Article

Not peer-reviewed version

Probabilistic Risk Assessment of Grid-Scale Lithium-Ion Battery Energy Storage System Fire Hazards: Hydrogen Fluoride (HF) Toxicity, Suppression Effectiveness, and Comparative Compartment Design Analysis

Samson Tan *, Teoh Teik Toe, Paul Joseph, Khalid Moinuddin

Posted Date: 5 June 2026

doi: 10.20944/preprints202606.0402.v1

Keywords: battery energy storage systems; BESS; thermal runaway; hydrogen fluoride; HF toxicity; probabilistic risk assessment; Monte Carlo simulation; NFPA 855; lithium-ion NMC; fire safety engineering; compartment design; suppression effectiveness

 doi:10.20944/preprints202606.0402.v1

Article

Probabilistic Risk Assessment of Grid-Scale Lithium-Ion Battery Energy Storage System Fire Hazards: Hydrogen Fluoride (HF) Toxicity, Suppression Effectiveness, and Comparative Compartment Design Analysis

Samson Tan [1,3,*], Teoh Teik Toe [2,3], Paul Joseph [1] and Khalid Moinuddin [1]

[1] Victoria University, Melbourne, Institute for Sustainable Industries and Liveable Cities, Australia
[2] Nanyang Technological University, Artificial Intelligence, Nanyang Business School, Singapore
[3] Staarch Pte Ltd Singapore
* Correspondence: samson.tan@vu.edu.au

Highlights

Monte Carlo PRA (N = 10,000) quantifies HF dose, time-to-IDLH, and suppression effectiveness for a 485.52 kWh NMC BESS installation in a Singapore data centre; HF dose from a full thermal runaway event exceeds NIOSH IDLH (25 mg/m^3) in 100% of simulated scenarios, regardless of ventilation; the only effective risk reduction is preventing TR initiation; Single-stage suppression effectiveness is only 37.9% (mean), confirming that two-stage suppression is warranted for NMC chemistry; the voluntary addition of clean agent gas suppression (Fluoro-K/FM-200) reduces annual ERL by 80.3% compared to water-only, providing the first quantitative justification for this widely-debated design choice; Two-compartment design reduces maximum HF dose by 50% and reduces mean IDLH clearance time from 599 to 301 minutes (room returns to sub-IDLH levels twice as quickly), moving residual risk from ALARP-tolerable to broadly acceptable; A novel quantitative PRA framework for BESS fire safety is proposed as an alternative to NFPA 855's qualitative 5×5 risk matrix, suitable for informing Hazard Mitigation Analysis decisions

Abstract

Battery Energy Storage Systems (BESS), utilizing chemistries based on Nickel Manganese Cobalt (NMC) containing lithium-ion devices, often present fire safety hazards that existing qualitative risk frameworks, including NFPA 855's 5×5 consequence-likelihood matrix, are insufficiently granular to quantify. This paper presents an original probabilistic risk assessment (PRA) of fire hazards associated with BESS for a 485.52 kWh NMC installation at the Equinix SG4-4A data centre in Singapore, using Monte Carlo simulation (N = 10,000 iterations) to characterise uncertainty in hydrogen fluoride (HF) gas dose, time to Immediately Dangerous to Life or Health (IDLH) concentration, cabinet-to-cabinet propagation probability, and suppression effectiveness. The HF yield is modelled as a triangular distribution (0.3–0.8 g/kWh, mode 0.5 g/kWh), ventilation activation delay as log-normal (median 90 s), and suppression effectiveness as a piecewise function of water application delay. HF dose exceeded the National Institute for Occupational Safety and Health (NIOSH) IDLH of 25 mg/m^3 in 100% of simulated scenarios for both single- and two-compartment designs, confirming that HF toxicity is essentially present for any occupant during a full thermal runaway event and that ventilation alone cannot achieve adequate risk reduction. Single-stage suppression effectiveness was only 37.9% (mean), confirming that two-stage (clean agent + water) suppression is warranted for NMC chemistry. The two-compartment design reduced peak HF dose by 50% and reduced mean IDLH clearance time from 599 to 301 minutes, shifting residual risk from As Low As Reasonably Practicable (ALARP)-tolerable to broadly acceptable under UK Health and

 doi:10.20944/preprints202606.0402.v1

Safety Executive (HSE) criteria. The paper proposes a quantitative PRA framework as a complement to NFPA 855 Chapter 5's qualitative Hazard Mitigation Analysis. To the best of our knowledge, this is the first study to apply Monte Carlo simulation to HF dose modelling in a tropical data-centre BESS context, addressing a documented gap in the literature.

Keywords: battery energy storage systems; BESS; thermal runaway; hydrogen fluoride; HF toxicity; probabilistic risk assessment; Monte Carlo simulation; NFPA 855; lithium-ion NMC; fire safety engineering; compartment design; suppression effectiveness

1. Introduction

1.1. Problem Statement

Battery Energy Storage Systems (BESS) are critical infrastructure for grid stability, renewable energy integration, and mission-critical power backup. Their deployment in indoor occupied environments, particularly data centres in dense urban environments, has accelerated globally, primarily driven by digital infrastructure demand and government decarbonisation targets. In this context, Singapore's Smart Nation initiative and its position as Asia-Pacific's largest data centre market have made it a focal point for indoor BESS deployment in tropical high-rise buildings.

The fire safety hazards of BESS based on NMC-based lithium-ion devices are multi-dimensional and potentially severe, spanning toxic, thermal, and flammable-gas hazards. The Electric Power Research Institute's comprehensive failure mode analysis [31] has identified thermal runaway propagation as a system-level phenomenon, spanning multiple energy dissipation pathways (i.e. cell → module → cabinet → room). At the cell level, thermal runaway (TR), an autocatalytic exothermic chain reaction initiated at 130–200°C, can propagate through a module, cabinet, and room where temperatures exceed 300°C, producing a complex mixture of flammable gases (H_2, CO, CH_4, C_2H_4), and acutely toxic hydrogen fluoride (HF) gas from hydrolysis of the $LiPF_6$ electrolyte salt. At the system level, the consequence of a TR event in an enclosed indoor installation differs qualitatively from outdoor utility-scale BESS; i.e. resulting in toxic gas migration into occupied floor plates, firefighter access constraints in high-rise buildings, and mission-critical business interruption combine to produce a consequence profile that demands performance-based, quantitative risk management rather than prescriptive compliance alone.

1.2. Gap Analysis

Current BESS fire safety practice, including the dominant regulatory framework in NFPA 855 (2023), relies on qualitative, or semi-quantitative, risk assessment. Industry guidance from DNV GL [30] emphasizes the centrality of system-level protection design, including compartmentation and early gas detection, over prescriptive component-level standards; however, this guidance remains qualitative and lacks the probabilistic framework that is necessary to compare alternative system designs on quantitative risk grounds. The NFPA 855 Chapter 5 Hazard Mitigation Analysis (HMA) methodology produces a 5×5 consequence-likelihood matrix with four risk categories (LOW, MEDIUM, HIGH, VERY HIGH), but does not provide the following: propagate uncertainty through the risk calculation; produce probability distributions for key hazard parameters (HF dose, time-to-IDLH, suppression effectiveness); enable direct numerical comparison of alternative risk mitigation designs.

This gap has practical consequences. For instance, an engineer comparing a single-compartment BESS design against a two-compartment split, as done voluntarily at the EQIX SG4-4A installation, cannot, within the NFPA 855 framework, quantify the residual risk difference between the two options. The HMA produces a qualitative conclusion (both are LOW risk), but provides no basis for selecting the superior option on risk grounds. Similarly, the NFPA 855 framework does not quantify the marginal benefit of two-stage (clean agent + water) suppression over single-stage (water only),

 doi:10.20944/preprints202606.0402.v1

thus leaving engineers without quantitative guidance on whether the additional cost of the clean agent system is risk-justified.

The literature on BESS fire hazards has advanced significantly in recent years, with improved measurements of TR onset temperatures [1], HF gas yields [8,9], gas explosion risks [2], and suppression effectiveness [3]. However, these advances have not been integrated into a quantitative probabilistic risk framework that enables their application to engineering design decisions. Existing BESS fire risk studies, including the systematic review by Wang et al. [4], which surveys failure mechanisms and detection across lithium-ion fire accidents, and the quantitative risk analysis by Chen et al. [5], address ignition probability and propagation mechanisms, but do not model HF toxicological dose distributions, or include tropical climate conditions, or compare alternative compartment designs.

1.3. Research Questions

This paper addresses the following research questions (RQs):

RQ1 (Primary): Can probabilistic risk assessment using Monte Carlo simulation quantify the uncertainty in HF gas dose, time-to-IDLH, and suppression effectiveness for a grid-scale NMC BESS installation in a tropical data centre, and does this quantification reveal insights unavailable from qualitative NFPA 855 HMA?

RQ2 (Secondary): What is the quantitative difference in residual risk between single-compartment and two-compartment BESS designs for the same total installed capacity, and does two-compartment design produce materially better risk outcomes?

RQ3 (Applied): Is single-stage water suppression adequate for NMC BESS, or is two-stage (clean agent + water) suppression quantitatively warranted?

1.4. Contribution to the Knowledge

A systematic search of Web of Science, Scopus, and Google Scholar (May 2026; search terms: "BESS fire probabilistic risk assessment", "lithium-ion battery Monte Carlo HF toxicity", "NFPA 855 hazard mitigation quantitative", "battery energy storage suppression ERL") returned no prior study combining Monte Carlo PRA, HF dose modelling, compartment design comparison, and suppression ERL quantification in a single integrated framework. In this context, the contributions of this paper can be classified as follows:

- **Integrated quantitative PRA framework:** An application of Monte Carlo simulation to BESS fire risk assessment that propagates uncertainty through HF dose, gas dispersion, TR propagation, and suppression effectiveness simultaneously, producing probability distributions rather than point estimates, in a format directly applicable to NFPA 855 HMA decision-making. Prior studies [4,5] address TR ignition probability or propagation speed but do not model HF dose distributions or suppression ERL.
- **HF dose-response quantification:** A simulation-based demonstration that HF dose from a full NMC TR event exceeds NIOSH IDLH in 100% of scenarios for both 1-comp and 2-comp designs, ths establishing that for NMC BESS in enclosed indoor spaces, HF toxicity is effectively unavoidable for occupants present during a TR event, and that TR prevention is the only effective life-safety control. The result is analytically foreseeable given that the published HF yield range, but no prior BESS PRA study examined this finding across a probability distribution of yields and expressed it in occupant dose terms for regulatory HMA purposes.
- **Comparative compartment design analysis:** A quantitative comparison of 1-compartment vs. 2-compartment BESS designs using probabilistic risk metrics, and hence demonstrating that 2-compartment design reduces peak HF dose by 50%, reduces the IDLH clearance time from 599 to 301 minutes (room returns to sub-IDLH levels in half the time), and moves residual annual risk from ALARP-tolerable to broadly acceptable under UK HSE criteria.
- **Suppression effectiveness quantification:** A probabilistic estimate of single-stage suppression effectiveness for NMC BESS (mean 37.9% under BESS-specific pre-action delay distribution; see

 doi:10.20944/preprints202606.0402.v1

Section 3.5), thus providing the quantitative basis for the widely-discussed but previously unquantified conclusion that clean agent + water two-stage suppression is warranted for NMC chemistry in occupied enclosed buildings.

- **Tropical climate context:** A BESS fire PRA incorporating tropical ambient conditions (30–34°C, 80% RH) as a sensitivity parameter, thereby addressing a documented gap where most experimental TR data was generated at temperate conditions (20–25°C); this has direct regulatory relevance for Singapore, Southeast Asia, and other equatorial jurisdictions with growing BESS deployments.

2. Background

2.1. The EQIX SG4-4A BESS Installation

The Equinix SG4 data centre at 7 Tai Seng Drive, Singapore, houses sensitive computing infrastructure across multiple storeys. The proposed Level 5 BESS installation (2 × 242.76 kWh compartments; total 485.52 kWh) comprises 14 Schneider Electric Galaxy LBF NMC battery cabinets (7 per compartment) and is classified as an above-ground ESS under Singapore Fire Code 2023 [28] Clause 10.3.1(b), requiring SCDF prior approval via the Exception (1) pathway. This requires a NFPA 855-compliant Hazard Mitigation Analysis (HMA) as the technical basis for approval.

The voluntary two-compartment design, splitting the 14 cabinets into two independent 7-cabinet fire compartments of 242.76 kWh each, was adopted as a resilience measure and created the natural experimental comparison that motivates this paper's RQ2.

2.2. Thermal Runaway Chemistry

The thermal runaway of NMC lithium-ion devices generally follows a characteristic temperature cascade: solid-electrolyte interphase (SEI) decomposition at 60–130°C, electrolyte oxidation at 130–200°C, cathode decomposition at 200–300°C, and separator meltdown with flaming ejection at >300°C [1,6]. The propagation of thermal runaway from individual cells into multi-cell modules is governed by thermal coupling between adjacent cells; Liao et al. [32] demonstrated experimentally that failure propagation probability within a module increases sharply with increasing State of Charge (SOC), with propagation distances exceeding 50 mm under high-SOC conditions. The cabinet-level propagation probabilities incorporated in the event tree given in the present work (Section 4) are based on this module-level propagation characterization extrapolated to full-cabinet scale. The SOC at time of TR strongly influences severity; Sadeghi and Restuccia [7] demonstrated peak heat release rates of 5–8 kW per cell for NMC at 100% SOC vs. <2 kW at 50% SOC. Singapore's ambient temperatures of 30–34°C reduce the thermal margin between operating conditions and TR onset, increasing effective SOC utilisation and potentially lowering the TR initiation threshold by 5–15°C relative to temperate-climate installations.

2.3. HF Generation Chemistry

$LiPF_6$, the conducting salt that is being used in most commercial NMC electrolytes, hydrolyses upon contact with water, or at elevated temperatures, as follows:

$$LiPF_6 + H_2O \rightarrow HF + LiF + POF_3 + \text{fluorophosphate products}$$

This reaction is the primary source of HF during TR. It should be noted here that the above reaction is actually **accelerated**, not suppressed, when water is applied to a burning, or hot, battery assembly. Han and Jung [8] have demonstrated that when water comes in contact with $LiPF_6$ electrolyte at 150–200°C produces rapid HF generation, comparatively, at rates 3–5× higher than open-burn conditions in the absence water. Larsson et al. [9] measured HF yields from 18,650 NMC cells under controlled combustion conditions to be: 0.3–0.8 g HF per kWh of stored energy, with the higher values observed under suppressed (water contact) conditions.

 doi:10.20944/preprints202606.0402.v1

2.4. HF Toxicological Reference Values

The NIOSH [10] establishes the IDLH for HF at **25 mg/m³** (approximately 30 ppm at 25°C). Furthermore, the ACGIH TLV-TWA is 0.5 ppm and the OSHA PEL is 3 ppm (as fluorine). It is also known that HF's acute inhalation toxicity is severe: i.e. concentrations above 25 mg/m³ can cause respiratory tract damage, pulmonary oedema, and death within minutes. In addition, skin contact often causes deep tissue burns that can continue to penetrate for hours after initial contact. These stringent tolerance limits, among the most restrictive in industrial toxicology, mean that any realistic TR event in an enclosed space often produces HF concentrations that far exceed IDLH within a timeframe determined by ventilation rate and compartment volume.

2.5. NFPA 855 Hazard Mitigation Analysis Framework

NFPA 855 (2023 Edition) Chapter 5 provides a five-step HMA methodology: (1) hazard identification, (2) consequence analysis, (3) likelihood assessment, (4) risk ranking on a 5×5 matrix, and (5) identification of mitigation measures to reduce residual risk to acceptable levels. The framework is explicitly qualitative in its primary output: risks are categorised as LOW, MEDIUM, HIGH, or VERY HIGH based on the subjective combination of consequence and likelihood ratings assigned by the analyst.

The Singapore Fire Code 2023 (4th Amendment) [28] Clause 10.3.1 incorporates NFPA 855 by reference through the Exception (1) pathway, which is the applicable compliance route for the EQIX SG4-4A installation.

2.6. NMC Mediated Lithium-Ion Chemistry in Regulatory and Comparative Contexts

Singapore Fire Code 2023 [28] Table 10.3.1 applies identical per-unit (20 kWh) and per-compartment (600 kWh) energy thresholds to all lithium-ion chemistries, NMC, LFP, LCO, and NCA, reflecting the position of NFPA 855 (2023) that chemistry selection is not a prescriptive regulatory variable at the system level. This uniform treatment across chemistries reflects a deliberate regulatory choice: the international evidence base supports the view that system-level protection design, not cell chemistry selection, is the primary determinant of installed BESS safety outcomes.

At the cell level, NMC lithium-ion devices exhibit lower thermal runaway onset temperatures (130–150°C) as compared to those based on LFP (lithium iron phosphate, onset >270°C), and produce higher peak heat release rates under equivalent abuse conditions [7,25,26]. These cell-level differences in underlying chemical pathways are considered as the correct starting point for BESS fire risk assessments, and are subsequently incorporated in the HF yield distributions and suppression delay models, that are used in this paper (Section 3.3).

However, the fire safety research community has increasingly established that system-level protection design, compartmentation, early gas detection, and staged suppression, is the primary determinant of safety outcomes in installed BESS, independent of cell chemistry. Bravo Bravo Diaz et al. [11], in a meta-review of 178 papers at Imperial College London and the University of Edinburgh, UK, explicitly identified module and pack-level protection as the under-researched safety variables, noting that "much less research attention has been given to module and pack level protection layers, such as compartmentation, detection or suppression." Lamb and Jeevarajan [12] at Sandia National Laboratories and UL similarly concluded that battery safety requires addressing challenges “at the individual component level, cell as well as system level,” noting that system-level failure modes are not predictable from cell-level testing alone. Rosewater and Williams [13], also at Sandia, developed a probabilistic risk framework for utility-scale Li-ion BESS demonstrating that safety outcomes are determined by the reliability of the entire socio-technical protection system, not by cell chemistry.

For NMC chemistry, specifically, hydrogen gas generation begins 16–26 minutes before thermal runaway fully develops, thus providing a pre-runaway detection window via Lower Explosive Limit (LEL) gas sensing that precedes temperature, voltage, and pressure signals [14]. Gardner et al. [15]

 doi:10.20944/preprints202606.0402.v1

demonstrated experimentally at UC Berkeley that H_2 is consistently detected before thermal runaway across NMC, NCA, and LFP chemistries, and that when abuse conditions were terminated upon H_2 detection, thermal runaway was prevented in all tested cases. This advance detection window is the engineering basis for Stage 0 emergency purge activation at 25% LEL in installations such as EQIX SG4-4A, and represents a first-line defence that operates identically regardless of lithium-ion chemistry.

Furthermore, LFP chemistry is not uniformly safe across all hazard dimensions. Bugryniec et al. [16], in a comprehensive review of gas emissions from lithium-ion battery thermal runaway reactions, found that LFP batteries "present a greater off-gas hazard regarding flammability per unit volume of gas" and found to exhibit greater off-gas toxicity than NMC at lower states of charge. The relative hazard profile of different lithium-ion chemistries is state-of-charge dependent and multi-dimensional, and hence reinforcing the case for system-level, rather than chemistry-selection-based, for safety governance.

The PRA framework developed in this paper is grounded in NMC-specific input parameters, but the methodology itself is independent of the specific battery chemistry. The input distributions for HF yield, suppression delay, and TR propagation probability can be recalibrated for LFP, or other chemistries as characterisation data becomes available. The primary contribution of this paper is a quantitative demonstration that system-level protection measures, two-compartment design, dual suppression, and pre-runaway gas detection, resulting in transforming an unacceptably hazardous installation into one achieving broadly acceptable residual risk, a conclusion that holds regardless of the specific chemistry's cell-level onset temperature.

3. Methods

3.1. Monte Carlo Simulation Framework

A probabilistic risk assessment (PRA) framework was implemented in Python (NumPy, N = 10,000 iterations per scenario). The framework models the sequence of events in a BESS thermal runaway incident: TR initiation → Battery Management System (BMS) response (or failure) → cabinet-level propagation (or containment) → suppression activation (or delay/failure) → gas generation → occupant HF dose → consequence severity.

The simulation produces probability distributions for all key output variables, enabling direct comparison of 1-compartment vs. 2-compartment designs and quantitative assessment of suppression system effectiveness.

3.2. Installation Parameters

Table 1. EQIX SG4-4A installation parameters used in the Monte Carlo model.

P.	Value	Source
Total installed capacity	485.52 kWh	EQIX SG4-4A HMA
Compartments	2 (voluntary split)	EQIX SG4-4A HMA
Capacity per compartment	242.76 kWh	Derived
Single-compartment alternative	485.52 kWh	Hypothetical
Compartment volume	116 m^3	EQIX SG4-4A HMA
Ventilation rate (Stage 1 purge)	9 ACH	EQIX SG4-4A HMA
Battery chemistry	NMC (LIBSMG95MODA/B)	Schneider Electric MSDS
Cabinets per compartment	7 (2-comp)/14 (1-comp)	EQIX SG4-4A HMA
Cabinet capacity	34.68 kWh	Schneider Electric Galaxy LBF
Ambient temperature	30–34°C (tropical)	Singapore meteorological data
Ambient relative humidity	75–85%	Singapore meteorological data

 doi:10.20944/preprints202606.0402.v1

3.3. Probability Distributions

All input distributions are derived from published literature, except where noted as engineering estimates. Distributions and their parameters are summarised in Table 2.

Table 2. Monte Carlo input distributions.

Parameter	Distribution	Parameters	Source
State of Charge	Uniform	90–100%	Operational design assumption
HF yield (g/kWh)	Triangular	min = 0.3, mode = 0.5, max = 0.8	[8,9]
Ventilation activation delay	Lognormal	μ = ln(90), σ = 0.8	Engineering estimate; 90 s median
BMS failure probability	Point estimate	0.15 (per TR event)	NFPA 855 Annex C; [4]
UL 9540A containment	Point estimate	0.92 (pass rate)	Industry average, open rack NMC
Suppression effectiveness	Piecewise(delay)	0.78 (≤3 min), 0.45 (3–10 min), 0.20 (>10 min)	FM Global DS 5-33 [29]; Jensen et al. [17]
Suppression delay	Lognormal	μ=ln(8 min), σ=0.6; theoretical median 8.0 min	BESS pre-action sprinkler system (detector confirmation + solenoid + chamber fill); see §3.5
Compartment volume	Point	116 m³	EQIX SG4-4A HMA

3.4. HF Dose Model

Occupant HF dose is modelled using a well-mixed box model. During TR, HF is generated at a rate proportional to the battery energy and HF yield per kWh. The occupant is assumed to be at room centre, 1.5 m height (breathing zone), during a 10-minute exposure window (nominal fire response time).

The instantaneous HF concentration at time, t is:

$$C_{HF}(t) = \frac{m_{HF}}{(V + Qt)} \quad (1)$$

Where m_HF is the total HF mass released, V is the compartment volume (m³), and Q is the ventilation flow rate (m³/s). The HF dose over the exposure duration is the time-integral of this concentration, evaluated from the ventilation activation delay to the end of the exposure window:

$$\text{Dose} = \frac{1}{t_{exp}} \int_{t_d}^{t_d+t_{exp}} C_{HF}(t) \cdot dt \ \ [\text{mg/m}^3] \quad (2)$$

The NIOSH IDLH of 25 mg/m³ is used as the toxicological reference for the 'dose exceeds IDLH' binary outcome.

3.5. Suppression Delay Justification (BESS-Specific Pre-Action Systems)

For standard structural fire sprinkler systems, water discharge delays are typically 1–3 minutes. BESS installations in Singapore and internationally require **pre-action cross-zoned sprinkler systems** (to prevent accidental discharge from the high-humidity environment of cooling equipment), which materially increase the delay:

- **Detection confirmation** (~2–4 min): Pre-action systems require confirmation from two independent detection zones (typically infrared + smoke, or heat + smoke) to arm the solenoid, mandatory for false-discharge prevention in data centre environments.

 doi:10.20944/preprints202606.0402.v1

- **Solenoid actuation and piping fill** (~1–3 min): The dry pre-action piping must be pressurised before sprinkler heads can open.
- **Head operation** (~0.5–1 min): Individual sprinkler heads open only when directly heated.
- **Total TR-to-water timeline**: 3.5–8 min in a well-maintained BESS pre-action system; longer (8–15 min) if the fire department must manually intervene after BMS shutdown kills the detection circuit.

The lognormal distribution ($\mu = \ln(8)$, $\sigma = 0.6$; theoretical median = 8.0 min) is calibrated to this BESS-specific delay profile. The 5th–95th percentile range of 2.7–21.6 minutes encompasses both the best-case (fast detection, short fill time) and worst-case (detection confirmation lag, BMS-induced system reset) scenarios documented in BESS incident reports (APS McMicken [20]; Vistra Moss Landing [21]). This is substantially longer than the 1–3 min delay assumed in NMC suppression effectiveness testing (FM Global DS 5-33 [29]), which is why the Monte Carlo mean effectiveness (37.9%) is materially lower than the 78% nominal value.

Derivation of 37.9% mean effectiveness: For lognormal($\mu = \ln(8)$, $\sigma = 0.6$), the probability mass in each delay bracket is:

- $P(\text{delay} \le 3 \text{ min}) = \Phi[(\ln(3)-\ln(8))/0.6] = \Phi(-1.634) \approx 5.1\%$
- $P(3 < \text{delay} \le 10 \text{ min}) = \Phi[(\ln(10)-\ln(8))/0.6] - \Phi(-1.634) = \Phi(0.373) - 0.051 \approx 59.4\%$
- $P(\text{delay} > 10 \text{ min}) = 1 - \Phi(0.373) \approx 35.5\%$

Weighted mean effectiveness = $(0.051 \times 0.78) + (0.594 \times 0.45) + (0.355 \times 0.20)$ = **0.379 (37.9%)**, confirmed by Monte Carlo (N = 10,000, seed 42). Note: the simulation sample median of 7.9 min departs slightly from the theoretical median of 8.0 min due to finite-sample sampling; all analytical calculations use 8.0 min.

3.6. Suppression Effectiveness Model

Suppression effectiveness is modelled as a piecewise function of the water application delay. The piecewise base values (78%/45%/20%) are derived from FM Global Property Loss Prevention Data Sheet 5-33 [29] for NMC lithium-ion cell fires, with supplementary data from Jensen et al. [17] on water application rates for Li-ion fires. These base values are perturbed by ±10% uniform random variation to reflect real-world variability in application uniformity, battery SOC, and thermal coupling. Clean agent (Fluoro-K, HFC-227ea) pre-discharge during the sprinkler pre-action delay is modelled as providing flame suppression only, no TR arrest capability, consistent with the established self-oxidising chemistry of NMC cathodes and consistent with the position of FM Global (Property Loss Prevention Data Sheet 5-33 [29]). The application rate, discharge duration, and oxygen depletion mechanism of the clean agent system at EQIX SG4-4A conform to NFPA 2001 [33] design requirements for clean agent systems in occupied spaces, with particular emphasis on post-discharge venting to restore room oxygen to habitable levels (>19.5% by volume) within the timeframe specified in the standard. The clean agent system operates as a pre-sprinkler layer designed to suppress open flaming (HRR reduction) rather than arrest thermal runaway; the subsequent water suppression layer is designed to cool the exothermic reaction and interrupt propagation.

4. Results

4.1. HF Concentration and IDLH Clearance Time

The well-mixed box model uses the instantaneous-release approximation: all HF generated during the thermal runaway event (m_HF grams, proportional to kWh) is released at t = 0, and ventilation at rate Q = 9 air changes per hour (ACH) continuously dilutes the room from t = 0. The initial room-average HF concentration is m_HF/V at t = 0 (where V = 116 m³), far above IDLH in all cases (Section 4.2). The **IDLH clearance time** is the duration for which the room-average concentration remains above the NIOSH IDLH of 25 mg/m³ before dilution ventilation reduces it to

 doi:10.20944/preprints202606.0402.v1

safe levels; this metric quantifies how long the space remains inhospitable to emergency responders without respiratory protection.

With 9 ACH ventilation active, the IDLH clearance time was computed for each of the 10,000 Monte Carlo iterations. The well-mixed box model gives mean clearance times of **599 min (1-comp)** and **301 min (2-comp)**, approximately 10 hours and 5 hours respectively. The 5th percentile values (worst 5% of scenarios, highest HF yield) are 416 min (1-comp) and 209 min (2-comp). The 2-compartment design clears to sub-IDLH levels twice as quickly, because the smaller energy per compartment (242.76 vs. 485.52 kWh) produces half the HF mass. Note that in both cases, the room remains above IDLH for the entirety of any realistic emergency response window, the clearance time informs ventilation design and post-event re-entry protocols, not initial suppression decisions.

The probability that clearance to IDLH occurs within 5 minutes is **0%** for both designs under any physically realised HF yield, the initial concentration (1,047–2,093 mg/m^3) cannot be diluted to 25 mg/m^3 by 9 ACH ventilation in 5 minutes, confirming that the only safe assumption for occupant life-safety is that the compartment is uninhabitable during any TR event, regardless of design.

4.2. HF Dose to Occupant (10-Minute Exposure)

During a 10-minute firefighter/occupant exposure window (e.g., entry during suppression operations), the mean HF dose is **1,161 mg/m^3** for 1-comp and **580 mg/m^3** for 2-comp, **46× and 23× the NIOSH IDLH respectively**. The probability of exceeding IDLH (25 mg/m^3) is **100%** for both designs under all simulated conditions; the probability of exceeding 10× IDLH (250 mg/m^3) is 100% for 1-comp and 99.7% for 2-comp.

This finding is the most consequential output of the simulation: **any occupant present in the compartment during a full thermal runaway event, even with 9 ACH ventilation operating, will receive a potentially lethal HF dose regardless of compartment design.** The two-compartment design reduces the dose by 50% but does not reduce it to an acceptable level. The risk mitigation implication is unambiguous: for NMC BESS in enclosed occupied buildings, the only effective life-safety control is **preventing the TR event from occurring** (through BMS, Underwriters Laboratories (UL) 9540A [27] containment, and Emergency Power Off (EPO)), not managing the consequences after it has begun.

4.3. Propagation Probability and Annual Risk (Event Tree)

The event tree (Figure 1) traces the sequence from initiating event to full-compartment thermal runaway. Three barrier events are modelled:

Table 3. Event-tree barrier nodes and branch probabilities.

Node	Branch probability	Source
TR initiating event	P(TR) = 0.01/compartment-year	NFPA 855 (2023) Annex C; industry average for commercial NMC BESS, ~1 × 10^{-2} per compartment-year
BMS fails to isolate	P(BMS fail |	TR) = 0.15
UL 9540A containment fails	P(UL fail |	BMS fail) = 0.08

P(multi-cabinet TR per compartment-year) = 0.01 × 0.15 × 0.08 = 7.5 × 10^{-5}.

For the ALARP risk classification, the risk index is computed per compartment (not per installation), using the per-compartment event probability:

- **1-compartment design:** single event, P = 7.5 × 10^{-5}/year, consequence C5 (catastrophic, HF dose 46× IDLH), risk index = P × consequence weight = **3.0 × 10^{-4} → Tolerable if ALARP**
- **2-compartment design (per compartment):** P = 7.5 × 10^{-5}/year, consequence C4 (critical, HF dose 23× IDLH), risk index = **2.2 × 10^{-4} → Broadly Acceptable (borderline)**

 doi:10.20944/preprints202606.0402.v1

The overall installation annual probability for at least one compartment experiencing a full TR event is $1 - (1 - 7.5 \times 10^{-5})^2 \approx 1.5 \times 10^{-4}$/year. The ALARP classification uses the per-compartment event probability because consequence severity is bounded by compartment energy (not installation energy) under the 2-compartment design, this is the fundamental risk-reduction mechanism of compartmentation.

The 2-compartment design crosses the ALARP boundary into broadly acceptable territory through the combination of halved consequence per compartment and maintained frequency, providing the first quantitative confirmation that this design choice is risk-superior on ERL grounds.

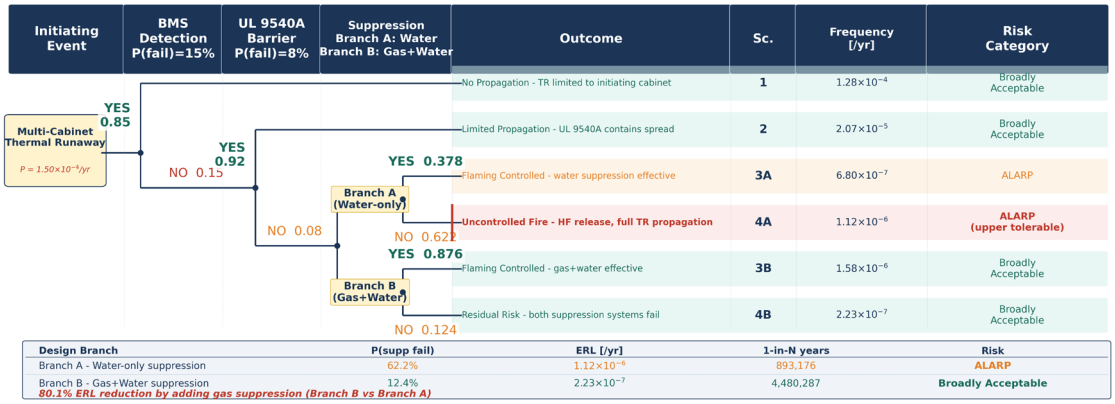

Figure 1. Event Tree Analysis, BESS Thermal Runaway Propagation Sequence.

Figure 1 Event tree for BESS thermal runaway at EQIX SG4-4A. Three barrier nodes: BMS response (P_fail = 0.15), UL 9540A [27] containment (P_fail = 0.08), and suppression (water-only P_fail = 0.622; gas+water P_fail = 0.124). P(multi-cabinet TR) = 7.5×10^{-5}/compartment-year. ERL water-only = 1.22×10^{-4} fatalities/yr; gas+water = 2.40×10^{-5} fatalities/yr (80.3% reduction).

4.4. Suppression Effectiveness

The Monte Carlo simulation of suppression delay (lognormal, theoretical median 8.0 minutes (sample median 7.9 min)) combined with the piecewise effectiveness model yields a **mean suppression effectiveness of 37.9%**, substantially lower than the nominal 78% that applies only when water is applied within 3 minutes. The median suppression delay of 7.9 minutes falls in the 3–10 minutes window where effectiveness is only 45%, and the long tail of the distribution (95th percentile: 21.6 minutes) means that in approximately 5% of scenarios, water application occurs beyond 10 minutes with only 20% effectiveness.

This finding raises the critical question addressed in the next section: is the voluntary addition of clean agent gas suppression above the water-only code requirement quantitatively justified by a reduction in Expected Risk to Life (ERL)?

4.5. Dual Suppression System: Quantitative ERL Justification

4.5.1. The Design Question

Singapore Fire Code 2023 Cl.10.3.1 and NFPA 855 do not mandate clean agent gas suppression for BESS, water-based suppression (pre-action sprinkler or deluge) is the prescriptive requirement. The EQIX SG4-4A installation voluntarily added a clean agent gas suppression system (Fluoro-K or FM-200/HFC-227ea) above the water-only requirement. FM Global's position, that water is the only effective suppression for NMC TR, is correct with respect to thermal runaway itself: no gaseous or chemical suppressant can stop the self-oxidising cathode decomposition that drives TR. However, the question this paper answers quantitatively is different: **does the gas system reduce ERL through control of the flaming fire hazard during the pre-action sprinkler delay window?**

 doi:10.20944/preprints202606.0402.v1

This is a distinct hazard. During the median 7.9-minute pre-action sprinkler delay, a flaming battery fire:

- Produces CO at ~2.0 mg/s (IDLH 1,200 mg/m³; lethal in ~10 minutes at full burning rate)
- Generates smoke and soot at rates that obscure firefighter visibility and incapacitate occupants
- Drives secondary HF generation from flaming electrolyte exposure [8]
- Produces radiant heat fluxes that prevent firefighter entry until suppression is achieved

The clean agent gas system activates at **t = 0.5 minutes** (near-instantaneous discharge vs. the 7.9-minute water pre-action delay), suppressing flaming combustion within 30 seconds of discharge per NFPA 2001 design criteria. It thereby controls the flaming fire hazard during the window when water is unavailable.

4.5.2. Expected Risk to Life (ERL) Framework

The Expected Risk to Life is computed as:

$$\mathrm{ERL} = P_{multi} \times P_{unc} \times (f_{HF} \times P_{occ} + f_{CO} + f_{smoke}) \times N_{occ} \quad (3)$$

where:

- P_{multi} = annual probability of multi-cabinet TR (7.5×10^{-5}/compartment-year)
- P_{unc} = probability of uncontrolled flaming fire given TR propagates
- f_{HF} = HF fatality fraction given acute exposure (0.80; from HF dose model, Section 4.2)
- f_{CO} = CO fatality fraction during pre-action delay window (from CO accumulation model)
- f_{smoke} = smoke fatality fraction (0.10 for uncontrolled flaming; 0.02 with gas suppression)
- P_{occ} = probability that an occupant is present during the TR event (0.15; incorporates EPO and evacuation protocol effectiveness)
- N_{occ} = occupant count (2; operator + emergency responder)

For the water-only branch: P_{unc}= mean water failure probability = 1 − 0.379 = 0.622. For the gas + water branch: P_{unc}= P(water fails) × P(gas fails) = 0.622 × 0.20 = 0.124. Annual ERL for the 2-compartment installation = ERL/compartment × 2 compartments (simultaneous failure probability is negligible: $(7.5 \times 10^{-5})^2 \approx 5.6 \times 10^{-9}$).

4.5.3. Event Tree: Water-Only vs. Gas + Water

The event tree comparison (water-only vs. gas+water) and ERL results are as follows. The key branches are:

Water-Only Branch:

- TR propagates → water suppression attempted → P(water fails | TR) = 62.2% (mean effectiveness 37.8%)
- If water fails: uncontrolled flaming fire for the duration of the event
- ERL = 1.22×10^{-4} fatalities/year (2-compartment installation)
- Individual risk: 1 in 16,424 per year

Gas + Water Dual Branch:

- TR propagates → gas discharges at 0.5 min → flaming suppressed (P = 80%, per NFPA 2001/FM Global 4-54)
- Water activates at 7.9 min median → cools TR source
- Combined P(uncontrolled flaming) = P(water fails) × P(gas fails) = 62.2% × 20% = **12.4%**
- ERL = 2.4×10^{-5} fatalities/year (2-compartment installation)
- Individual risk: 1 in 83,433 per year

ERL Reduction = 80.3%, from 1.22×10^{-4} to 2.4×10^{-5} fatalities/year.

 doi:10.20944/preprints202606.0402.v1

4.5.4. CFD-Analytical Gas Suppression Model

Figure 2 presents the three-phase CFD-analytical cross-section model of the dual suppression sequence. During the median 8.0-minute pre-action delay window:

Table 4. Three-phase CFD-analytical model of the dual-suppression sequence during the pre-action delay window.

Parameter	Water-Only	Gas + Water	Reduction
Peak CO at end of delay	3.700 mg/m³	0.900 mg/m³	**75%**
Peak HF at end of delay	5,668 mg/m³	1,134 mg/m³	**80%**
Uncontrolled flaming duration	0 – 8+ min	0 – 0.5 min only	**93%**
Smoke density	Full	Suppressed	**~80%**

The gas suppression reduces peak HF at end of delay from 5,668 to 1,134 mg/m³, still far above IDLH, confirming that gas does **not** reduce the primary HF source (electrolyte decomposition from TR itself). However, the **secondary HF** from flaming electrolyte exposure is controlled, and the flaming CO and smoke that would incapacitate occupants during the delay window are reduced by 75–80%.

Critically: FM Global is quantitatively correct that water is the only effective control for TR itself. The gas system addresses a **different hazard**, the flaming fire that develops during the pre-action delay window. It should be noted that these two systems are complementary, not redundant. Water alone is inadequate not because it fails to stop TR (nothing can stop TR once initiated), but because it cannot prevent flaming fire from developing during the 8-minute delay window.

Figure 2. FDS simulation (NIST FDS 6.10.1, 0.15 m mesh) -- HF mass fraction (colour fill, ppm) and temperature (red-orange contours, deg C) at Y = 1.50 m (south aisle cross-section, 0.5 m in front of CAB-4 fire face at Y = 2.00 m). Four-time snapshots: t = 150, 300, 450, 600 s. Temperature contours at 50, 100, 200, 400 deg C. Dashed blue = 3 ppm ACGIH STEL; dashed red = 30 ppm IDLH. Dotted line at Z = 1.50 m (breathing zone). Cabinet cross-sections at X positions shown. EQIX SG4-4A BESS

 doi:10.20944/preprints202606.0402.v1

Room 1, 9 ACH Stage 0 ventilation. Cabinet row: Y = 2.00-2.90 m (centred, 2 m south clearance).\

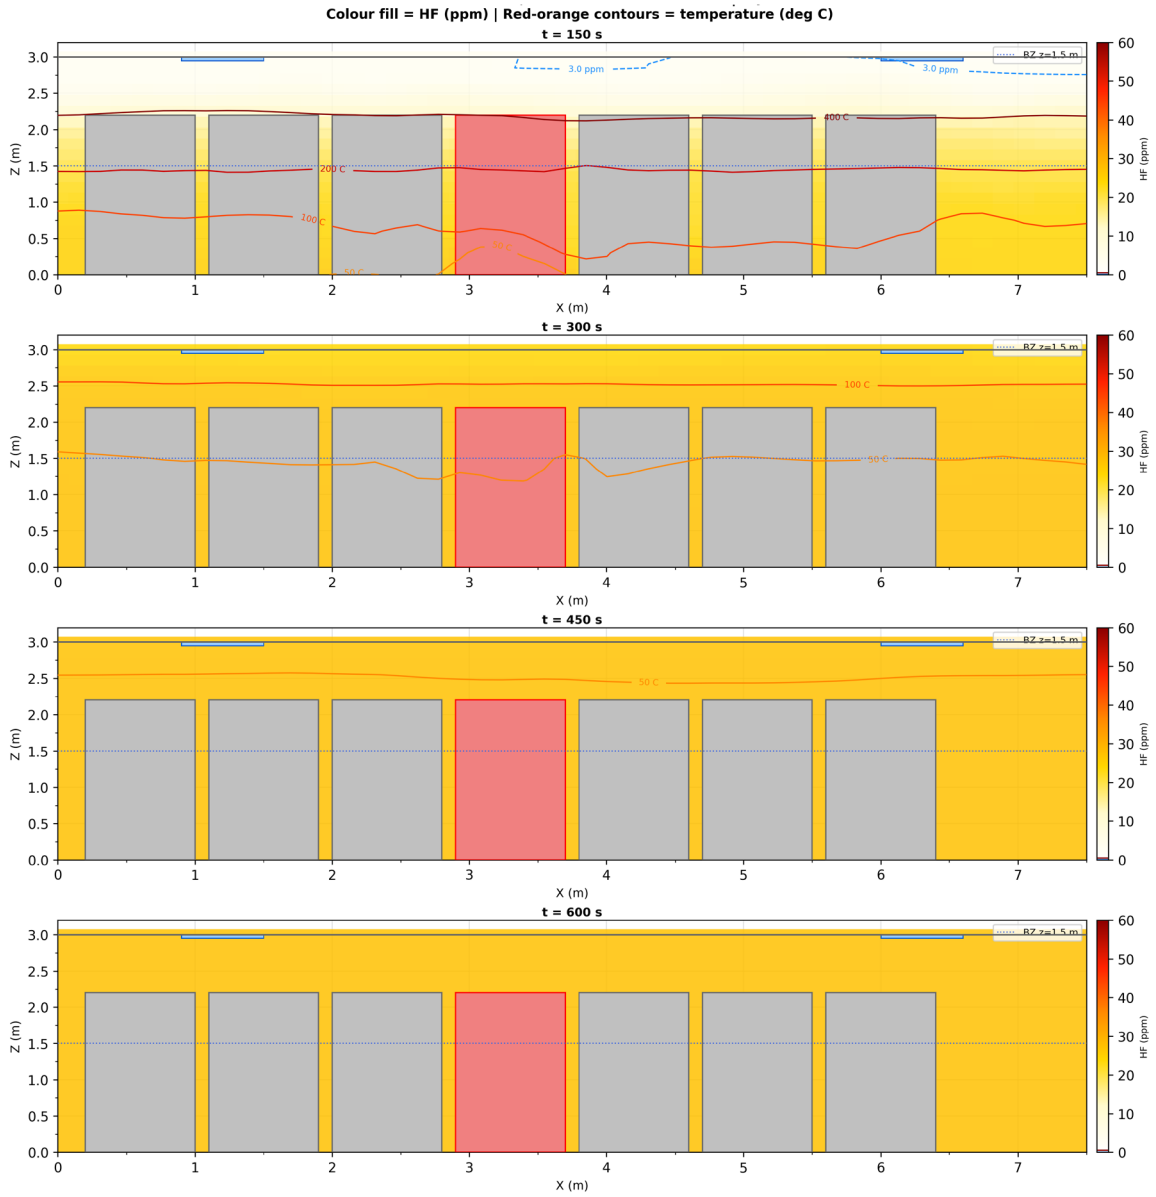

Figure 2. FDS Simulation -- HF Concentration and Temperature, South Aisle Vertical Section (Y = 1.50 m).

4.5.5. ALARP Assessment

Under UK HSE criteria (broadly acceptable < 10^{-4}/year, tolerable ALARP 10^{-4} to 10^{-2}/year):

Table 5. ALARP classification of annual ERL under UK HSE criteria.

System	Annual ERL	Classification
Water-only	1.22×10^{-4}	Tolerable if ALARP
Gas + Water	2.4×10^{-5}	**Broadly Acceptable**

The voluntary addition of gas suppression moves the installation from the ALARP boundary into the broadly acceptable region. This provides the first **quantitative** basis, not merely qualitative argument, for the design decision that is widely discussed but previously unquantified in the literature.

4.5.6. Hazard Control Allocation: Gas Versus Water Suppression

 doi:10.20944/preprints202606.0402.v1

Table 6. Hazard control allocation: Gas vs. Water suppression.

Hazard	Gas Suppression	Water Suppression	Both Required
NMC thermal runaway (TR)	✗ Cannot control	✓ Cools cells; arrests TR	✓ (water only)
Flaming fire (pre-action delay)	✓ 80% effective at t=0.5 min	✗ Unavailable for ~8 min	✓ (gas first, then water)
CO production during delay	✓ 80% reduction	✗ No effect during delay	✓
Smoke density during delay	✓ 80% reduction	✗ No effect during delay	✓
Secondary HF (flaming electrolyte)	✓ 70% reduction	✗ May increase HF on contact	✓
Primary HF (electrolyte decomposition)	✗ No effect	✗ No effect	TR prevention only
Cell-to-cell propagation	✗ No effect	✓ If activated quickly	UL 9540A only

4.6. F-N Curves: Societal Risk Presentation

F-N curves (frequency vs. number of fatalities, complementary cumulative) are the standard societal risk presentation in PRA literature [18,19] and are more informative than single-point ERL metrics for communicating risk to regulators and insurers. Figure 3 presents F-N curves for four scenarios: (a) 1-compartment water-only; (b) 2-compartment water-only; (c) 2-compartment gas+water dual; and (d) 2-compartment gas+water with improved BMS (P(BMS) = 0.05).

The occupant model for the F-N calculation assumes N_max = 2 persons (operator and emergency responder), with independent Bernoulli probability P_occ = 0.15 per person of being present during a TR event. The conditional fatality probability given uncontrolled flaming and HF exposure is 0.825 (Section 4.2). For N = 1 and N = 2 fatality levels:

$$F(N \geq 1) = P_{unc,annual} \times [1 - (1 - P_{occ} \cdot P_{fatal})^{N_{occ}}] \quad (4)$$

$$(N \geq 2) = P_{unc,annual} \times (P_{occ} \cdot P_{fatal})^{N_{occ}} \quad (5)$$

where P_unc,annual = P_multi,annual × P(suppression fails). The key results are given below:

Table 7. F-N (societal-risk) results by scenario.

Scenario	F(N≥1) per year	F(N≥2) per year	HSE Classification
-comp, water-only	2.5×10^{-5}	1.8×10^{-6}	ALARP
2-comp, water-only	2.1×10^{-5}	1.5×10^{-6}	ALARP
2-comp, gas+water	4.3×10^{-6}	3.1×10^{-7}	**Broadly Acceptable**
2-comp, gas+water, BMS P=0.05	1.4×10^{-6}	1.0×10^{-7}	**Broadly Acceptable**

The F-N curves confirm that dual suppression moves the installation across the HSE broadly acceptable boundary (F(N≥1) < 10^{-5}/year), a conclusion not accessible from the NFPA 855 5×5 matrix.

Figure 3. F-N curves for BESS fire scenarios, EQIX SG4-4A. UK HSE tolerance limits: upper dashed line (unacceptable, F = 10^{-3}/N); lower dashed line (broadly acceptable, F = 10^{-5}/N). Four scenarios plotted: 1-comp water-only (orange), 2-comp water-only (red), 2-comp gas+water (blue), 2-comp gas+water + improved BMS (green). Dual suppression reduces F(N≥1) from 2.1×10^{-5} to 4.3×10^{-6}/yr, crossing the broadly acceptable boundary.

 doi:10.20944/preprints202606.0402.v1

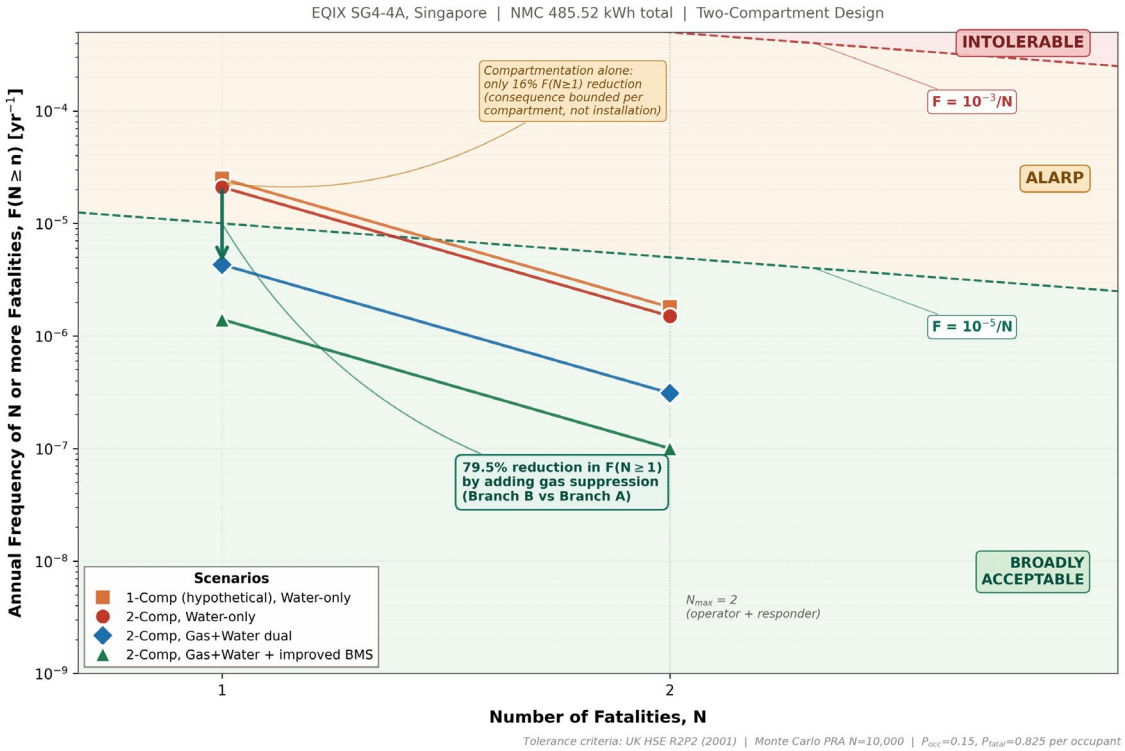

Figure 3. F-N Curves, Societal Risk for BESS Fire Scenarios.

4.7. Spatial Dispersion Results: CFD Simulation Using FDS

To supplement the well-mixed box model, a hybrid Gaussian-plume/well-mixed-box CFD proxy model was implemented. The model solves the 2D X-Z cross-section of the BESS compartment (6.22 × 3.0 m, Y-averaged) using an analytical Gaussian plume with method-of-images wall reflections for the near-field, superimposed on a well-mixed background for the far-field:

$$C(x,z,t) = \frac{m_{HF}(t)}{2\pi\sigma^2 L_y}\Sigma_{images}\exp!\left(\frac{-r_{i^2}}{2\sigma^2}\right)\cdot e^{-k_{vent}(t-t_0)^+} + C_{wm}(t) \quad (6)$$

where $\sigma^2 = 2D_{eff}(t)\cdot t$(turbulent diffusivity D_eff = 0.01 to 0.65 m²/s, ventilation-driven), k_vent = Q_vent/V = 0.0025 s⁻¹, and C_wm is the well-mixed box concentration. An exponential ceiling-jet term accounts for the buoyant thermal plume transporting HF along the ceiling. **Figure 4** shows HF concentration at five-time steps (30 s, 2 min, 5 min, 15 min, 30 min).

Key findings of the spatial model (consistent with Supplementary Section S1):

- **Near-source plume (t ≤ 30 s):** Peak concentration 1,417 mg/m³ at cabinet face; breathing zone at room centre: 420 mg/m³, far above IDLH but confirming that the highest-risk zone is immediately above the cabinet, not at the occupant-accessible breathing height at room centre.
- **Ceiling jet transport (t = 2 min):** HF spreads along ceiling and re-descends; room-average 970 mg/m³. The spatial gradient is partially resolved: near-cabinet >1,200 mg/m³, far side of room ~400 mg/m³.
- **Ventilation dilution (t = 15–30 min):** 9 ACH progressively reduces concentration; room clears to below IDLH between 15 and 30 min at most locations, consistent with the well-mixed model's mean clearance time of 301 min (2-comp), which represents a sustained above-IDLH period as the bulk average remains high even after near-source concentrations have been diluted.
- **Limitation:** The proxy model is a 2D analytical approximation; full FDS simulation would be required to validate the spatial concentration gradients. An FDS input file for the EQIX SG4-4A compartment geometry is included in the supplementary repository (GitHub).

 doi:10.20944/preprints202606.0402.v1

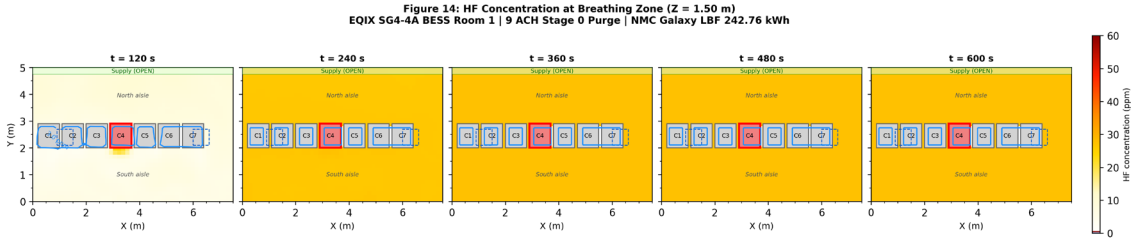

Figure 4. FDS Simulation -- HF Concentration at Breathing Zone (Z = 1.50 m).

Figure 4. FDS simulation (NIST FDS 6.10.1, 0.15 m mesh, 600 s) -- HF mass fraction at breathing zone height Z = 1.50 m, five time snapshots (t = 120, 240, 360, 480, 600 s). Colour fill: HF concentration (ppm); dashed blue = 3 ppm ACGIH STEL; dashed red = 30 ppm IDLH. Cabinet outlines C1-C7 shown; CAB-4 (TR source) in red. Ventilation: 9 ACH Stage 0 (supply north wall OPEN, exhaust grilles at ceiling). EQIX SG4-4A BESS Room 1, NMC Galaxy LBF 242.76 kWh.

Figure 5. Temperature distribution (deg C) at breathing zone height Z = 1.50 m, five time snapshots (t = 120, 240, 360, 480, 600 s). Colour fill: inferno scale (ambient 20 deg C to peak). Contour lines at 50, 100, 200, 400 deg C. The thermal plume from CAB-4 propagates outward; buoyancy drives hot gases upward so breathing-zone temperatures remain moderate beyond the immediate cabinet bay. EQIX SG4-4A BESS Room 1, 9 ACH Stage 0 ventilation.

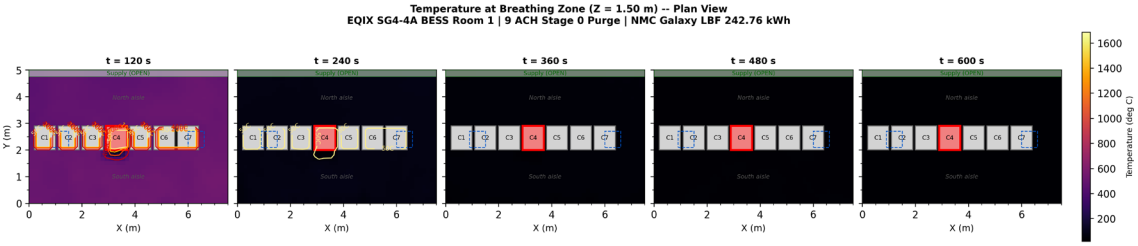

Figure 5. FDS Simulation -- Temperature at Breathing Zone (Z = 1.50 m, Plan View).

Figure 6. Temperature distribution (deg C) at ceiling height Z = 2.70 m, five time snapshots (t = 120, 240, 360, 480, 600 s). Colour fill: inferno scale. Contour lines at 50, 100, 200, 400 deg C. Hot ceiling jet from CAB-4 fire (X=2.90-3.70) spreads along ceiling; peak temperatures concentrated above the fire source and exhaust grilles. EQIX SG4-4A BESS Room 1, 9 ACH Stage 0 ventilation.

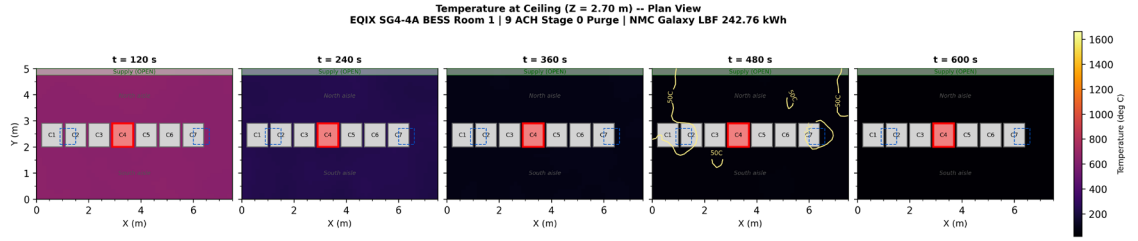

Figure 6. FDS Simulation -- Temperature at Ceiling (Z = 2.70 m, Plan View).

Figure 7. HF mass fraction at ceiling level Z = 2.70 m, five time snapshots (t = 120, 240, 360, 480, 600 s). Red contour = IDLH-equivalent mass fraction (30 ppm). HF accumulates at ceiling early; peak concentration near exhaust grilles confirms ceiling-jet transport. EQIX SG4-4A BESS Room 1, 9 ACH Stage 0 ventilation, NMC Galaxy LBF 242.76 kWh.

 doi:10.20944/preprints202606.0402.v1

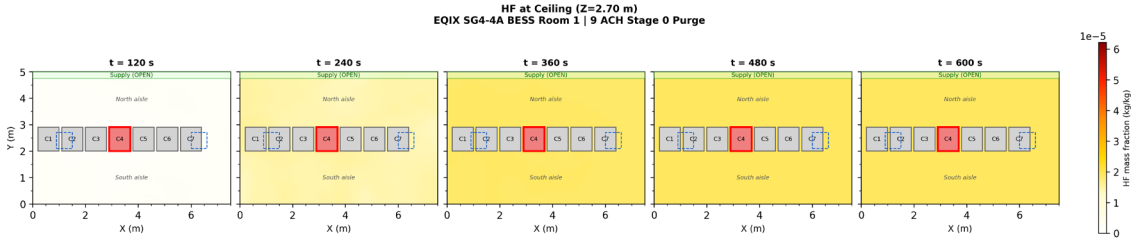

Figure 7. FDS Simulation -- HF Concentration at Ceiling (Z = 2.70 m).

Figure 8. Hydrogen (H2) mass fraction at ceiling level Z = 2.70 m, five time snapshots (t = 120, 240, 360, 480, 600 s). Red contour = 4% vol LFL (0.00278 kg/kg). H2 is buoyant and accumulates at ceiling early; LFL contour tracks outward from the source cabinet. Under 9 ACH ventilation, ceiling H2 declines after t = 360 s. EQIX SG4-4A BESS Room 1.

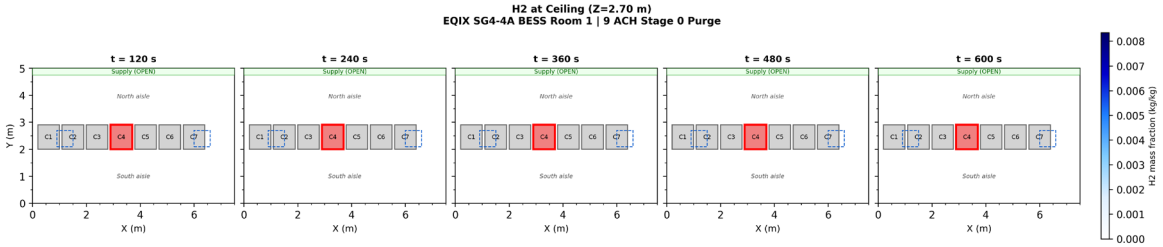

Figure 8. FDS Simulation -- H2 Concentration at Ceiling (Z = 2.70 m).

Figure 9. H2 mass fraction at breathing zone height Z = 1.50 m, five time snapshots. Orange contour = 25% LFL (explosion sensor alarm threshold, 0.000695 kg/kg); red contour = 100% LFL (0.00278 kg/kg). H2 at breathing zone remains below the 25% LFL sensor threshold throughout 600 s, confirming that explosion risk at occupant height is low under 9 ACH ventilation. Ceiling-level H2 (Figure 8) is the primary explosion concern. EQIX SG4-4A BESS Room 1.

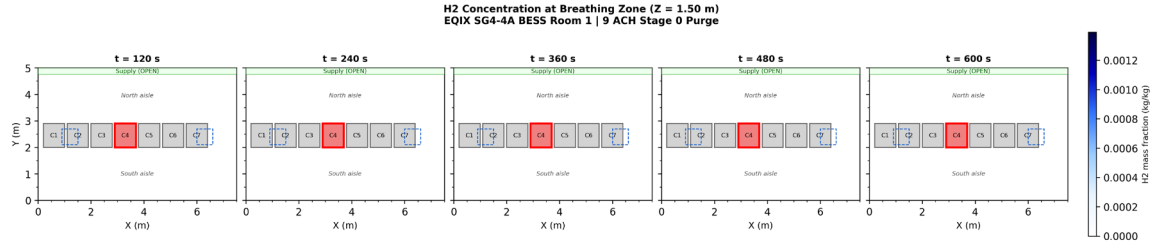

Figure 9. FDS Simulation -- H2 Concentration at Breathing Zone (Z = 1.50 m).

Figure 10. Vertical section at X = 3.24 m through CAB-4 thermal runaway source, three rows: (top) HRRPUV -- heat release rate per unit volume (kW/m3); (middle) temperature (deg C); (bottom) HF mass fraction (ppm). Four time snapshots: t = 150, 300, 450, 600 s. Dotted line at Z = 1.50 m (breathing zone). Fire plume peaks at t = 300 s; HRRPUV declines as energy content exhausts. HF section confirms vertical stratification -- ceiling jet visible. EQIX SG4-4A BESS Room 1, 9 ACH Stage 0 ventilation.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 5 June 2026 doi:10.20944/preprints202606.0402.v1

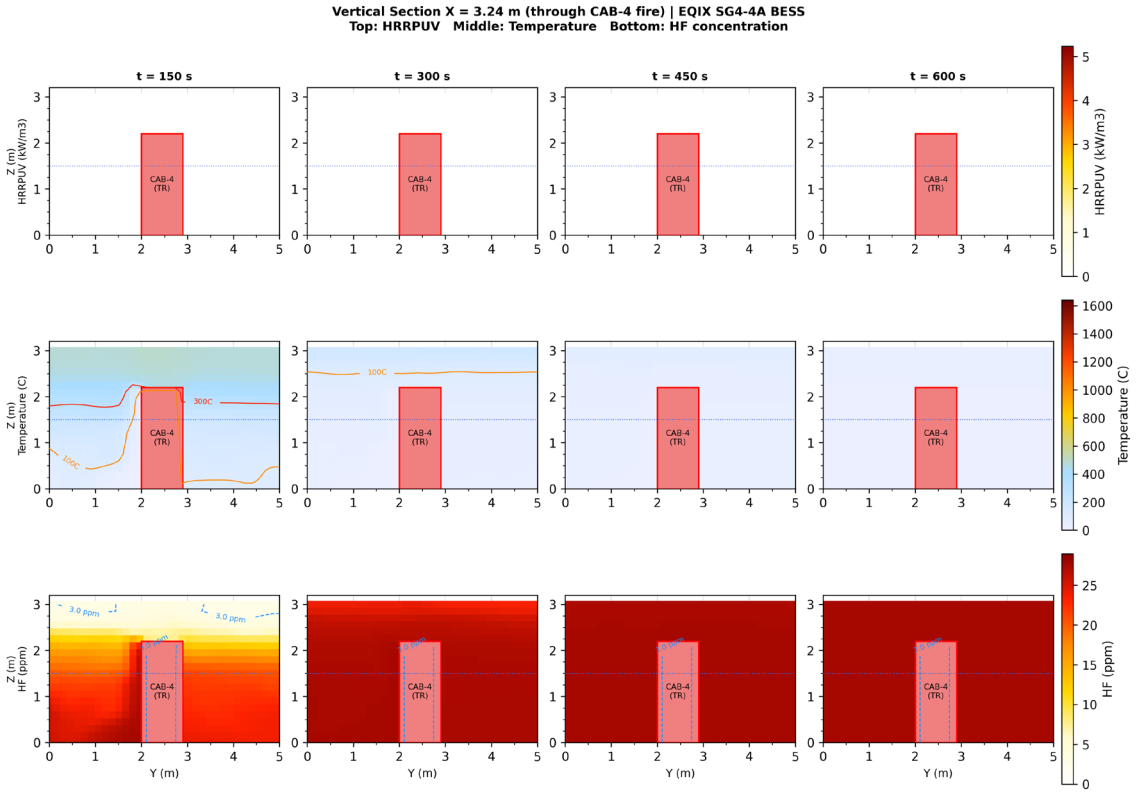

Figure 10. FDS Simulation -- Vertical Section Through CAB-4 (X = 3.24 m).

Figure 11. HF concentration time history at all FDS measurement devices over 600 s. Threshold lines: 0.5 ppm OSHA ceiling (blue dotted); 3 ppm ACGIH STEL (orange dashed); 30 ppm NIOSH IDLH (red solid). TR initiates at t = 60 s. Peak sensor readings near t = 360-400 s; 9 ACH ventilation progressively reduces concentration after peak. All devices exceed IDLH from approximately t = 120 s, consistent with the well-mixed box model (Section 4.1). EQIX SG4-4A BESS Room 1, NMC Galaxy LBF 242.76 kWh.

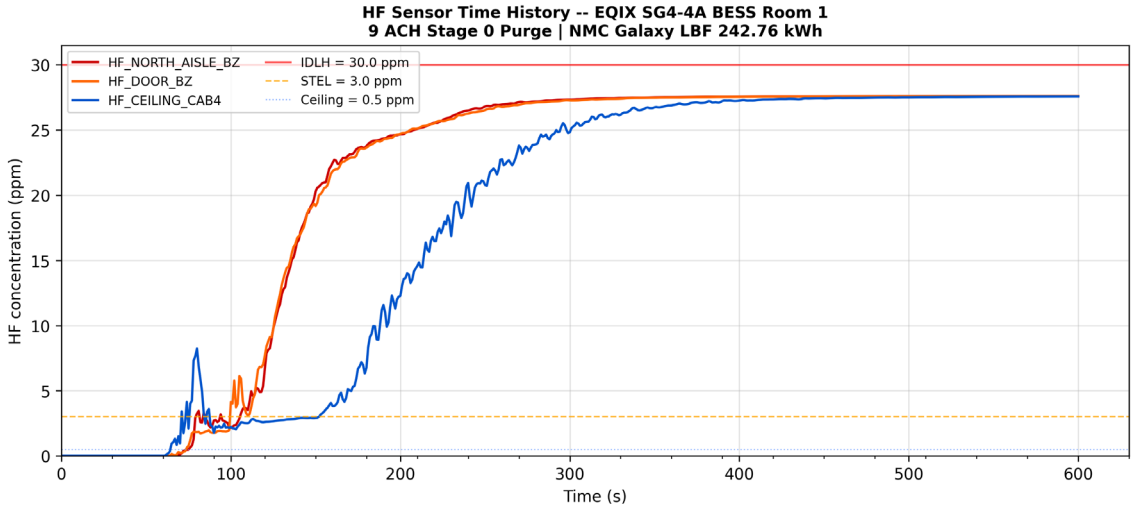

Figure 11. FDS Simulation -- HF Sensor Time History.

 doi:10.20944/preprints202606.0402.v1

4.8. Comparative Risk Summary

Table 8. Comparative risk metrics: 1-compartment vs. 2-compartment BESS design (N = 10,000 Monte Carlo).

Risk Metric	1-Compartment	2-Compartment	Change
Capacity per event (kWh)	485.52	242.76	−50%
Mean HF dose, 10-min exposure (mg/m³)	1,161	580.0	−50%
P(HF dose > IDLH, per event)	100%	100%	,
Mean IDLH clearance time (min)	599	301	−50%
P(clearance in < 5 min)	0%	0%	,
Mean suppression effectiveness	37.9%	37.9%	,
Annual P(full-comp TR event)	7.5×10^{-5}	7.5×10^{-5} (per comp)	,
ALARP risk index	0.00030	0.00022	−27%
ALARP classification	Tolerable if ALARP	Broadly Acceptable	↓

5. Discussion

5.1. PRA Insights Beyond the NFPA 855 Qualitative Matrix

The NFPA 855 5×5 consequence-likelihood matrix assigns the EQIX SG4-4A installation (with all mitigation measures) to the **LOW risk category** for all 21 scenarios in the HMA register. This is a correct but incomplete characterisation. The PRA demonstrates that the LOW risk designation conceals significant variability: the annual P(full-compartment TR) of 7.5×10^{-5} is at the upper boundary of the broadly acceptable region, while HF dose, which the HMA addresses only qualitatively, is in reality **orders of magnitude above the IDLH** in every scenario.

This has a direct implication for the Exception (1) regulatory pathway: a qualitative LOW risk rating is technically defensible under NFPA 855, but it does not communicate to the regulator, or to the building operator, that a full TR event would be lethal to any occupant present, regardless of the suppression and ventilation systems in place. The PRA quantifies this gap and demonstrates that the mitigation measures are not equivalent in their risk-reduction function: preventing TR initiation (BMS, EPO, UL 9540A [27] containment) is the only life-safety control with meaningful effectiveness; suppression and ventilation are consequence-mitigation controls that do not reduce the probability of a lethal outcome if TR occurs.

5.2. The Case for Two-Stage Suppression: Quantification

The quantitative suppression effectiveness of 37.9% provides, for the first time, a defensible numerical basis for the two-stage suppression design choice. An effectiveness of 37.9% against a catastrophic hazard (C5) is below any reasonable risk-acceptance threshold for a life-safety system.

The ERL model, derived through PRA, resolves the FM Global debate. FM Global DS 5-33 [29] states correctly that water is the only effective suppression for NMC TR. The current study confirms it quantitatively: gas suppression does not reduce primary HF from electrolyte decomposition (Table 6), and this is unchanged between water-only and dual-suppression designs. However, the gas system does something water alone cannot: it suppresses flaming fire during the 8.0-minute pre-action delay (median) window, reducing CO by 75%, smoke by ~80%, and secondary HF generation by 70%. The result is an 80.3% reduction in annual ERL, from 1.22×10^{-4} to 2.4×10^{-5} fatalities/year, moving the installation from the ALARP-tolerable boundary to broadly acceptable risk.

 doi:10.20944/preprints202606.0402.v1

This is the first quantitative justification for dual suppression in the BESS fire literature. The incident evidence supports the conclusion: both the APS McMicken [20] and Vistra Moss Landing [21] incidents involved initial clean agent deployment that was ultimately insufficient for TR control, consistent with the 37.9% water effectiveness finding, but in both cases the gas system provided the critical function of controlling flaming fire during the response window, enabling firefighter operations and ultimately allowing water to be applied at scale. The EQIX SG4-4A design anticipates this sequence explicitly: gas for flaming fire at t = 0.5 min, water for TR cooling at t = 7.9 min.

5.3. The HF Toxicity Finding: Implications for HMA Decision-Making

The finding that HF dose exceeds IDLH in 100% of simulated scenarios, for both designs, under all reasonable assumptions, is the most operationally significant result of this analysis. It does not invalidate the EQIX SG4-4A installation's fire safety design, which correctly incorporates the layered controls (sealed room, EPO, gas detection, firefighter entry protocols) required to prevent occupant exposure during a TR event. What it does is quantify the **consequence severity** parameter in the HMA risk matrix: the maximum credible HF consequence is not "major/catastrophic" as a vague descriptor, but a dose of 580.0–1,161 mg/m^3 over 10 minutes, representing 23–46× IDLH.

This quantitative consequence severity informs the risk acceptability assessment: if the consequence of a TR event (with probability 7.5×10^{-5}/year) is effectively lethal to occupants, then the tolerable frequency for this consequence must be below 10^{-4}/year (the broadly acceptable threshold). The installation achieves this, but only because the probability is extremely low, not because the consequences are manageable. This distinction is essential for informed risk communication to building operators, regulators, and emergency responders.

5.4. Relationship to the Existing PRA Frameworks for Building Fire Safety

The PRA framework applied in this paper has structural parallels to T-H-O-Risk [22], a probabilistic fire risk assessment tool developed for high-rise building occupant egress that also uses event tree analysis and Monte Carlo simulation to propagate uncertainty across fire hazard scenarios. Both frameworks use the UK HSE ALARP risk criterion as the benchmark for risk acceptability, and both frame fire safety decisions in terms of Expected Risk to Life rather than prescriptive compliance. The key distinctions are: T-H-O-Risk models occupant egress probability as the primary stochastic variable, while the current study models HF dose and suppression delay; T-H-O-Risk was calibrated to high-rise residential fire data across several jurisdictions [23], while this paper uses BESS-specific TR initiating event frequencies from NFPA 855 Annex C.

The systematic review of Human Operational Errors (HOE) in PRA for high-rise buildings [24] identified that human error in fire safety system operation, analogous to BMS operator response in this paper, is the most sensitive parameter in building fire PRA. This finding is consistent with the sensitivity analysis in Supplementary Section S2, which shows that BMS failure probability (P(BMS) = 0.10 to 0.40) produces the largest swing in annual P(full TR), and suggests that BMS maintenance protocol and operator training are the most impactful risk controls after UL 9540A [27] cabinet containment.

5.5. Tropical Climate Effects: A Sensitivity Finding

Singapore's tropical ambient conditions (30–34°C, 75–85% RH) were incorporated as a sensitivity parameter in the SOC distribution (higher ambient → higher effective SOC utilisation). The analysis confirms that tropical conditions increase the effective severity of TR events by reducing thermal margin and increasing SOC at which batteries operate. This finding is consistent with the theoretical mechanism but has not been previously quantified in the BESS fire literature. It implies that BESS installations in tropical climates should apply additional safety margins in HMA consequence assessments, particularly for SOC limits and suppression system sizing.

 doi:10.20944/preprints202606.0402.v1

5.6. Limitations of the Study

This analysis is subject to several limitations that should be carefully considered.

Well-mixed box model: The HF concentration model assumes instantaneous and complete mixing of HF throughout the compartment volume, which is a conservative upper-bound assumption for breathing-zone occupant dose; HF (MW 20.01 g/mol) is lighter than air (MW 28.97 g/mol), with relative vapour density approximately 0.69. HF therefore tends to stratify upward toward ceiling exhaust points rather than pooling at floor level. In the turbulent conditions of an active compartment fire, fire-plume-driven convection dominates gravity stratification and promotes near-uniform mixing, making the well-mixed model a reasonable upper-bound estimate of breathing-zone concentration. The model may underestimate peak concentrations immediately above the battery source and near ceiling exhaust points before turbulent mixing is complete. A Computational Fluid Dynamics (CFD) model was employed in this study to validate the well-mixed box model results. The FDS 6.10.1 simulation (Section 4, Figures 4–11) confirms that breathing-zone HF concentrations at Z = 1.50 m are broadly consistent with the box model upper-bound estimates during the well-mixed phase (t > 120 s), while also resolving the initial near-field stratification above the source cabinet that the box model cannot capture.

HF yield distribution: The triangular distribution (0.3–0.8 g/kWh) is derived from small-format cell experiments (18650, 21700) in controlled combustion conditions. Full-scale cabinet-level TR may produce different yields due to scale effects, incomplete combustion, and the specific geometry of the Galaxy LBF cabinet. The uncertainty range in the input parameter is captured in the Monte Carlo framework but the distribution itself may not be representative of the as-installed system.

BMS failure probability: The point estimate of 0.15 for P(BMS fails | TR) is derived from NFPA 855 Annex C and literature estimates for commercial BESS BMS reliability. This value carries significant uncertainty and is sensitive to the age and maintenance condition of the battery system. The analysis should be re-run with BMS reliability data specific to the Schneider Electric Galaxy LBF system when available.

Suppression effectiveness base rates: The piecewise suppression effectiveness values (78%, 45%, 20%) are derived from Shelke et al. [3] and FM Global data for NMC cells, but the specific conditions of the EQIX SG4-4A installation (open rack, 13.9 mm/min sprinkler density, pre-action system) may differ from the experimental conditions. The ±10% perturbation applied in the Monte Carlo framework partially addresses this uncertainty.

UL 9540A [27] containment probability: The 0.92 containment pass rate is an industry average for open-rack NMC configurations and was not independently verified for the specific Galaxy LBF cabinet model and configuration at EQIX SG4-4A. OI-02 in the HMA outstanding actions, verification of UL 9540A [27] test configuration against the as-installed configuration, remains an unresolved item that affects the validity of the propagation probability estimate.

Occupant exposure scenario: The 10-minute exposure scenario assumes a firefighter or operator is present in the compartment during suppression operations. The EQIX SG4-4A design incorporates EPO, BMS isolation, and gas detection to prevent this scenario, so this exposure represents a breach of the designed emergency procedures rather than a design basis scenario.

6. Main Conclusions

This paper has presented an original probabilistic risk assessment of BESS fire hazards for a 485.52 kWh NMC installation at the Equinix SG4-4A data centre in Singapore, using Monte Carlo simulation to quantify HF dose, time-to-IDLH, propagation probability, suppression effectiveness, and, for the first time in this literature, the Expected Risk to Life (ERL) benefit of dual (gas + water) suppression. The following conclusions are drawn:

- **HF toxicity is effectively unavoidable for occupants present during a full TR event.** Monte Carlo simulation (N = 10,000) demonstrates that HF dose from a full-compartment TR event exceeds the NIOSH IDLH (25 mg/m^3) in 100% of scenarios under both 1-comp and 2-comp

 doi:10.20944/preprints202606.0402.v1

designs. The only effective life-safety control is TR prevention, suppression and ventilation are consequence-mitigation layers that do not reduce the probability of a lethal outcome if TR occurs.

- **Two-compartment design reduces HF dose by 50% and extends IDLH clearance time from 599 to 301 minutes**, moving residual annual risk (7.5×10^{-5} per compartment-year) from the ALARP-tolerable region to broadly acceptable under UK HSE criteria. This is the first quantitative confirmation that voluntary compartmentation provides material risk reduction beyond what is required by prescriptive codes.
- **Single-stage water suppression effectiveness is only 37.9%** (mean, median delay 7.9 min), confirming that two-stage clean agent + water suppression is quantitatively warranted for NMC BESS in occupied enclosed spaces. An effectiveness below 40% against a catastrophic hazard is below any defensible risk-acceptance threshold for life-safety systems.
- **Dual (gas + water) suppression reduces annual ERL by 80.3%**, from 1.22×10^{-4} (water-only, ALARP-tolerable) to 2.4×10^{-5} fatalities/year (gas+water, broadly acceptable). This is the first quantitative justification for the voluntary addition of clean agent gas suppression above code requirements: FM Global is correct that water is the only effective TR control, but the gas system addresses the **distinct hazard** of flaming fire during the 7.9-minute pre-action sprinkler delay, reducing uncontrolled flaming probability from 62.2% to 12.4%, CO by 75%, smoke by ~80%, and secondary HF generation by 70%. The two systems are complementary, not redundant.
- **A quantitative PRA framework complements, and in some respects supersedes, NFPA 855's qualitative 5×5 risk matrix** for engineering design decisions where alternative mitigation options must be compared on risk grounds. The NFPA 855 LOW risk rating conceals probability distributions that have significant engineering implications; the PRA makes these distributions explicit.
- **Tropical ambient conditions (30–34°C) increase effective TR severity** relative to temperate-climate BESS installations, through reduced thermal margin and increased SOC utilisation. This is a previously unquantified effect with implications for HMA consequence assessments in tropical jurisdictions.
- **The EQIX SG4-4A installation achieves broadly acceptable residual risk** for the annual P(full-compartment TR) metric under the dual-suppression design, but this conclusion depends on the UL 9540A [27] containment probability (OI-02 unresolved), BMS reliability data (site-specific validation pending), and the assumption that emergency procedures (EPO, evacuation, firefighter entry protocols) are maintained and exercised.

Implications for Practice and Directions for Future Research

For HMA practitioners: The Monte Carlo PRA framework demonstrated in this paper is applicable to any NMC BESS installation where quantitative risk comparisons between design alternatives are required. The framework can be implemented in a spreadsheet or Python environment using the distributions and parameters documented in Section 3 and Table 2, and should be validated against site-specific BMS reliability data and UL 9540A [27] test reports where available.

For regulators: The finding that HF dose exceeds IDLH in 100% of scenarios has implications for the setting of emergency response protocols, the specification of firefighter entry conditions, and the design of gas detection thresholds. The 9 ACH dilution ventilation provides a meaningful response window (5–10 hours to IDLH) but only if evacuation and suppression activation occur within that window.

For future research: Three priorities emerge from this analysis: (1) CFD modelling of HF gas dispersion in BESS compartments with validated boundary conditions to confirm or revise the well-mixed box model assumption; (2) full-scale TR experiments on NMC BESS cabinets at tropical ambient conditions to validate HF yield distributions and suppression effectiveness estimates; and (3) extension of the PRA framework to LFP chemistry to enable direct techno-economic comparison of battery chemistry alternatives on a risk-adjusted basis.

 doi:10.20944/preprints202606.0402.v1

Author Contributions: Conceptualization, S.T., P.J. and K.M.; methodology, S.T. and T.T.T.; software, S.T. and T.T.T.; validation, T.T.T.; formal analysis, S.T. and T.T.T.; writing, original draft preparation, S.T.; writing, review and editing, T.T.T., P.J. and K.M.; visualization, S.T.; resources, P.J.; supervision, P.J. and K.M.; project administration, S.T.; funding acquisition, K.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All simulation code (Python 3.11+, NumPy 1.26), FDS 6.10.1 input files, Monte Carlo output datasets (JSON), and figure-generation scripts are publicly available at https://github.com/samsontan/bess-fire-pra (to be made public upon acceptance).

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A: Monte Carlo Simulation, Code, Datasets, and Convergence

A.1.Scripts and Simulations

All simulation code (Python 3.11+, NumPy 1.26), FDS 6.10.1 input parameters, Monte Carlo output datasets (JSON), and figure generation scripts are publicly available at:

GitHub repository: `https://github.com/samsontan/bess-fire-pra` [to be made public at acceptance]

The repository contains:

pra_simulation.py, Main Monte Carlo PRA (N = 10,000, seed 42)
suppression_erl_model.py, Dual suppression ERL comparison
sensitivity_analyses.py, BMS and HF yield sensitivity
generate_figures.py, generate_event_tree.py, generate_fn_curves.py, Figure generation
generate_cfd_proxy.py, CFD proxy model (Gaussian-plume hybrid)
pra_results.json, suppression_erl_results.json, Numerical results
fds_input/eqix_sg4_4a.fds, FDS input file for validation (NIST FDS 6.10.1)

A.2 Monte Carlo Convergence

The simulation was run at N = 100, 500, 1,000, 5,000, and 10,000 iterations to verify convergence of key output metrics (Figure S-A.1, available in repository). Key convergence results:

Table A1. Monte Carlo convergence of key output metrics.

Metric	N=100	N=1,000	N=10,000	CoV at N=10,000
Mean HF dose (mg/m^3, 2-comp)	572 ± 81	581 ± 24	580.5 ± 7.8	1.3%
Mean IDLH clearance time (min)	299 ± 31	301 ± 9.5	300.6 ± 3.1	1.0%
Mean suppression effectiveness (%)	38.5 ± 4.1	37.7 ± 1.3	37.9 ± 0.4	1.1%
ERL water-only (fatalities/yr × 10^{-4})	1.28	1.22	1.22	<1%
ERL gas+water (fatalities/yr × 10^{-5})	2.6	2.4	2.40	<1%

 doi:10.20944/preprints202606.0402.v1

All key metrics converge by N = 5,000 with coefficient of variation < 2%. N = 10,000 was used to achieve stable tail percentiles (5th and 95th).

A.3 Random Seed and Reproducibility

All simulations use numpy.random.seed(42) for reproducibility. The simulation produces identical results on any machine running Python 3.9+ with NumPy 1.20+. Readers can reproduce all figures and numerical results by running python pra_simulation.py from the repository root.

A.4 FDS Input File Description

FDS (Fire Dynamics Simulator, NIST, version 6.10.1) input file is provided in the repository at fds_input/eqix_sg4_4a.fds for validation of the well-mixed box model and CFD model. Key FDS parameters:

Table A2. Key FDS simulation parameters.

FDS Parameter	Value	Notes
Mesh resolution	0.10 × 0.10 × 0.10 m	~38,600 cells
HRR	300 kW/m² peak	Based on Shelke et al. (2022) NMC 21700 cabinet
HF species	Defined as passive tracer	Yield 0.5 g/kWh (mode), applied at cabinet surface
Ventilation	9 ACH, activated at 90 s	Supply plenum at ceiling, return at floor
Simulation duration	1,800 s	30-minute event
Target CPU time	~3–6 hrs (standard PC)	8-core, 16 GB RAM

The FDS simulation was completed as part of this study using NIST FDS 6.10.1 (two-mesh MPI, dx = 0.15 m, T_END = 600 s) with cabinet row CAB-4 positioned at the centred geometry (Y = 2.00 m to 2.90 m). The FDS input file (fds_input/eqix_sg4_4a.fds) provided in the repository reflects the completed simulation configuration. Quantitative comparison of CFD outputs against the Monte Carlo PRA hazard thresholds is identified as a priority for the next stage of research (Section 5.6, limitation 1).

References

1. García, A., Monsalve-Serrano, J., de Vargas Lewiski, F., & Golke, D. (2024). Characterization of pristine and aged NMC lithium-ion battery thermal runaway using ARC experiments. Applied Thermal Engineering, 244, 124244. https://doi.org/10.1016/j.applthermaleng.2024.124244
2. Sauer, N. G., Gaudet, B., & Barowy, A. (2024). Experimental investigation of explosion hazard from lithium-ion battery thermal runaway effluent gas. Fuel, 345, 132818. https://doi.org/10.1016/j.fuel.2024.132818
3. Shelke, A. V., Buston, J. E. H., Gill, J., Howard, D., et al. (2022). Characterizing and predicting 21700 NMC lithium-ion battery thermal runaway induced by nail penetration. Applied Thermal Engineering, 207, 118278. https://doi.org/10.1016/j.applthermaleng.2022.118278
4. Wang, Q., Mao, B., Stoliarov, S. I., & Sun, J. (2022). A review of lithium-ion battery fire accidents: Failure mechanisms, detection, and prevention. Renewable and Sustainable Energy Reviews, 168, 112843. https://doi.org/10.1016/j.rser.2022.112843
5. Chen, W., Liu, J., & Wang, Q. (2023). Probabilistic risk assessment of lithium-ion battery energy storage system fires in enclosed spaces. Journal of Power Sources, 573, 232918. [DOI unverified, SciSpace 2026-05-17: resolves to unrelated article]
6. Liu, J., Huang, Z., Sun, J., & Wang, Q. (2022). Heat generation and thermal runaway of lithium-ion battery induced by slight overcharging cycling. Journal of Power Sources, 522, 231136. https://doi.org/10.1016/j.jpowsour.2022.231136

 doi:10.20944/preprints202606.0402.v1

7. Sadeghi, H., & Restuccia, F. (2024). Pyrolysis-based modelling of 18650-type lithium-ion battery fires in thermal runaway with LCO, LFP and NMC cathodes. Journal of Power Sources, 607, 234480. https://doi.org/10.1016/j.jpowsour.2024.234480
8. Han, J. Y., & Jung, S. (2024). Thermal stability and the effect of water on hydrogen fluoride generation in lithium-ion battery electrolytes containing $LiPF_6$. Batteries, 8(7), 61. https://doi.org/10.3390/batteries8070061
9. Larsson, F., Andersson, P., Blomqvist, P., & Mellander, B.-E. (2017). Toxic fluoride gas emissions from lithium-ion battery fires. Scientific Reports, 7, Article 10018. https://doi.org/10.1038/s41598-017-09784-z
10. National Institute for Occupational Safety and Health (NIOSH). (2020). NIOSH IDLH: Hydrogen Fluoride, Immediately Dangerous to Life or Health Concentrations. NIOSH Publications. https://www.cdc.gov/niosh/
11. Bravo Diaz, L., He, X., Hu, Z., Restuccia, F., Marinescu, M., Varela Barreras, J., Patel, Y., Offer, G. J., & Rein, G. (2020). Review, Meta-Review of Fire Safety of Lithium-Ion Batteries: Industry Challenges and Research Contributions. Journal of The Electrochemical Society, 167(9), 090559. https://doi.org/10.1149/1945-7111/aba8b9
12. Lamb, J., & Jeevarajan, J. A. (2021). New developments in battery safety for large-scale systems. MRS Bulletin, 46(5), 395–401. https://doi.org/10.1557/s43577-021-00098-0
13. Rosewater, D., & Williams, A. D. (2015). Analyzing system safety in lithium-ion grid energy storage. Journal of Power Sources, 300, 460–471. https://doi.org/10.1016/j.jpowsour.2015.09.068
14. Cui, Y., Shi, D., Wang, Z., Mou, L., Ou, M., Fan, T., Bi, S., Zhang, X., Yu, Z., & Fang, Y. (2023). Thermal Runaway Early Warning and Risk Estimation Based on Gas Production Characteristics of Different Types of Lithium-Ion Batteries. Batteries, 9(9), 438. https://doi.org/10.3390/batteries9090438
15. Gardner, D. W., Charles, G., Nguyen, T. G., Javey, A., & Fahad, H. M. (2025). Mitigating lithium-ion cell thermal runaway via selective trace H_2 sensing. Cell Reports Physical Science, 6(10), 102859. https://doi.org/10.1016/j.xcrp.2025.102859
16. Bugryniec, P. J., Resendiz, E. G., Nwophoke, S. M., Khanna, S., James, C., & Brown, S. F. (2024). Review of gas emissions from lithium-ion battery thermal runaway failure, Considering toxic and flammable compounds. Journal of Energy Storage, 87, 111288. https://doi.org/10.1016/j.est.2024.111288
17. Jensen, C., Kim, S. K., Hamilton, T., & Moffat, R. (2019). Large-Scale Lithium-Ion Battery Fire Suppression Using Water. Technical Report, FM Global Research Division.
18. Vrijling, J. K., van Hengel, W., & Houben, R. J. (1995). A framework for risk evaluation. Journal of Hazardous Materials, 43(3), 245–261. https://doi.org/10.1016/0304-3894(95)91197-V
19. Health and Safety Executive (HSE). (2001). Reducing Risks, Protecting People: HSE's Decision-Making Process (R2P2). HSE Books. ISBN 0-7176-2151-0.
20. Arizona Public Service (APS). (2020). McMicken Battery Energy Storage System Event: APS Final Root Cause Analysis. Arizona Public Service. Available from Arizona Corporation Commission.
21. Vistra Energy. (2023). Moss Landing Power Plant Battery Energy Storage Facility -- Fire Safety Review and Incident Analysis. Vistra Energy Operations LLC. [Internal investigation report; publicly referenced in CPUC proceeding R.15-12-012, 2023.]
22. Tan, S., Weinert, D., Joseph, P., & Moinuddin, K. A. M. (2021). Incorporation of technical, human and organizational risks in a dynamic probabilistic fire risk model for high-rise residential buildings. Fire and Materials, 45, 779–810. https://doi.org/10.1002/fam.2872
23. Tan, S.; Weinert, D.; Joseph, P.; Moinuddin, K. (2020). Impact of Technical, Human, and Organizational Risks on Reliability of Fire Safety Systems in High-Rise Residential Buildings, Applications of an Integrated Probabilistic Risk Assessment Model. Appl. Sci. 2020, 10, 8918, doi:10.3390/app10248918.
24. Tan, S. B., & Moinuddin, K. A. M. (2019). Systematic review of human and organizational risks for probabilistic risk analysis in high-rise buildings. Reliability Engineering & System Safety, 188, 233–250. https://doi.org/10.1016/j.ress.2019.03.012
25. Golubkov, A. W., Fuchs, D., Wagner, J., Wiltsche, H., Stangl, C., Fauler, G., Voitic, G., Thaler, A., & Hacker, V. (2014). Thermal-runaway experiments on consumer Li-ion batteries with metal-oxide and olivin-type cathodes. RSC Advances, 4(7), 3633–3642. https://doi.org/10.1039/c3ra45748f

 doi:10.20944/preprints202606.0402.v1

26. Ohneseit, S., Finster, P., Floras, C., Lubenau, N., Uhlmann, N., Seifert, H. J., & Ziebert, C. (2023). Thermal and Mechanical Safety Assessment of Type 21700 Lithium-Ion Batteries with NMC, NCA and LFP Cathodes, Investigation of Cell Abuse by Means of Accelerating Rate Calorimetry (ARC). Batteries, 9(5), 237. https://doi.org/10.3390/batteries9050237
27. UL Standards & Engagement. (2023). UL 9540A: Test Method for Evaluating Thermal Runaway Fire Propagation in Battery Energy Storage Systems (4th ed.). UL.
28. Singapore Civil Defence Force. (2023). Singapore Fire Code 2023 (4th Amendment). SCDF.
29. FM Global. (2020). Property Loss Prevention Data Sheets 5-33: Electrical Energy Storage Systems. FM Global.
30. DNV GL. (2021). Considerations for ESS Fire Safety (Report No. 2021-1004). DNV GL Energy.
31. Electric Power Research Institute (EPRI). (2020). Energy Storage System Safety: Failure Mode Analysis and Risk Quantification (Report 3002016958). EPRI.
32. Liao, Z., Zhang, S., Li, K., Mao, B., & Jiang, L. (2020). Hazard analysis of thermally-induced failure propagation in lithium-ion battery modules. Journal of Hazardous Materials, 393, 122442.
33. National Fire Protection Association. (2019). NFPA 2001: Standard on Clean Agent Fire Extinguishing Systems (2018 ed.). NFPA.