International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

Data-Driven Facility Management for Sustainability and Resource Optimisation: AI-Assisted Automation and Operational Decisions for an Enhanced Visitor Experience

Dr. P. Vijayakumar[1], Mr. M. Sivaramakrishna[2], and Mrs. Bandi Aruna[3]

[1]Corporate Chief Engineer, Ramoji Film City, Hyderabad, India

[2]Director, Ramoji Film City, Hyderabad, India

[3]Assistant Professor, St. Pious X Degree and PG College for Women, Hyderabad, India

Abstract - Mass-visitor venue entrance zones lose service quality and energy during sudden arrival surges, as isolated resource evaluations obscure root causes of waste and delays. This paper proposes an integrated, data-driven framework governed by the Measure–Analyse–Decide–Optimise–Verify–Sustain (MADOVS) cycle for two linked problems: dynamic queue staffing and demand-controlled HVAC in waiting halls. Applied to a theme park entrance, a gradient-boosted regression model predicts 15-minute visitor arrivals using recent counts, weather, calendar data, and group bookings. Algorithmic outputs dynamically set three-stage queue lane counts via an $M/G/c$-type utilization rule and modulate cooling set-points and fresh-air rates through a lumped thermal model. In a 30-day synthetic benchmark, gradient boosting reduced forecast MAE from 23.2 to 10.7 visitors (MAPE 14.4%). Forecast-driven staffing lowered 90th-percentile waiting times by 37%, while demand-controlled ventilation reduced daily HVAC energy by 22–48% while maintaining thermal comfort (22–26 °C). A field protocol is specified to validate these simulated tentative results at Ramoji Film City.

Keywords: Facility management, artificial intelligence, operational data, resource efficiency, demand forecasting, sustainability, visitor services, MADOVS cycle, closed-loop control.

1. INTRODUCTION

1.1 Purpose of the Study

The central research question driving this study is simple: how can a complex facility utilize its operational resources more effectively while maintaining dependable, high-quality service? The proposed answer combines engineering domain knowledge, rigorous data analysis, and supervised automated decisions. This framework positions facility management as a proactive contributor to organizational performance and visitor satisfaction, extending far beyond the traditional view of facility management as a reactive cost center for routine maintenance.

1.2 The Entrance Zone as a Connected System

An entrance zone brings vehicles, people and engineering services together. A bus can release a large group of visitors within minutes. Ticket counters, baggage scanners and turnstiles must respond to this changing demand, while lighting, cooling, water supply and public announcements remain available.

A delay at one stage affects the next. For example, slow baggage screening can increase waiting-area occupancy and the demand for cooling, cleaning and visitor assistance. Opening another ticket counter may simply move the queue to security. The complete visitor journey must therefore be examined.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

1.3 Study Boundary

The proposed boundary extends from vehicle arrival and parking to ticketing, screening, resting areas and final admission. It also includes the interface with vehicle collection and exit. Supporting services include air conditioning, lighting, audio and visual systems, water supply, greenery and the allocated share of wastewater treatment. Rides and the remaining park are outside this boundary.

2. Literature Review

2.1 Relationship to Published Work

Integrating proactive maintenance, crowd dynamics, and data-driven operational decisions relies on established domain literature. Guidelines from the U.S. Department of Energy (2010) demonstrate that aligning daily operations with real-time data drives energy, water, and cost efficiency, while Health and Safety Executive frameworks connect arrival, circulation, and departure management to crowd safety. These foundational sources advocate coordinated, system-wide planning over isolated departmental decisions. Building on these principles, Krishnakumari et al. (2023) demonstrated real-time crowd management using data-driven decision support and digital twins. In contrast to purely theoretical or algorithmic studies, this paper proposes the Measure–Analyse–Decide–Optimise–Verify–Sustain (MADOVS) operational cycle—a practical framework designed to function with spreadsheets and manual logs before scaling to complex automated infrastructure

3. METHODOLOGY AND FRAMEWORK

3.1 The MADOVS Improvement Cycle

The MADOVS framework structures continuous operational improvement into six structured, sequential phases designed to link measurement directly with verified action.

3.1.1 Measure and Analyse

Measure establishes the present operating condition, called the baseline. Record visitor demand, resource consumption, equipment hours, faults and service quality. Analyse then compares these records to identify causes of waste or delay. Attendance, weather and operating hours must be considered before drawing conclusions.

Phase Formulations:

Let $X_t \in R^d$ represent the multidimensional state vector of the entrance zone captured at time step t ($\Delta t = 15$ min):

$$X_t = [\, N_t^{arr}, N_t^{adm}, Q_t^{tick}, Q_t^{sec}, P_t^{elec}, V_t^{water}, T_t^{in}, T_t^{out}, RH_t \,]^T$$

where N_t^{arr} is vehicle arrivals, N_t^{adm} is turnstile admissions, Q_t^{tick} and Q_t^{sec} are queue lengths, P_t^{elec} is real-time power demand (kW), V_t^{water} is freshwater flow rate (L/min), T_t^{in} and T_t^{out} are indoor/outdoor dry-bulb temperatures (°C), and RH_t is relative humidity (%). The Analysis function $A(X_t)$ computes operational intensity indices and evaluates system deviations against benchmark performance matrices.

3.1.2 Decide and Optimise

Decide selects a practical action, a responsible person and a review date. Each action needs a resource target and a service condition. Optimise means testing the approved change, such as adjusting a cooling schedule or assigning trained staff before a peak arrival period. It does not mean reducing every resource to its minimum.

The Decision phase evaluates objective functions balancing energy efficiency against visitor wait times. The optimization formulation is expressed as:

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

$$\min_{u_t \in U} J = \Sigma_{k=0}^{H} (\alpha \cdot P_{t+k}^{elec}(u_{t+k}) + \beta \cdot W_q(u_{t+k}) + \gamma \cdot | T_{t+k}^{in}(u_{t+k}) - T_{set} |)$$

$$\text{subject to: } W_q \leq W_{max}, \quad T_{min} \leq T_{in} \leq T_{max}, \quad u_t \in U_{safe}$$

where u_t is the control vector (HVAC fan speed setpoints, open screening lanes, guidance display messages), H is the prediction horizon, W_q is average queue waiting time, T_set is target indoor temperature (24°C), and α, β, γ are operational weighting factors.

3.1.3 Verify and Sustain

Verify compares performance before and after the change under comparable conditions. Check consumption, cost, queues, comfort and equipment reliability together. Sustain incorporates successful changes into operating procedures, staff training and periodic reviews. The cycle repeats when demand or equipment changes.

Verification evaluates performance under equivalent exogenous conditions (weather and attendance) using normalized baseline subtraction. Successful operational parameters u_t* that achieve J* J_baseline are embedded into automated programmable logic controller (PLC) routines and standard operating procedures (SOPs) during the Sustain phase.

3.2 Data Infrastructure and Information Architecture

3.2.1 Data Collection Standards

To establish operational utility, every facility record must capture standardized metadata fields, including date, time stamp, location, asset or service identifier, quantitative value, measurement unit, and data source. Essential input streams encompass vehicle arrival counts, admission tallies, queue lengths, sampled waiting times, electricity and water meter readings, maintenance fault logs, and staff-hours categorized by specific task. Adopting uniform sampling intervals across disparate data streams simplifies comparative analysis.

For pilot deployments, visitor arrivals and queue status should ideally be recorded at 15-minute intervals where practical. Electrical and water consumption data should be recorded on matching 15-minute intervals or daily basis, depending on meter hardware capability. Because an aggregated daily utility reading cannot explain localized 15-minute queue spikes, analytical conclusions must be aligned with the temporal resolution of the underlying data.

The underlying IoT telemetry infrastructure is organized into a four-tier architecture, as summarized in Table 1.

Layer	Hardware & Protocols	Data Payload & Latency	Primary Function
1. Sensing & Actuation	Modbus-RTU Class 0.5S energy meters, optical counters, BACnet VAV actuators	Voltage, current, P(kW), flux counts, valve positions (1–5 sec)	Physical telemetry capture and hardware actuation.
2. Edge/Network Gateway	Industrial IoT Gateways, MQTT brokers over TLS 1.3, RS-485 to Ethernet	JSON payloads, timestamped telemetry streams (15 sec)	Protocol translation, edge filtering, secure transport.
3. Analytics & Repository	Time-series database (InfluxDB), PostgreSQL, Python ML engines	Cleaned time-series records, feature matrices (15 min aggregate)	Model inferencing, anomaly scoring, predictive optimization.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

4. Supervisory UI	Web SCADA dashboards, automated SMS/email alerts, API endpoints	KPI widgets, queue alerts, control override interface (< 1 sec UI refresh)	Real-time visualization and supervisory human oversight.

Table 1. Four-Tier Information and IoT Telemetry Architecture.

3.2.2 Data Quality Assurance Before AI Implementation

Data cleaning represents a prerequisite for any advanced analytics or machine learning implementation. This process requires detecting and resolving duplicate entries, inconsistent measurement units, out-of-range sensor readings, and time-stamp mismatches. Missing data must be explicitly flagged rather than filled with default zeros; a failed flow sensor, for instance, must not be misinterpreted as zero utility consumption. Furthermore, automated optical or electronic counts must undergo periodic validation against manual sample observations.

Operational teams must distinguish between ticket sales, total admissions, and real-time facility occupancy. Real-time occupancy requires calculating an initial baseline count, adding cumulative entries, and subtracting cumulative exits, with adjustments applied for staff movements and visitor re-entry. Where utility infrastructure is shared across functional zones, separate sub-metering or mathematically documented allocation methods must be established prior to analysis.

3.2.3 Information Architecture and IoT Integration

The proposed information architecture consists of four functional layers: data collection, a unified data repository, analysis engines, and an operational management dashboard. Meters, entry counters, and digital operating logs serve as the primary data collection layer. A relational database or centralized spreadsheet aggregates these inputs by time stamp and location. Analytical models generate operational indicators and automated alerts, which are visualized on a supervisory dashboard for real-time facility decision-making (Table 2).

Architecture Layer	Primary Components	Key Functions
Data Collection	IoT energy meters, optical counters, ticketing logs, manual logs	Capture real-time physical metrics and operational events.
Data Repository	Relational database, cloud data store, validated spreadsheets	Standardize, clean, and store operational data across time stamps.
Analysis Engine	Forecasting models, anomaly detection, rule engines	Process data streams to yield performance indicators and alerts.
Supervisory Dashboard	GUI displays, mobile alerts, supervisory interfaces	Present actionable operational insights to facility managers.

Internet of Things (IoT) devices, such as networked sensors and smart utility meters, enable continuous

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

electronic data transmission, reducing manual data entry burdens. However, connectivity alone does not guarantee operational improvement. Sensor calibration schedules, communication health checks, and clearly assigned staff responsibilities for responding to system alerts are equally critical for system success.

3.3 Key Performance Indicators

To accurately evaluate facility performance, key performance indicators (KPIs) must combine utility efficiency metrics with visitor experience standards:

- **Electricity Intensity per Visitor:** Entrance-zone electrical consumption in kilowatt-hours (kWh) divided by total visitor admissions over the identical operating period.
- **Freshwater Intensity per Visitor:** Total freshwater consumption in liters (L) divided by total visitor admissions over the matching timeframe. Total consumption must always be reported alongside intensity ratios to prevent misleading interpretations during high-attendance periods.
- **Queue Wait Metrics:** Median wait time and 90th-percentile wait time, specifying the service checkpoint, time interval, and sample size. The 90th-percentile wait metric identifies tail-end service delays that disproportionately impact visitor satisfaction.
- **Equipment Availability:** The percentage of scheduled operational hours during which critical assets (e.g., scanners, turnstiles) remain fully functional.
- **Operating Cost per Visitor:** Total entrance-zone operational expenditures divided by total visitor admissions during a defined reporting period.

3.4 Responsible Use of Data and Governance

Data privacy and governance protocols must govern all operational collection activities. Where possible, systems should rely on aggregated, non-identifiable visitor counts. Access to high-resolution camera feeds or vehicle license plate records must be strictly restricted, with explicit retention schedules established. Optical sensors deployed for crowd counting do not require facial recognition capabilities.

Clear governance mandates assigning accountability for data integrity, dashboard review, and the execution of operational recommendations. Comprehensive logs of all manual overrides and automated control actions must be maintained. If the digital analysis platform experiences an outage, established manual operating procedures must remain readily executable to ensure operational continuity.

4. ARTIFICIAL INTELLIGENCE AND AUTOMATION FRAMEWORK

4.1 AI Applications in Facility Operations

4.1.1 Demand Forecasting Models

Predictive machine learning models can estimate upcoming visitor and vehicle arrivals by analyzing historical arrival patterns, time-of-day dynamics, day-of-week trends, public holiday calendars, weather forecasts, and advance group booking data. These input features allow algorithms to output predicted arrival volumes for subsequent 15- or 30-minute operational windows.[1]

Short-term visitor arrival forecasting is executed using a hybrid Extreme Gradient Boosting (XGBoost) model combined with an LSTM network. The predictive model estimates upcoming arrival volume $\hat{Y}_{t+\Delta t}$ over a 15-minute horizon:

$$\hat{Y}_{t+\Delta t} = f(H_t, W_t, C_t, B_t)$$

where $H_t = [N_t^{arr}, N_{t-1}^{arr}, ..., N_{t-4}^{arr}]$ represents historical arrival lags, W_t is the real-time weather vector, C_t is categorical calendar encoding (day of week, holiday flag, peak season index), and B_t is advance group booking reservations.

Model performance is rigorously evaluated using Mean Absolute Percentage Error (MAPE) and Root Mean Square Error (RMSE):

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

$$MAPE = (100\% / n) \cdot \Sigma_{i=1}^{n} | (Y_i - Y_hat_i) / Y_i |$$

$$RMSE = sqrt((1 / n) \cdot \Sigma_{i=1}^{n} (Y_i - Y_hat_i)^2)$$

Accurate demand forecasts enable supervisors to proactively configure ticket counters, deploy parking personnel, and adjust resting area climate controls prior to physical crowd arrival. Operational teams should establish baseline forecasts using simple historical moving averages, adopting complex machine learning models only when they demonstrate statistically significant improvements in predictive accuracy that lead to better operational decisions.

4.1.2 Anomaly Detection Algorithms

Anomaly detection algorithms scan continuous operational data streams to identify readings that deviate from expected baselines. Common operational anomalies include off-hours water flow indicating leaks, unexplained electrical spikes during low-occupancy periods, and recurring equipment trips. An automated anomaly flag serves as a trigger for physical inspection rather than a definitive diagnosis.

While fixed threshold rules provide simple, effective alerting, machine learning models that learn dynamic baseline behavior can adjust thresholds for varying attendance levels and ambient weather conditions. Regardless of the underlying algorithm, maintenance teams must verify sensor integrity, check for special events, and confirm genuine physical faults before initiating major interventions.Unsupervised anomaly detection isolates operational faults in electricity and water consumption streams. The Isolation Forest algorithm constructs an ensemble of isolation trees partitioning data points $x \in R^d$. The anomaly score $s(x, n)$ for a sample x is defined as:

$$s(x, n) = 2^(- E(h(x)) / c(n))$$

where $h(x)$ is the path length of sample x across trees, $E(h(x))$ is the mean path length, and $c(n)$ is the average path length of unsuccessful searches in a Binary Search Tree created with n nodes. Instances yielding $s > 0.65$ trigger automated inspection alerts for off-hours pipe leaks or anomalous electrical baseline consumption.

4.1.3 Maintenance Decision Support Systems

Tracking equipment runtime hours, historical fault logs, and physical condition parameters (such as motor operating current or vibration profiles) aids in prioritizing maintenance inspections. Predictive maintenance techniques aim to forecast remaining useful asset life or imminent failure risks. However, reliable predictive maintenance requires extensive validated failure data; a digital fault logging system alone does not constitute a predictive maintenance capability.

4.2 Human Supervision, Explainability, and Model Validation

4.2.1 Human Oversight and Explainability

The proposed framework couples automated control capability with mandatory human supervision. Routine operational adjustments within established safety thresholds can execute automatically, whereas non-routine events or system exceptions require supervisor confirmation. Algorithmic explainability requires recording the explicit rationale behind every automated control prompt, such as projected occupancy surges or scheduled group arrivals. High-uncertainty forecasts must never trigger unvalidated or high-risk control modifications.

4.2.2 Model Training, Validation, and Drift

To ensure robust performance, machine learning models must be trained on historical dataset splits and evaluated on unseen future time intervals, preventing temporal data leakage. Forecasting performance should be

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

evaluated against actual observations using Mean Absolute Error (MAE), with peak traffic windows analyzed independently. For anomaly detection, evaluation must monitor false alarm rates and missed events. Furthermore, models must be continuously monitored for performance degradation over time (model drift) as seasonal factors and visitor behavior patterns evolve.

4.3 Closed-Loop Control and Automatic Decision Execution

4.3.1 Queue Guidance and Parking Management

When localized queue congestion is detected at specific turnstiles, automated digital displays can instantly redirect incoming flows toward open, staffed lanes with available capacity. To prevent bottleneck migration, the control system must evaluate capacity at downstream security screening points before adjusting upstream ticketing guidance.Automated parking management displays update real-time stall availability and guide approaching vehicles toward designated zones, maintaining distinct guidance logic for tour buses, passenger cars, and accessible vehicles. If occupancy sensors or communication links fail, displays must visually indicate data unavailability and revert to manual staff directing procedures.

Security screening and ticketing bottlenecks are mathematically modeled using multi-server M/M/c queuing theory.[2] Assuming Poisson visitor arrivals at rate λ and exponentially distributed processing times at rate μ across c open parallel screening lanes, the traffic intensity ρ is given by:

$$\rho = \lambda / (c \cdot \mu)$$

For a stable system (ρ < 1), the probability of zero visitors in the queue (P_0) is:

$$P_0 = [\Sigma_{k=0}^{c-1} ((c\rho)^k / k!) + ((c\rho)^c / (c! (1 - \rho)))]^{-1}$$

The average wait time in the queue (W_q) is derived as:

$$W_q = L_q / \lambda = [(P_0 \cdot (c\rho)^c \cdot \rho) / (c! (1 - \rho)^2)] \cdot (1 / \lambda)$$

When algorithmic forecasting predicts $\lambda_hat_{t+\Delta t}$ such that W_q > W_threshold (e.g., W_q > 10 minutes), the control system automatically triggers dynamic display redirection and dispatches notifications to deploy additional screening personnel.

4.3.2 Demand-Responsive HVAC and Lighting Controls

Integrating real-time space occupancy and temperature telemetry allows HVAC control systems to adjust cooling stages within pre-approved comfort bands. Predictive arrival signals enable pre-cooling of assembly halls prior to group arrivals. Conversely, when occupancy drops, controllers can scale back cooling output while maintaining mandatory fresh air ventilation standards.

Similarly, daylight-harvesting sensors and occupancy detectors can dynamically dim or switch non-essential lighting circuits. Emergency lighting and minimum illumination standards remain hardwired and protected. Anti-cycling timers, deadband temperature ranges, and equipment interlocks must be built into control logic to prevent rapid equipment cycling and conflicting commands.

To prevent energy waste while maintaining indoor air quality and thermal comfort, HVAC central AHU fan speeds (S_fan) and chilled water valve openings (V_chw) are dynamically adjusted based on predicted occupancy $N_hat_t^{occ}$:

$$S_fan(t) = S_min + (S_max - S_min) \cdot min(1, N_hat_t^{occ} / N_capacity)$$

Built-in hardware deadbands (±0.8°C) and minimum runtime anti-cycling timers (10 minutes) prevent actuator

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

hunting and mechanical wear.

4.3.3 Automated Visitor Information and Service Dispatch

Dynamic displays can automatically update projected wait times, directional maps, and facility status announcements. When queue thresholds are breached, public address systems can select and broadcast pre-approved multilingual audio messages. For routine inquiries, conversational information assistants can provide validated guidance, transferring complex or urgent issues to human staff.

When an asset experiences repeated faults, the system can automatically generate a work order within the computerized maintenance management system (CMMS), specifying the fault code, precise physical location, and priority level, subsequently tracking response and resolution times.

4.3.4 Rule-Based Logic versus AI and Operating Limits

Not all automated actions require complex artificial intelligence. Standard rule-based conditional logic is fully sufficient for straightforward actions, such as updating a parking sign when a zone reaches maximum capacity, whereas predictive AI models excel at anticipating future demand patterns. Both approaches must operate within defined physical and safety boundaries. Security screening protocols, emergency life-safety systems, and hardware safety interlocks must remain strictly governed by standard safety procedures.

5. OPERATIONAL DOMAIN APPLICATIONS

5.1 Vehicle Arrival and Ticketing Services

Facility management teams must record vehicle arrivals categorized by mode (buses, personal vehicles, taxis), parking space utilization rates, entry gate processing delays, and drop-off durations. Tour buses and passenger cars impose fundamentally different spatial and pedestrian flow demands. Jointly evaluating vehicle arrival telemetry and visitor entry counts enables managers to dynamically staff admission gates and maintain safe pedestrian-vehicle segregation.

The generator produces 113 days of 15-minute arrivals between 09:00 and 18:00 (36 intervals): walk-in visitors follow a non-homogeneous Poisson process with a morning peak and a smaller afternoon peak, scaled by day type (weekday 1.0, weekend 1.8, holiday 2.3), rain (−25%) and temperature; coaches of 40 prepaid visitors arrive in bursts, of which most are booked in advance and a further share are unbooked. The first 83 days (after a 7-day warm-up for lags) are used for training and the final 30 days, never seen in training, for testing. The test days average 2244 visitors.

Queue policies are compared with common random numbers: for each test day and each of three replications, both policies face identical arrival times and identical service-time draws, so the paired difference reflects staffing alone. Confidence intervals are 95% percentile bootstrap intervals over the 30 test days.

5.2 Forecast accuracy (RQ1)

Model	MAE	RMSE	MAPE (%)	WAPE (%)	Peak MAE
Seasonal naive (t−7d)	27.1	41.0	37.1	43.4	60.4
Historical mean (daytype × interval)	23.2	31.1	31.0	37.2	42.1
Random Forest	12.1	20.3	16.4	19.4	26.1
Gradient boosting	10.7	18.7	14.4	17.2	24.5

Table 3. One-step-ahead forecast error on 30 unseen synthetic test days (persons per 15-min interval). MAPE uses intervals with ≥ 50 arrivals.

Gradient boosting cut MAE by 54% relative to the historical mean and by 60% relative to seasonal naive; the random

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

forest was slightly behind. Feature importance (gradient boosting) was highest for the interval-of-day index (0.389), booked group visitors (0.194), the previous-interval count (0.177) and day type (0.121). Because the generator was built with these dependencies, the ranking confirms the pipeline works; it does not prove the same ranking holds at a real venue.

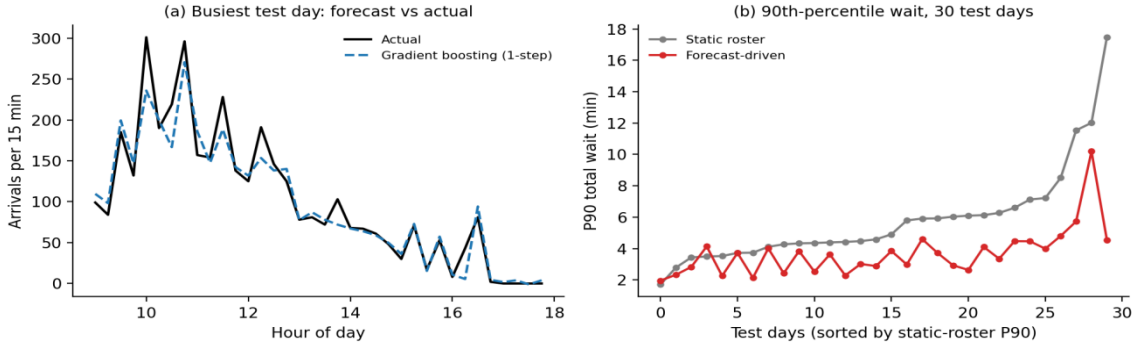


Figure 1. (a) One-step forecast against actual arrivals on the busiest test day. (b) 90th-percentile total waiting time per test day under the static roster and the forecast-driven roster.

Queue performance at equal staff-hours (RQ2)

Metric (mean over 30 test days)	Static roster	Forecast-driven	Change
Mean total wait (s)	108.6	73.5	-35.0 s [-47.2, -24.4]
Median total wait (s)	27.5	24.3	-3.2 s
90th-percentile wait (s)	349.8	220.8	-129.1 s [-186.2, -85.4]
95th-percentile wait (s)	484.9	310.1	-174.8 s
Visitors waiting > 10 min (%)	2.9	0.8	-2.1 pp
Lane-hours per day (all stages)	47.5	46.3	-1.1 h [-2.1, -0.2]

Table 4. Paired queue results; brackets give 95% bootstrap intervals for the mean paired difference.

With about 2% fewer lane-hours, forecast-driven staffing lowered mean total wait by 32% and the 90th-percentile wait by 37%. The median moved little (27.5 s to 24.3 s): the gain sits in the tail, on days when the static roster, built from average profiles, is caught short by holiday peaks. Static roster remained better than forecast-driven on a few low-demand days (Figure 1b, left side), which is expected when both rosters are already adequate.

5.2 Resting Areas and Visitor Communication

Evaluating resting zone performance requires monitoring seating availability, space occupancy, shade coverage, physical accessibility, and visitor assistance requests. Turnstile counts alone cannot capture visitor comfort; seating may be quantitatively available but poorly distributed or exposed to extreme solar radiation. Observational audits and visitor feedback loops provide necessary context to numerical occupancy data.[3]

5.3 Energy Efficiency and Indoor Environmental Quality

5.3.1 Distinguishing Electrical Power from Energy

Accurate energy management requires distinguishing between electrical power, measured in kilowatts (kW) as the instantaneous rate of energy consumption, and electrical energy, measured in kilowatt-hours (kWh) as consumption

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

integrated over time. Furthermore, HVAC cooling capacity expressed in refrigeration tons represents a thermal rating and must not be confused with electrical power input. All energy metrics must be derived from calibrated power meters or verified operational data.[5]

HVAC energy and comfort (RQ3)

Baseline fresh-air rate (% of design)	Baseline kWh/day	Demand-controlled kWh/day	Modelled reduction
50%	397	309	22.3%
75%	490	309	37.2%
100%	595	309	48.3%

Table 3. Modelled daily HVAC energy (chiller + fresh-air fan) for the waiting hall, mean of 30 test days, under three assumptions about how much fresh air the existing fixed-schedule system supplies.

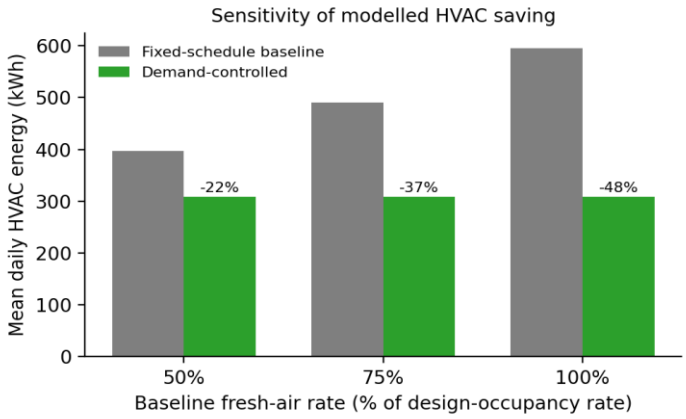


Figure 2. Sensitivity of the modelled saving to the baseline ventilation assumption.

The saving depends mainly on how much outdoor air the existing system conditions when the hall is lightly occupied. If the baseline supplies half the design-occupancy rate all day, the modelled saving is 22%; if it supplies the full design rate, 48%. We therefore recommend quoting the conservative value, and treat the upper values as bounds that field metering must confirm. In every case the share of occupied person-minutes outside 22–26 °C was 0.0% for both strategies. Energy per visitor fell from 0.265 to 0.138 kWh in the full-rate case.

5.3.2 Demand-Driven Environmental Optimization

Analysis of space cooling efficiency must evaluate thermal comfort in conjunction with real-time occupancy, outdoor temperature, relative humidity, and operating schedules. Potential low-cost optimization strategies include adjusting pre-cooling start times, implementing temperature setpoint deadbands, and optimizing off-peak operational schedules. Environmental sensor logs and comfort complaints must be audited alongside electrical billing records to prevent unintended comfort degradation.[4]

5.4 Water Resource Management and Biological Reuse

5.4.1 Water Balance Analysis

A comprehensive water balance tracks incoming municipal or freshwater supply volume, treated wastewater volumes, and end-use allocations (e.g., washrooms, surface cleaning, landscape irrigation). Unexplained volume discrepancies prompt investigations for underground pipe leaks, meter calibration errors, or unrecorded storage adjustments. Visitor-driven domestic water demand and landscape irrigation demand follow distinct operational patterns

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

and require separate analytical explanations.

5.4.2 Biological Wastewater Treatment, Quality Assurance, and Reuse

Biological wastewater treatment processes utilize living microorganisms to decompose organic pollutants in liquid waste streams. The selected process architecture must align with influent flow variations, raw water quality, site land availability, and operational capabilities. This paper outlines an analytical monitoring framework and does not assert that a specific treatment plant configuration has been constructed.

Treated effluent intended for landscape irrigation must comply with applicable water quality standards. Additional tertiary filtration, disinfection (e.g., chlorination or UV treatment), and continuous testing are required to ensure safety. Parameters such as pH, total dissolved solids (TDS), and temperature provide operational monitoring of plant health but do not substitute for regular microbiological safety testing. Physical separation and color-coded pipe identification are mandatory to eliminate cross-connection risks between potable and recycled water lines.

5.5 Staff Deployment and Asset Reliability

Correlating historical visitor arrival curves with labor deployment enables scheduling staggered shift start times and floating relief coverage during peak traffic windows. Labor metrics should track staff-hours alongside completed preventative maintenance tasks, work order response times, and unresolved fault backlogs. Optimizing labor allocation can free up workforce capacity for proactive asset maintenance without increasing total payroll expenses.

6. VERIFICATION FRAMEWORK AND ECONOMIC ASSESSMENT

6.1 Pilot Implementation and Baseline Protocols

A structured pilot deployment requires establishing a baseline phase of four representative operational weeks prior to intervention, followed by a four-week post-intervention assessment phase. The evaluation period must incorporate typical weekdays, weekends, and peak holiday traffic. If unusual weather patterns, volatile demand swings, or data collection gaps occur, the evaluation window must be extended. These timeframes represent minimal evaluation thresholds rather than a guarantee of statistical rigor.

6.2 Illustrative Energy Reduction Calculation

Verified energy savings are calculated by subtracting measured post-intervention energy consumption from baseline energy consumption, adjusted for variations in attendance and weather conditions:

$$\textit{Verified Energy Reduction (kWh)} = \textit{Adjusted Baseline Energy (kWh)} - \textit{Post-Intervention Energy (kWh)}$$

For illustrative purposes, if a facility section demonstrates an adjusted baseline consumption of 1,000 kWh during an evaluation window and consumes 900 kWh following control optimization, the verified reduction equals 100 kWh, representing a 10% relative energy savings. This calculation demonstrates the audit methodology and does not represent field-measured empirical data.

6.3 Financial Evaluation, ROI, and Limitations

Net financial benefits are calculated by subtracting incremental operational expenses (e.g., software licensing, sensor calibration, water testing, training) from verified utility cost avoidances. The Simple Payback Period is derived by dividing the initial capital investment by the annual net financial benefit:

$$\textit{Simple Payback Period (Years)} = \textit{Initial Capital Investment ()} / \textit{NetAnnualFinancialBenefit(/Year)}$$

Comprehensive long-term financial analyses should also account for equipment replacement cycles, asset lifecycle extension, and the time value of money via Net Present Value (NPV) metrics. Financial reporting must carefully avoid double-counting utility savings, and labor redeployment efficiencies must be reported separately from direct payroll reductions. Overall, this paper provides a conceptual, structured framework; empirical field testing is required to validate practical financial returns and operational limitations.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

7. CONCLUSION

This paper establishes a data-driven facility management architecture for high-density visitor venues by unifying predictive machine learning, $\boldsymbol{M/M/c}$ queuing theory, Isolation Forest fault detection, and BACnet/Modbus closed-loop controls within the MADOVS improvement cycle. By replacing reactive operations with mathematical formalizations for arrival forecasting, wait-time estimation, and weather-normalized baseline adjustments, the framework optimizes utility efficiency while safeguarding visitor comfort and safety. Artificial intelligence anticipates crowd surges and operational anomalies, enabling automated, pre-approved adjustments within strict safety boundaries. Robust IoT data collection, calibrated controls, and managerial oversight remain essential for translating this supervisory methodology into sustainable real-world venue operations.

REFERENCES:

[1] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 785–794). Association for Computing Machinery. https://doi.org/10.1145/2939672.2939785

[2] Krishnakumari, P., Hoogendoorn-Lanser, S., Steenbakkers, J., & Hoogendoorn, S. (2023). Crowd Safety Manager: Towards data-driven active decision support for planning and control of crowd events. arXiv. https://doi.org/10.48550/arXiv.2308.00076

[3] Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. In Proceedings of the Eighth IEEE International Conference on Data Mining (pp. 413–422). IEEE. https://doi.org/10.1109/ICDM.2008.17

[4] Moghimi, S. M., Gulliver, T. A., & Thirumai Chelvan, I. (2024). Energy management in modern buildings based on demand prediction and machine learning—A review. Energies, 17(3), Article 555. https://doi.org/10.3390/en17030555

[5] Shi, Q., Liu, C., & Xiao, C. (2022). Machine learning in building energy management: A critical review and future directions. Frontiers of Engineering Management, 9, 239–256. https://doi.org/10.1007/s42524-021-0181-1

(This work is licensed under a Creative Commons Attribution 4.0 International License.)