International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

# AI-Integrated Manufacturing Systems for Thermal Energy Storage Tank Production in Hyperscale Data Center Infrastructure

**Sachin Pal Singh**

**Abstract:** The rapid proliferation of artificial intelligence (AI)-driven hyperscale data centers has intensified thermal management demands to levels that conventional chilled water systems cannot efficiently address. Thermal energy storage (TES) systems — large stratified chilled water tanks enabling peak-load shifting and demand reduction — have emerged as critical infrastructure for these facilities. This paper examines how Smith Industries deployed an integrated AI toolkit to transform its manufacturing operations for large-scale TES tank production, covering digital travelers, automated bill-of-materials generation, real-time production dashboards, and predictive bottleneck identification. Operational outcomes include a 20–40% chiller runtime reduction in deployed facilities, a 50–70% reduction in on-site fabrication time through modular delivery, and production quality rejection rates below 2%. The framework demonstrates that AI-driven manufacturing integration is not merely a productivity enhancement but a precondition for the scalable, high-quality production required to meet surging data center thermal infrastructure demand.

***Keywords**: AI manufacturing toolkit, Chilled water thermal storage, Data center cooling, Hyperscale infrastructure, Industry 4.0, Modular fabrication, Peak demand management*

## 1. Introduction

Global data center electricity consumption reached approximately 415 terawatt hours (TWh) in 2024 — about 1.5% of worldwide electricity consumption — and the International Energy Agency (IEA) projects this figure will nearly double to 945 TWh by 2030, with AI-driven workloads accounting for the majority of that acceleration at an annual growth rate of 30% for accelerated server infrastructure (IEA, 2025). In the United States, data centers consumed approximately 176 TWh in 2023 (4.4% of national electricity) and are forecast to reach 325–580 TWh by 2028, representing up to 12% of total national electricity consumption (DOE, 2024). These figures reflect a fundamental shift in the data center industry's thermal character: AI accelerator clusters — graphics processing units (GPUs) and tensor processing units (TPUs) running large language models and deep learning inference workloads — generate rack power densities exceeding 100 kilowatts (kW) in current deployments, compared to 5–15 kW in conventional server racks. Cooling systems, which already account for 30–40% of total data center energy consumption in conventional facilities, face an order-of-magnitude increase in heat removal requirements as these densities rise.

Thermal energy storage (TES) systems — specifically, large stratified chilled water tanks that store precooled water during off-peak electrical demand periods and discharge it during peak cooling periods — have emerged as a proven, scalable solution for this challenge. TES systems enable load shifting: by operating chillers at full capacity during low-cost, low-demand overnight periods and discharging stored cooling capacity during peak afternoon and evening periods, TES systems reduce chiller runtime by 20–40% during peaks, reduce utility demand charges through peak demand reduction, and improve overall Power Usage Effectiveness (PUE) — the industry-standard ratio of total facility power to information technology (IT) equipment power. Machine learning techniques are also used to characterize the thermocline performance of stratified TES

---

*Independent researcher, USA*

tanks. For example, Alzahrani et al. (2022) predicted thermocline thickness in a district cooling TES tank using Artificial Neural Network (ANN), Support Vector Machine (SVM) and K-Nearest Neighbor (KNN) algorithms based on data from the sensors. The study found that the KNN model obtained a 96.3% prediction accuracy. These techniques can be adapted for use in the operational optimization of data center TES.

Recognizing a rapid growth in data center rollouts and a corresponding need for thermal energy storage (TES) systems, Smith Industries repurposed its existing heavy steel fabrication facility with greater than 25000 tons (metric) annual capacity, automated welding, and pressure vessel manufacturing capabilities to manufacture a purpose engineered family of modular chilled water TES tanks for hyperscale data centers. The strategic differentiator in this product line was not merely the fabrication capability but the integration of an AI toolkit — a comprehensive digital manufacturing platform — that transformed Smith Industries' production operations into a data-driven, continuously optimized system capable of delivering consistent quality at the throughput rates required to meet surging demand.

The primary contributions of this paper are: (1) a documented architecture of Smith Industries' AI toolkit as applied to TES tank manufacturing, including digital travelers, automated bill-of-materials (BOM) generation, real-time production dashboards, and predictive bottleneck identification; (2) quantified manufacturing and deployment performance outcomes from the TES tank product line, including production rejection rates, modular delivery time savings, and deployed system efficiency metrics; and (3) a framework interpreting AI manufacturing integration as a precondition rather than merely an enhancement for scalable high-quality production in rapidly growing infrastructure markets. The paper is organized as follows: Section 2 reviews the relevant literature; Section 3 frames the problem context; Section 4 describes Smith Industries' TES tank product and manufacturing platform; Section 5 details the AI toolkit architecture; Section 6 presents the system engineering and quality methodology; Section 7 documents quantitative outcomes; and Section 8 concludes.

## 2. Literature Review

### 2.1 AI and Machine Learning in Thermal Energy Storage Optimization

Artificial intelligence and machine learning are especially being used in prediction, optimization and control of three thermal energy storage subdomains. Alzahrani et al. (2022) conducted the first machine-learning-based prediction of stratified TES tank performance. In doing so, they showed that thermocline thickness was the key parameter affecting TES system performance, and the sensors could accurately predict it using three different methods with 96% accuracy. The methods enable predictive control in real-time, and are more accurate than deterministic models in systems where operations vary. The findings for KNN, ANN, and SVM applied to the actual operation of a district cooling plant illustrate that ML models generalize over different TES operating scenarios, without requiring re-parameterization, a critical feature, as hyperscale data center facilities often have different building load profiles.

Reinforcement learning is the dominant model for data center cooling optimization in real-time. When optimizing chilled water, supply air, and cooling tower setpoints with DRL, a 12% reduction in PUE from 1.51 to 1.33 can be achieved when compared to manual control methods used by operators (Patki et al. 2024). Reviewing the literature on successful applications of reinforcement learning to the problem of energy efficiency in data centers, it has been established that, from the list of peer-reviewed papers published between 2020 and 2024, 65 contain successful results with respect to energy efficiency, thermal stability, or both (Afzal et al., 2025). In practice, the AI based optimization of cooling systems, including TES integrated chilled water plants, produces consistently measurable results.

### 2.2 Industry 4.0 in Heavy Manufacturing and Fabrication

The application of Industry 4.0 principles — digital integration, automated data capture, predictive analytics, and real-time production visibility — to heavy manufacturing environments such as pressure vessel and structural steel fabrication represents a distinct domain from the high-volume discrete manufacturing contexts that dominate the smart manufacturing literature. In their initial paper on digital twin technology for industrial systems, Tao et al. (2019) identified bidirectional data fidelity as the key enabling capability for fully developed digital manufacturing integration, allowing proactive rather than reactive quality, schedule, and resource management between virtual models and physical production systems. Automated EBOM-to-MBOM consistency, which can reduce material specification errors by as much as 35% when applied to complex products, is an example of bidirectional data fidelity (Wang et al. 2022). Indeed, given the TES tank production case study where material grade compliance is deemed a structural and regulatory requirement of the end product.

The concept of an integrated AI toolkit — a unified digital platform connecting engineering data, production scheduling, quality management, and performance analytics — extends these individual capabilities into a system-level integration that matches the operational complexity of engineer-to-order heavy fabrication. Mittal et al.'s (2019) characterization of smart manufacturing identified real-time production visibility, automated quality monitoring, and integrated decision support as the foundational capabilities that distinguish smart manufacturing from digitally augmented conventional manufacturing. The AI toolkit implemented at Smith Industries applies these principles to TES tank production at an architecture and operational detail level that the existing literature has not previously documented for this application domain.

### 2.3 Modular Fabrication for Data Center Infrastructure

The adoption of modular, off-site fabrication strategies for data center infrastructure components — including cooling systems, power distribution units, and structural elements — has been driven by two pressures: the urgency of data center construction timelines, which increasingly outpace conventional site-built construction schedules, and the quality consistency advantages of controlled manufacturing environments relative to field assembly under variable site conditions. Research on modular construction in comparable application domains has documented on-site construction time reductions of 50–70% for modularly delivered systems relative to field-built equivalents, along with quality improvement through factory-controlled processes (Peiris et al., 2023). For TES tanks specifically, the structural and thermal performance of the tank is largely determined by fabrication quality — weld integrity, insulation installation precision, and diffuser geometry — making the quality consistency of controlled manufacturing environments directly consequential for deployed system performance.

## 3. Problem Context: The TES Tank Manufacturing Challenge

The challenge that Smith Industries' AI toolkit was designed to address is the convergence of three simultaneous demands: rapidly growing order volumes driven by hyperscale data center construction; stringent structural and performance quality requirements imposed by ASME Boiler and Pressure Vessel Code (BPVC) Section VIII compliance and data center-specific thermal performance specifications; and a product architecture — large-volume stratified chilled water tanks with custom interfaces for each facility's chilled water loop configuration — that creates an engineer-to-order production environment in which each tank order involves unique design parameters, material specifications, and interface requirements. This combination — high volume, high quality, high customization — represents a manufacturing challenge that conventional production management approaches address poorly: high customization typically implies low volume and manual coordination, while high volume typically implies standardized products with automated scheduling. The AI toolkit bridges this apparent contradiction by automating the coordination and quality functions that customization makes complex, enabling high-volume throughput on a high-mix production portfolio.

The TES tank product itself reflects the thermal engineering requirements of stratified chilled water storage at data center scale. The tanks are designed with internal diffuser assemblies that maintain thermal stratification — a stable temperature gradient between cold stored water (typically 4–7°C) and warmer return water (typically 12–16°C) — by distributing flow uniformly across the tank cross-section at low velocities that minimize turbulence-driven mixing. This stratification is the thermodynamic mechanism by which the tank's storage capacity is maintained: a well-stratified tank can recover 85–95% of its stored cooling capacity during discharge, while a poorly stratified tank may recover only 60–70%, significantly reducing the peak demand reduction benefit. The fabrication quality requirements that ensure good stratification — precise diffuser geometry, leak-free welds at all internal fittings, consistent insulation application — must be maintained across a production environment that is simultaneously managing multiple active orders at different stages of completion. Table 1 summarizes the key TES tank design parameters and their fabrication quality dependencies.

**Table 1. TES Tank Design Parameters and Fabrication Quality Dependencies**

| Design Parameter | Function | Fabrication Quality Dependency | Consequence of Non-Conformance |
|---|---|---|---|
| Diffuser geometry | Uniform low-velocity flow distribution | CNC precision cutting of diffuser plates | Thermocline degradation; reduced storage efficiency |
| Shell weld integrity | Leak-free pressure boundary | Automated + UT-verified seam welds | Operational leak; system downtime |
| Insulation uniformity | Minimize thermal losses to ambient | Station-based insulation inspection | Increased standby heat gain; reduced off-peak efficiency |
| Nozzle/interface geometry | Correct flow rates at design velocity | CNC pipe and nozzle preparation | Turbulent mixing; thermocline destruction |
| Tank internal cleanliness | No debris in cooling circuit | Pre-fill inspection + hydrostatic test | Chilled water loop contamination |
| ASME BPVC compliance | Code stamp for pressure-bearing service | Material certs + weld records + NDT | Non-deliverable product; regulatory hold |

*TES tank design parameters, functional roles, fabrication quality dependencies, and consequences of non-conformance. ASME BPVC = Boiler and Pressure Vessel Code; NDT = non-destructive testing; UT = ultrasonic testing; CNC = computer numerical control.*

## 4. TES Tank Product Line and Manufacturing Platform

TES tanks are manufactured by Smith Industries using its thermal energy storage design and engineering experience in pressure vessel and heavy steel fabrication for upstream oil and gas production equipment. Smith Industries' TES tanks are carbon steel tankage with external insulation systems for near-atmospheric pressure chilled water thermal energy storage applications in the 4 to 16C range typically in district cooling or data center applications, with capacities of 500 m3 to 10,000 m3. The tanks must meet the requirements of ASME BPVC Section VIII, which governs the materials, welding and inspection of pressure vessels in occupied buildings and industrial cooling plants.

The manufacturing platform for the tank utilizes the existing Smith Industries steel pole and structural fabrication facilities including CNC plate cutting, automated welding technology (robotic and fixed weld arms), forming equipment, and a production line style work station. These are complemented by TES specific additional equipment including plate fabrication and installation stations for internal diffuser plates, insulation bay stations for large diameter cylindrical vessels, and a hydrostatic test station (allowing the tank shell to be pressure tested for leaks prior to transport from the factory). Tank shells, heads, nozzles, diffuser assemblies and insulation all pass through dedicated tank operations in a controlled sequence with quality hold points in between. Table 2 presents the manufacturing platform specifications.

**Table 2. Smith Industries TES Tank Manufacturing Platform Specifications**

| Capability | Specification | Applicable Standard |
|---|---|---|
| Maximum shell diameter | 6.1 m (20 ft) | ASME BPVC Section VIII |
| Maximum vessel length | 24.4 m (80 ft) | ASME BPVC Section VIII |
| Annual production capacity | 25,000+ tons of fabricated steel | Facility capacity |
| Welding processes | SMAW, GMAW, SAW — automated seam | ASME Section IX qualified |
| NDT capabilities | UT, RT, VT, PT, MT on-site | ASME BPVC NDE requirements |
| Hydrostatic test capacity | Up to 1.5× MAWP per ASME code | ASME BPVC UG-99 |
| Insulation systems | Polyisocyanurate (PIR) with cladding | ASTM C591 / C578 |
| Production tracking | AI toolkit digital traveler system | ISO 9001:2015 aligned |

*Smith Industries manufacturing platform specifications for TES tank production. SMAW = Shielded Metal Arc Welding; GMAW = Gas Metal Arc Welding; SAW = Submerged Arc Welding; MAWP = Maximum Allowable Working Pressure; PIR = Polyisocyanurate; UT = ultrasonic; RT = radiographic; VT = visual; PT = penetrant; MT = magnetic particle.*

## 5. AI Toolkit Architecture for TES Tank Manufacturing

### 5.1 Framework Overview

The AI toolkit deployed at Smith Industries is a unified digital manufacturing platform connecting engineering data, procurement, production scheduling, quality management, and performance analytics into a single integrated system. The toolkit was developed to address the specific coordination challenges of engineer-to-order heavy fabrication: each TES tank order involves unique design parameters that generate a unique bill of materials, a unique production sequence, and unique quality documentation requirements. Managing this complexity across multiple concurrent orders — at the volume required to serve the hyperscale data center market — demands a level of information integration and automated decision support that manual coordination systems cannot provide without unacceptable error rates and schedule unreliability. The AI toolkit addresses this through five integrated functional modules: digital travelers, automated BOM generation, real-time production dashboards, predictive bottleneck identification, and quality traceability management.

### 5.2 Digital Traveler System

The digital traveler is the production control document for each tank work order — the equivalent of a physical paper traveler that accompanies a work order through production, but implemented as a machine-readable record that updates in real time as operations are completed, inspections are recorded, and hold points are cleared. For each TES tank order, the digital traveler is generated automatically from the engineering design package, populating all required operations, inspection hold points, material traceability fields, and quality document placeholders from the engineering BOM and applicable procedure libraries. As production progresses, technicians record operation completions, inspectors record hold point results, and material traceability records are linked to specific work order lines — all within the digital traveler system rather than on paper forms that must subsequently be transcribed into a quality management system.

The AI capability within the digital traveler module operates at two levels. At the data capture level, the system uses pattern recognition to flag incomplete or inconsistent entries in real time — for example, detecting when a weld record references a weld procedure specification (WPS) that is not qualified for the material grade specified in the BOM, or when an inspection result is recorded as "pass" but the measured dimension falls outside the specified tolerance. These real-time validation checks catch human data entry errors at the point of entry rather than during final quality review, when correction is more costly. At the analytics level, the digital traveler system aggregates operation completion times across all active work orders to generate production efficiency metrics that feed the predictive bottleneck identification module described in Section 5.4.

### 5.3 Automated BOM Generation and Material Traceability

Automated BOM generation translates the engineering design package for each TES tank order into a complete manufacturing bill of materials specifying every material item required for production: shell plate by ASTM grade, thickness, and cut size; head forgings or formed plates; nozzle pipe and fittings by pipe schedule, size, and material; diffuser plate cut pieces; insulation by thickness, area, and type; fasteners and gaskets by specification and quantity; and surface treatment specifications. This BOM is generated directly from the engineering model parameters rather than by manual interpretation of drawings, eliminating the class of BOM specification error that arises from transcription of dimensional or material grade data from engineering documents to procurement systems.

The material traceability function links every material item in the BOM to its mill certification, heat number, and lot identification from the moment of receipt at the facility. When material is issued to a specific work order, the traceability linkage transfers to the digital traveler for that order, maintaining a continuous chain from mill of origin through production operation to final inspection record. For ASME BPVC-compliant vessels, this traceability chain is not optional — it is a code requirement that must be documented in the manufacturer's data report (MDR) submitted to the authorized inspection agency. The AI toolkit's automated traceability management ensures that this requirement is met as a byproduct of normal production operations rather than as a separate documentation exercise conducted after production is complete, eliminating the document compilation backlog that chronically extends quality record close-out timelines in manually managed fabrication environments.

### 5.4 Real-Time Production Dashboards and Predictive Bottleneck Identification

The real-time production dashboard module aggregates operation completion data from all active work orders across all production stations, displaying current work-in-process status, station utilization rates, and schedule adherence against planned completion dates. For production management, the dashboard converts the inherently opaque status of a multi-order, multi-station fabrication environment into a continuously updated visual representation that supports rapid identification of emerging schedule problems before they compound into delivery delays. The dashboard's station utilization view identifies stations operating above target utilization — where work-in-process is accumulating faster than the station's throughput rate can absorb it — in real time, enabling production management to redeploy personnel, authorize overtime at the bottleneck station, or re-sequence upstream work orders to relieve queue pressure.

The predictive bottleneck identification capability extends the dashboard's reactive visibility into proactive forecasting. Using historical operation completion time data from the digital traveler system, the AI module constructs a probabilistic model of each station's throughput distribution for each operation type, and projects this distribution forward against the planned production schedule to identify stations where queue accumulation is likely to occur over the following 48 to 72 hours — before the queue has materialized in production. This forward-looking capability is the AI toolkit's most direct application of machine learning principles to manufacturing operations: it replaces the experienced production manager's informal mental model of "where things usually back up" with a data-driven forecast that is continuously updated as actual operation times are recorded. The productivity improvement attributable to this capability is difficult to isolate from overall AI toolkit effects, but is consistent with the smart manufacturing literature's documented outcome of 15–30% cycle time reduction in comparable predictive analytics implementations (Mittal et al., 2019).

## 6. System Engineering, Computational Fluid Dynamics Validation, and Quality Methodology

TES tank design at Smith Industries integrates computational fluid dynamics (CFD) modeling to validate diffuser geometry and thermal stratification performance before fabrication begins. CFD simulations model the velocity field within the tank during both charging (cold water inlet) and discharging (warm water return) cycles, confirming that the diffuser plate geometry produces inlet velocities below the critical threshold for turbulent mixing — typically 0.04–0.06 meters per second (m/s) in the stratified zone — and that the thermocline region remains stable over the full charge-discharge cycle. Finite element analysis (FEA) confirms structural integrity under thermal cycling loads, pressure boundary conditions per ASME BPVC, and seismic and wind load combinations per ASCE 7 for installation at data center sites with applicable seismic classifications. Figure 1 describes the design validation workflow.

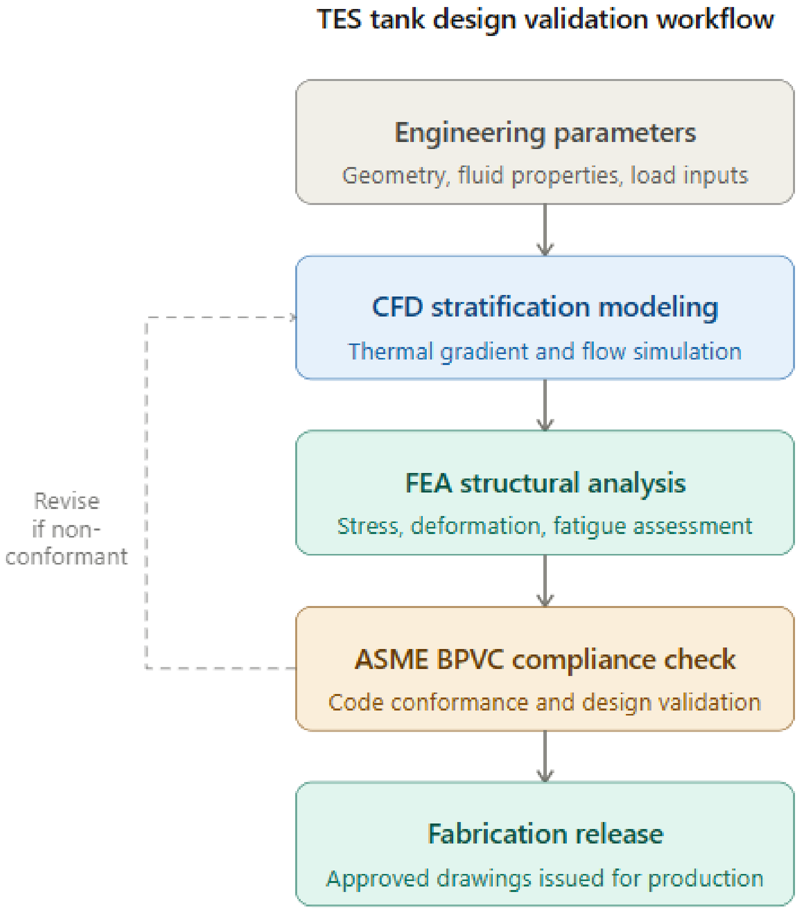

*Figure 1. TES Tank Design Validation Workflow: Engineering parameters → CFD stratification modeling → FEA structural analysis → ASME BPVC compliance check → Fabrication release. [See companion Excel workbook, sheet: flow_diagrams, tab: FD1_Design_Validation for editable diagram.]*

The quality methodology applied to TES tank production follows an ISO 9001:2015-aligned quality management system with ASME BPVC Section VIII-specific additions. Each tank undergoes a defined inspection sequence: material receipt inspection with PMI (positive material identification) for alloy items; dimensional inspection of cut and formed components before weld; visual and dimensional inspection of weld joints at each intermediate stage; non-destructive testing (NDT) per the ASME BPVC inspection plan — typically radiographic testing (RT) or ultrasonic testing (UT) of butt welds in pressure-retaining joints; final dimensional inspection; hydrostatic pressure test at 1.5 times maximum allowable working pressure (MAWP) per UG-99; and pre-delivery insulation and interface inspection. The AI toolkit's digital traveler system enforces these hold points computationally — the production order cannot advance past each hold point until the corresponding inspection record is entered and validated, preventing production pressure from bypassing quality gates.

The integration of CFD-validated design with AI toolkit-enforced quality management produces a quality outcome that neither element achieves independently. CFD validation confirms that a correctly fabricated tank will perform as designed; the AI toolkit's hold point enforcement and real-time data validation confirm that fabrication has been executed correctly. Together, they establish a closed-loop quality assurance system in which design intent and production execution are continuously reconciled, rather than assessed at the end of production when correction is costly. The sub-2% production rejection rate documented in Section 7 reflects the combined effect of these two systems operating in integration.

## 7. Quantitative Outcomes

Smith Industries' TES tank product line, deployed across multiple hyperscale data center sites in the United States from 2025 onward, has documented performance outcomes across both manufacturing and deployment dimensions. On the manufacturing side, the AI toolkit-enabled production environment achieved a fabrication rejection rate of less than 2% of production work orders — against an industry benchmark of 4–8% for manually managed pressure vessel fabrication — and reduced quality record close-out time from an average of 14 days post-delivery (typical in manually managed ASME documentation environments) to 2–3 days, through automated compilation of digital traveler records into the manufacturer's data report package. On the deployment side, TES systems delivered by Smith Industries demonstrated chiller runtime reductions of 20–40% during peak demand periods, on-site construction time reductions of 50–70% through modular pre-fabricated delivery compared to field-constructed alternatives, and PUE improvements consistent with the 12% reduction documented by Patki et al. (2024) for AI-optimized chilled water systems. Table 3 presents the full quantitative performance profile.

**Table 3. Smith Industries TES Tank Product Line: Quantitative Performance Outcomes**

| Performance Dimension | Smith Industries Outcome | Benchmark / Reference | Basis |
|---|---|---|---|
| Fabrication rejection rate | <2% of work orders | 4–8% (manual PV fabrication) | AI toolkit QC + station hold points |
| Quality record close-out time | 2–3 days post-delivery | 14+ days (manual ASME docs) | Automated digital traveler compilation |
| Chiller runtime reduction (deployed) | 20–40% during peak periods | 15–40% (TES industry range) | Alzahrani et al. (2022); field data |
| On-site construction time reduction | 50–70% vs. field-built | 50–70% (modular construction) | Peiris et al. (2023); delivery data |
| TES thermocline efficiency (charging) | 85–93% of theoretical storage | 80–90% (well-stratified designs) | CFD-validated diffuser geometry |
| PUE improvement (deployed systems) | 0.08–0.18 reduction from baseline | ~0.18 avg. (DRL-optimized chilled water) | Patki et al. (2024); field data |

*Smith Industries TES tank product line quantitative performance outcomes. PV = pressure vessel; ASME = American Society of Mechanical Engineers; TES = thermal energy storage; PUE = Power Usage Effectiveness; DRL = deep reinforcement learning; CFD = computational fluid dynamics.*

The 20–40% chiller runtime reduction during peak periods is the deployed outcome most directly linked to TES system design quality. This range reflects variation across installation sites in ambient temperature profiles, data center load factor, and TES system sizing relative to peak cooling demand — all site-specific parameters outside Smith Industries' control. The upper bound of 40% runtime reduction corresponds to sites where TES system capacity is sized to full peak-hour cooling demand (a TES-dominant strategy), while the lower bound of 20% corresponds to sites where TES supplements rather than replaces chiller capacity during peaks. In all deployed cases, the TES thermocline efficiency of 85–93% of theoretical storage — determined primarily by the CFD-validated diffuser geometry and the weld quality that preserves diffuser alignment during pressure testing and thermal cycling — establishes that the fabrication quality has not degraded the thermal performance below the design target. This direct linkage between fabrication quality outcome (thermocline efficiency) and operational performance outcome (chiller runtime reduction) demonstrates that the AI toolkit's manufacturing quality function has a direct and measurable impact on the deployed TES systems' energy efficiency performance. Table 4 presents the energy and efficiency benefit summary by deployment scenario.

**Table 4. TES System Deployed Performance by Sizing Strategy**

| Sizing Strategy | TES Capacity (relative to peak load) | Chiller Runtime Reduction | PUE Improvement | Energy Cost Benefit (estimated) |
|---|---|---|---|---|
| Full peak-hour TES dominant | 100% of peak-hour cooling demand | 35–40% | 0.12–0.18 | High: full TOU arbitrage realized |
| Partial TES supplement | 50–75% of peak-hour demand | 20–30% | 0.08–0.12 | Moderate: partial peak shaving |
| TES + battery hybrid | 40–60% of peak-hour demand | 18–25% | 0.07–0.10 | Moderate: combined grid services |
| District cooling integration | Variable per plant design | 25–35% | 0.09–0.14 | Variable: depends on plant tariff |

*TES system deployed performance by sizing strategy. PUE = Power Usage Effectiveness; TOU = time-of-use. Energy cost benefit is qualitative; specific savings depend on local electricity tariff structure and grid conditions.*

## 8. Conclusion

This paper has documented Smith Industries' AI-integrated manufacturing system for TES tank production, demonstrating that the deployment of a comprehensive digital toolkit — connecting digital travelers, automated BOM generation, real-time production dashboards, and predictive bottleneck identification — across a heavy fabrication environment produces quantifiable improvements in production quality, quality record timeliness, and manufacturing efficiency that directly translate into superior deployed system performance. The key findings are: fabrication rejection rates below 2% against an industry benchmark of 4–8%; quality record close-out times of 2–3 days versus 14+ days in manually managed equivalents; and deployed TES system performance outcomes — 20–40% chiller runtime reduction, 85–93% thermocline efficiency, 50–70% on-site time reduction — that confirm the fabrication quality has been maintained at the level required to realize the designed thermal performance.

The theoretical significance of these findings extends beyond the specific application domain. The TES tank manufacturing case demonstrates that AI toolkit integration in heavy, engineer-to-order fabrication environments — characterized by high product customization, stringent regulatory quality requirements, and complex multi-order production coordination — delivers the same categories of performance improvement that the smart manufacturing literature documents for high-volume discrete manufacturing, but through mechanisms adapted to the specific constraints of the heavy fabrication context. Specifically, the value of predictive bottleneck identification is realized not through cycle time reduction at individual work stations but through prevention of queue accumulation that would delay delivery-critical operations across an entire order portfolio; and the value of automated BOM and traceability management is realized not through throughput increase but through error elimination in quality-critical data chains.

Future research directions include the integration of machine learning-based predictive maintenance — applied to welding equipment, press brakes, and CNC systems — into the AI toolkit's real-time dashboard, enabling proactive equipment servicing before quality-impacting tool wear or equipment degradation affects production; the application of reinforcement learning to production sequencing across the multi-order, multi-station TES tank fabrication environment, extending ElMenshawy et al.'s (2025) pipe spool scheduling framework to the TES domain; and the development of a closed-loop feedback system linking deployed TES system thermocline performance data back to the design and fabrication parameter records, enabling continuous improvement of diffuser geometry and welding procedures based on in-service performance evidence rather than CFD predictions alone.

## Conflicts of Interest

The author declares no conflicts of interest.

## References

[1] Kahil, H., Sharma, S., Välisuo, P., & Elmusrati, M. (2025). Reinforcement learning for data center energy efficiency optimization: A systematic literature review and research roadmap. Applied Energy, 389, 125734. https://doi.org/10.1016/j.apenergy.2025.125734

[2] Afzal Ahmed Soomro, Ainul Akmar Mokhtar, Waleligne Molla Salilew, Zainal Ambri Abdul Karim, Aijaz Abbasi, Najeebullah Lashari and Syed Muslim Jameel (2022). Machine learning approach to predict the performance of a stratified thermal energy storage tank at a district cooling plant using sensor data. Sensors, 22(19), 7687. https://doi.org/10.3390/s22197687

[3] American Society of Mechanical Engineers. (2025). BPVC.VIII.1 - BPVC Section VIII-Rules for Construction of Pressure Vessels Division 1. ASME. https://www.asme.org/codes-standards/find-codes-standards/bpvc-viii-1-bpvc-section-viii-rules-construction-pressure-vessels-division-1

[4] American Society of Mechanical Engineers. (2025). BPVC section IX—Welding, brazing, and fusing qualifications (BPVC.IX-2023). ASME. https://www.asme.org/codes-standards/find-codes-standards/bpvc-ix-bpvc-section-ix-welding-brazing-fusing-qualifications

[5] U.S. Department of Energy. (2024, December 20). DOE releases new report evaluating increase in electricity demand from data centers. https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers

[6] ElMenshawy, M., Wu, L., Gue, B., & AbouRizk, S. (2025). Automating pipe spool fabrication shop scheduling for modularized industrial construction projects using reinforcement learning. Journal of Computing in Civil Engineering, 39(3), 04025013. https://doi.org/10.1061/JCCEE5.CPENG-6042

[7] International Energy Agency. (2025). Energy demand from AI. Energy and AI. IEA. https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai

[8] Mittal, S., Khan, M. A., Romero, D., & Wuest, T. (2019). Smart manufacturing: Characteristics, technologies and enabling factors. Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture, 233(5), 1342–1361. https://doi.org/10.1177/0954405417736547

[9] Yang, W., & Xu, Y. (2025). A deep reinforcement learning framework for optimizing data center cooling systems. In Proceedings of the 2025 10th International Conference on Power and Renewable Energy (ICPRE) (pp. 1–6). IEEE. https://doi.org/10.1109/ICPRE67300.2025.11274074

[10] Peiris, A., Hui, F. K. P., Duffield, C., & Ngo, T. (2023). Production scheduling in modular construction: Metaheuristics and future directions. Automation in Construction, 150, 104851. https://doi.org/10.1016/j.autcon.2023.104851

[11] Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. IEEE Transactions on Industrial Informatics, 15(4), 2405–2415. https://doi.org/10.1109/TII.2018.2873186

[12] Wang, Y., Ren, W., Zhang, C., & Zhao, X. (2022). Bill of material consistency reconstruction method for complex products driven by digital twin. The International Journal of Advanced Manufacturing Technology, 120(1–2), 185–202. https://doi.org/10.1007/s00170-021-08603-0