TRACK 4 - INTERNET OF THINGS (IoT)

⚠️ **IMPORTANT: Hardware Requirements & Disqualification Rules**

Bring Your Own Hardware: Teams must supply all necessary components. No hardware will be provided at the venue.

100% On-Site Assembly: Components must arrive completely unassembled. Any team found bringing a partially or fully pre-built project will face **immediate disqualification** upon inspection.

HTH-IT-01: Capacity-Aware Smart Building Security Orchestrator

Category	Details
Industry	Facilities / Corporate Security
Scenario	A building security team has multiple simulated IoT devices (door sensors, cameras, locks) but is overwhelmed by raw alerts and needs prioritized, explainable ones instead.
Problem Statement	Aggregate multi-device IoT signals into a single dashboard, detect unusual activity patterns (e.g., door access at odd hours combined with no corresponding badge swipe), and prioritize alerts by risk level given limited security-staff response capacity.
What Teams Should Build	Multi-device data ingestion/simulation, cross-signal correlation rule/ML engine, risk-scored alert prioritization under fixed responder capacity.
24-Hour MVP	Simulated multi-device event stream (door, badge, camera-motion), correlation-based anomaly detection, prioritized alert dashboard with only N alerts actionable per shift.
Data / Input	Simulated/synthetic IoT event streams generated by the team (documented generation logic).
Key Constraints	Must correlate across device types, not alert on single-sensor triggers alone; false-alarm rate must be visibly tracked.
Expected Output	Live prioritized alert dashboard with correlation rationale.
Evaluation Criteria	Cross-signal correlation quality, prioritization under capacity constraint, false-alarm handling.
Bonus Features	Auto-lockdown simulation for highest-severity tier with human confirmation step.

HTH-IT-02: Irrigation ROI Optimizer for Water-Constrained Farms

Category	Details
Industry	AgriTech
Scenario	Farms in water-scarce regions must allocate a limited water budget across multiple fields/crops with different water-sensitivity and yield value.
Problem Statement	Use simulated soil/weather sensor data to recommend an irrigation schedule that maximizes yield outcome under a fixed total water budget across multiple fields, not just per-field optimal watering.
What Teams Should Build	Sensor-data ingestion/simulation, per-field irrigation need model, budget-constrained multi-field water allocation optimizer.
24-Hour MVP	Simulated multi-field sensor data (5+ fields, different crop types), generate a water-allocation schedule under a fixed daily water budget with projected yield impact.
Data / Input	Simulated soil-moisture/weather data (documented generation method) or public agri sensor datasets.
Key Constraints	Total water is fixed and insufficient for all fields' ideal need — must trade off; crop water-sensitivity must differ by field.
Expected Output	Multi-field irrigation schedule with allocation rationale and projected yield impact.
Evaluation Criteria	Allocation optimization quality, realistic constraint modeling, projected impact clarity.
Bonus Features	Weather-forecast-responsive re-allocation mid-demo.

HTH-IT-03: Wearable Health Triage & Escalation Simulator

Category	Details
Industry	Digital Health / Remote Patient Monitoring
Scenario	Remote-monitoring nurses cover many patients simultaneously and cannot review every reading manually; they need a triage system that escalates only genuinely concerning patterns.
Problem Statement	Simulate wearable vitals (heart rate, SpO2, temperature) for multiple patients, detect abnormal patterns (not just threshold breaches — trend-based deterioration), and generate a triage queue under a limited nurse review capacity.
What Teams Should Build	Multi-patient simulated vitals stream, trend-based anomaly detection (not single-point thresholds), capacity-constrained triage queue ranked by clinical urgency.
24-Hour MVP	Simulated multi-patient vitals stream (10+ patients), trend-based deterioration detection, triage dashboard ranking patients under a fixed nurse-capacity limit (e.g., 5 reviews/hour).
Data / Input	Simulated wearable vitals generator (documented, physiologically plausible ranges).
Key Constraints	Must detect trend deterioration, not just static thresholds; nurse capacity fixed — must rank by urgency, not alert on everyone.
Expected Output	Triage dashboard with ranked patients, trend charts, and escalation rationale.
Evaluation Criteria	Trend-detection quality, capacity-aware triage logic, clinical plausibility.
Bonus Features	Simulated nurse acknowledgment loop with response-time tracking.

HTH-IT-04: Fleet-Wide Waste Collection Route Optimizer Under Budget

Category	Details
Industry	Municipal Services / Waste Management
Scenario	A city has a fixed number of collection trucks/hours per day and must decide which bins to service to minimize overflow risk city-wide.
Problem Statement	Use simulated fill-level sensor data across many bins to generate a daily collection route that stays within a fixed truck-hour budget while minimizing overflow risk and total distance.
What Teams Should Build	Simulated multi-bin fill-level data, overflow-risk prediction, budget-constrained route optimizer (vehicle routing under time/capacity limits).
24-Hour MVP	Simulated fill-level data for 30–50 bins, generate optimized routes for a fixed number of trucks/hours, compare against naive fixed-schedule collection.
Data / Input	Simulated ultrasonic sensor fill-level data (documented generation).
Key Constraints	Truck-hours/trucks fixed and insufficient to service every bin daily — must prioritize by overflow risk; must show improvement vs. naive baseline.
Expected Output	Daily route plan per truck with overflow-risk-based prioritization and comparison to baseline.
Evaluation Criteria	Optimization quality, realistic constraint handling, quantified improvement over baseline.
Bonus Features	Real-time re-routing when a bin's fill-rate spikes mid-demo.

HTH-IT-05: Network Access Risk Scorer for Unmanaged IoT Fleets

Category	Details
Industry	Enterprise IT / OT Security
Scenario	Enterprises with many IoT devices (cameras, sensors, badge readers) struggle to tell a new legitimate device from a rogue/compromised one joining the network.
Problem Statement	Fingerprint devices by traffic pattern/MAC/type, and instead of a binary "known/unknown" flag, produce a risk score that weighs device behavior anomaly, criticality of the network segment, and confidence — then recommend quarantine/monitor/allow.
What Teams Should Build	Device fingerprinting from simulated traffic logs, behavioral-anomaly scoring, segment-criticality-weighted risk scoring, recommended action tiering.
24-Hour MVP	Simulated network traffic logs with injected rogue-device scenarios, fingerprinting + risk scoring + tiered recommendation dashboard.
Data / Input	Simulated/synthetic network traffic logs (documented generation) or public IoT network traffic datasets (e.g., IoT-23).
Key Constraints	Must distinguish "new legitimate device" from "intrusion" with more than a static allowlist; segment criticality must affect scoring (a rogue device on a badge-reader segment is worse than on a lobby-display segment).
Expected Output	Risk-scored device inventory dashboard with recommended action per device.
Evaluation Criteria	Fingerprinting accuracy, risk-scoring logic sophistication, false-positive handling.
Bonus Features	Live simulation of a rogue device joining mid-demo and triggering escalation.

HTH-IT-06: Peak-Shaving Energy Control System for Multi-Device Facilities

Category	Details
Industry	Energy Management
Scenario	Facilities pay demand charges for peak power draw; operators need to actively shed/shift load in real time, not just monitor consumption.
Problem Statement	Track real-time simulated power draw from multiple devices, predict imminent peak-threshold breach, and automatically (with override) shed/shift non-critical device load to stay under a configurable peak cap.
What Teams Should Build	Real-time simulated multi-device power stream, peak-prediction logic, automated load-shedding decision engine respecting critical-vs-shiftable device tags.
24-Hour MVP	Simulated multi-device power stream (10+ devices, some marked critical), live dashboard with automatic shed/shift actions and manual override, before/after peak comparison.
Data / Input	Simulated device power-draw data (documented generation, plausible load profiles).
Key Constraints	Critical devices must never be auto-shed; peak cap configurable; must show quantified $ savings from demand-charge avoidance.
Expected Output	Live control dashboard with shed/shift log and quantified peak-charge savings.
Evaluation Criteria	Real-time responsiveness, constraint respect (critical-device protection), quantified business impact.
Bonus Features	Predictive pre-emptive shedding (act before breach, not after).

HTH-IT-07: Dynamic Multi-Lot Parking Revenue & Availability Optimizer

Category	Details
Industry	Smart City / Parking Operations
Scenario	Parking operators managing multiple lots want to both show real-time availability and dynamically price/redirect demand to underused lots to maximize revenue and reduce search congestion.
Problem Statement	Use simulated occupancy sensor data across multiple lots to show real-time availability, predict near-term fill trend per lot, and recommend dynamic pricing/redirection to balance utilization across the lot network.
What Teams Should Build	Simulated multi-lot occupancy stream, short-horizon occupancy forecasting, dynamic pricing/redirection recommendation engine.
24-Hour MVP	Simulated occupancy data for 4–6 lots, real-time dashboard + forecast + pricing/redirection suggestions to balance load.
Data / Input	Simulated occupancy sensor data (documented generation) or public parking occupancy open datasets.
Key Constraints	Must optimize across a network of lots, not single-lot monitoring; pricing recommendations must respond to predicted trend, not just current state.
Expected Output	Network dashboard with per-lot availability, forecast, and dynamic price/redirect suggestions.
Evaluation Criteria	Forecast quality, network-level optimization logic, practical UI.
Bonus Features	Driver-facing suggestion ("Lot B is 5 min further but guarantees a spot").

HTH-IT-08: Multi-Zone Air Quality Exposure & Response Advisor

Category	Details
Industry	Environmental / Public Health
Scenario	Cities/campuses with multiple air-quality sensors need to translate raw readings into concrete public-health actions (school recess cancellation, traffic advisories) rather than just displaying numbers.
Problem Statement	Aggregate simulated multi-location air-quality data (CO2, PM2.5, humidity), predict short-term trend, and generate zone-specific response recommendations tiered by health-risk threshold and population sensitivity (e.g., near a school vs. an industrial zone).
What Teams Should Build	Multi-zone simulated AQ data stream, short-horizon trend prediction, population-sensitivity-weighted response-tier recommender.
24-Hour MVP	Simulated AQ data across 5+ zones with different population-sensitivity tags, live dashboard + trend prediction + tiered response recommendations.
Data / Input	Simulated air-quality sensor data (documented generation) or public open AQ datasets (e.g., OpenAQ).
Key Constraints	Response recommendation must weight population sensitivity, not just raw pollutant level; must predict trend, not just alert on current threshold breach.
Expected Output	Zone-tiered dashboard with trend charts and recommended public actions.
Evaluation Criteria	Trend-prediction quality, sensitivity-aware recommendation logic, practical clarity.
Bonus Features	Historical exposure accumulation tracking per zone (chronic exposure risk, not just instantaneous).

HTH-IT-09: Leak-Loss Prioritized Water Network Monitor

Category	Details
Industry	Utilities / Water Management
Scenario	Water utilities lose significant revenue/resource to undetected leaks across large pipeline networks and must prioritize repair crews.
Problem Statement	Use simulated flow/level sensor data across a pipeline network to detect leak/overflow signatures, estimate water-loss rate per detected leak, and prioritize repair dispatch under a limited number of repair crews per day.
What Teams Should Build	Simulated multi-node flow/level data stream, leak-signature detection (flow imbalance patterns), loss-rate estimation, crew-capacity-constrained dispatch prioritization.
24-Hour MVP	Simulated pipeline network (10+ nodes) with injected leak scenarios, detection + estimated loss rate + prioritized dispatch list under a fixed crew count.
Data / Input	Simulated flow/level sensor data (documented generation with plausible leak-injection logic).
Key Constraints	Crew capacity fixed and smaller than detected issues — must prioritize by estimated loss rate × time-to-worsen; must distinguish leak signatures from normal demand fluctuation.
Expected Output	Ranked repair-dispatch list with estimated loss rate and network map.
Evaluation Criteria	Detection accuracy, loss-estimation realism, capacity-aware prioritization.
Bonus Features	Cumulative water/revenue loss tracker if a leak goes unaddressed.

HTH-IT-10: Fleet-Wide IoT Firmware Risk & Patch Prioritization Dashboard

Category	Details
Industry	IT/OT Security Operations
Scenario	Enterprises with hundreds of IoT devices can't patch everything at once and need to prioritize by actual risk, not just "outdated firmware = bad."
Problem Statement	Track firmware version, uptime, and health across a simulated device fleet, cross-reference known vulnerability severity per firmware version, and generate a patch-priority queue under a limited weekly patching capacity, weighted by device criticality and exposure.
What Teams Should Build	Simulated fleet inventory with firmware/version metadata, vulnerability-severity mapping (using a synthetic or public CVE-style reference list), criticality/exposure-weighted patch-priority ranking under fixed patching capacity.
24-Hour MVP	Simulated fleet of 50+ devices with varied firmware versions, vulnerability mapping, ranked patch queue under a fixed weekly patch-slot constraint.
Data / Input	Simulated device fleet metadata + synthetic/public vulnerability-severity reference data.
Key Constraints	Patch capacity fixed and smaller than vulnerable-device count — must prioritize by risk (severity × criticality × exposure), not just "oldest firmware first."
Expected Output	Ranked patch-priority dashboard with risk rationale per device.
Evaluation Criteria	Risk-scoring sophistication, capacity-aware prioritization, practical dashboard clarity.
Bonus Features	Simulated "patch now" action that updates fleet risk posture live.