requires innovative cooling approaches:

- **Hybrid Cooling:** Seamless integration of air and liquid cooling within the same facility
- **Direct-to-Chip:** Advanced liquid cooling supports up to 190kW per rack for the most demanding AI accelerators
- **Closed-Loop Efficiency:** Water-side economization reduces mechanical cooling requirements by up to 70%
- **Precision Control:** Granular temperature and humidity management at the rack level
- **Minimal Footprint:** In-row CDUs maximize compute space while ensuring optimal cooling delivery

Beyond Infrastructure: The Complete AI Ecosystem

Software Integration

Flexscale includes comprehensive software integration capabilities:

- **Hyperscaler Ready:** Pre-configured for seamless integration with AWS, Google Cloud, Azure, and Alibaba Cloud
- **AI Frameworks:** Optimized environments for TensorFlow, PyTorch, and other leading ML frameworks
- **DCIM Excellence:** Advanced Data Center Infrastructure Management with predictive analytics
- **API-Driven:** Programmable infrastructure supporting Infrastructure as Code methodologies

Security & Compliance

Enterprise-grade security built into every layer:

- **Physical Security:** Multi-factor authentication, biometrics, and comprehensive monitoring
- **Regulatory Compliance:** Designs adaptable to PCI-DSS, HIPAA, GDPR, and other frameworks
- **Supply Chain Integrity:** Verifiable component sourcing and secure manufacturing processes
- **Zero Trust Architecture:** Security principles embedded throughout the infrastructure

Professional Services

Comprehensive support for the entire lifecycle:

- **Capacity Planning:** Expert guidance on right-sizing infrastructure for AI workloads
- **Migration Support:** Seamless transition from existing environments to Flexscale
- **24/7 Operations:** Optional fully-managed services with guaranteed SLAs
- **Training & Knowledge Transfer:** Comprehensive programs for your operations team

The San Van approach goes beyond simply providing physical infrastructure—we deliver a comprehensive ecosystem that supports your entire AI journey from initial deployment through ongoing operations and future expansion.

The Road to AI Infrastructure Excellence

info@sanvanhpc.com

pwl@sanvanhpc.com

adam@sanvanhpc.com

Taking the Next Step

As AI continues to transform every industry, the organizations that succeed will be those with infrastructure that can keep pace with rapidly evolving compute demands. San Van Flexscale provides the agility, efficiency, and scalability needed to stay ahead in this dynamic landscape.

We invite you to explore how our solution can address your specific challenges through:

Schedule a Demo

Experience Flexscale capabilities through a virtual or in-person demonstration at one of our showcase installations

Request a Site Assessment

Our technical team can evaluate your current and planned locations to determine optimal Flexscale configurations

Engage in Capacity Planning

Work with our AI infrastructure experts to develop a tailored roadmap for your organization's specific compute needs

The future of AI infrastructure is modular, efficient, and deployed at the speed of innovation—let us show you how San Van Flexscale can transform your approach to AI data center strategy.