Intelligent Data Tiering Design for Optimizing Data Availability and System Load in Nuclear Activity Detection Platforms

Sueyeon Lee, Gayeon Ha, Yonhong Jeong*
Center for Nuclear Nonproliferation Strategy & Technology, Korea Institute of Nuclear Nonproliferation and Control (KINAC), 1418 Yuseong-daero, Yuseong-gu, Daejeon 34101, Republic of Korea
Corresponding author: jyh1404@kinac.re.kr

Keywords: Data Tiering, Information Lifecycle Management (ILM), Nuclear Security, Big Data Engineering, OSINT

1. Introduction

The surveillance and early detection of nuclear activities in neighboring countries represent a critical challenge for regional and international security. Recent advances in commercial satellite technology and the growth of multilingual Open Source Intelligence (OSINT) data have enabled high-precision surveillance. However, these advancements have simultaneously created severe information overload challenges for next-generation nuclear activity detection systems. Current systems employ uniform storage strategies, treating all collected data equally regardless of analytical value or temporal relevance. This causes computational bottlenecks, degraded access latency, and inefficient resource allocation.

This study proposes a dynamic 3-tier data management framework that determines data lifecycle and value in real time. By optimizing processing paths through intelligent tiering, we aim to minimize system load and maximize analytical efficiency.

2. Case Study and Methodology

2.1 Case Study

Case A: Cyber Threat Intelligence (CTI) Systems and Data Lifecycle Management

CTI systems face challenges analogous to nuclear activity monitoring: managing massive volumes of heterogeneous data while identifying critical threats in real time. Sun et al. (2023) surveyed CTI mining frameworks, revealing a progressive refinement pipeline where raw threat indicators flow through collection, processing, analysis, and dissemination stages [1]. Critically, each stage exhibits distinct performance requirements—immediate collection demands millisecond-scale access, validation requires batch processing capacity, and archived intelligence serves long-term analysis.

Wagner et al. (2019) documented that organizations processing 10,000+ threat indicators daily maintain differentiated data pathways [2]. High-priority indicators of compromise (IOCs) receive sub-second processing in hot storage, while contextual threat data undergoes batch validation before archival. This staged approach enables rapid response to critical threats while efficiently managing the full data lifecycle.

Connection to nuclear monitoring: This progression directly maps to nuclear activity detection, where urgent satellite anomalies require immediate analysis (Tier 1), validation against proliferation indicators occurs through automated and expert review (Tier 2), and confirmed intelligence feeds long-term trend analysis (Tier 3).

However, CTI systems primarily handle structured cyber indicators (IOCs, TTPs), whereas nuclear activity detection requires processing heterogeneous, multimodal data including satellite imagery and geospatial signals. Direct methodological transfer is therefore limited without domain-specific adaptation.

Case B: Adaptive Storage Management for AI Workloads

Pang et al. (2023) proposed Adaptive Intelligent Tiering (AIT), which uses deep learning to predict data access patterns and reinforcement learning to optimize data placement across storage tiers [3]. Their experiments with 3-tier storage systems demonstrated performance improvements of up to 85% compared to traditional frequency-based migration policies such as Least Recently Used (LRU) or Least Frequently Used (LFU).

The key innovation in AIT is real-time adaptability to changing workload patterns. Unlike legacy Hierarchical Storage Management (HSM) systems that rely solely on file age for tier demotion, AIT continuously learns which data elements will be needed next and proactively migrates them to appropriate performance tiers. This predictive capability is particularly valuable in environments where data value depends on context rather than simple access frequency.

Connection to nuclear monitoring: In nuclear activity detection, data value varies dramatically based on context. For example, thermal imagery of a facility captured months ago is low-priority during routine monitoring but becomes immediately critical when suspicious procurement patterns are detected. AIT's context-aware tiering strategy provides a framework for responding to such dynamic value changes.

Nevertheless, AIT optimizes data placement based on access-frequency patterns derived from enterprise IT workloads. In nuclear monitoring, data value is

governed less by frequency and more by geopolitical context and PIR alignment, necessitating a domain-specific scoring function beyond standard access-pattern learning.

2.2 Proposed Tiering Methodology

The system employs a 3-tier hierarchical architecture to manage data based on analytical urgency and verified value:

Tier 1 – Transient Layer (Hot): Processes real-time indicators including thermal anomalies, OSINT keyword triggers, and steam/cooling water discharge. NVMe SSDs storage provides sub-millisecond latency for immediate GPU-accelerated LLM inference.

Tier 2 – Validation Layer (Warm): Validation layer for pattern matching, geospatial contextualization, and expert review. SSD/HDD hybrid storage provides second-scale latency; data routing by Bayesian confidence scores (S).

Tier 3 – Persistent Layer (Cold): Archives verified intelligence including facility profiles and proliferation networks. HDD arrays provide minute-scale retrieval for long-term trend analysis and AI training.

Table 1. Three-Tier Data Management Architecture

Attribute	Tier 1	Tier 2	Tier 3
Retention	7 days	30–90 days	Indefinite
Storage	NVMe SSD	SSD/HDD Hybrid	HDD Array
Latency	Sub-millisecond	Second-scale	Minute-scale
Score (S)	Pre-scoring (input)	$0.30 \leq S < 0.75$	$S \geq 0.75 \rightarrow$ Tier 3; $S < 0.30 \rightarrow$ Purge
Data Volume	5–10%	Transitional	85–90%
Key Content	Satellite anomalies, OSINT triggers, thermal imagery	Pattern matching, geospatial context, expert review	Facility profiles, proliferation networks

Migration Logic and Decision Thresholds: Data placement is dynamically managed by a multi-factor scoring function:

$$S(d) = W_1 \cdot f_{temporal}(d) + W_2 \cdot f_{confidence}(d) + W_3 \cdot f_{access}(d) + W_4 \cdot f_{PIR}(d)$$

where w_1, w_2, w_3, w_4 are weights optimized via reinforcement learning to adapt to the threat landscape, and $f_{temporal}$, $f_{confidence}$, f_{access} and f_{PIR} denote temporal decay, Bayesian confidence, access frequency, and Priority Intelligence Requirement (PIR) correlation scores, respectively.

To ensure objective and automated tiering, each component of the scoring function S(d) is normalized to the range [0, 1] as follows:

(1) Temporal Decay:

$$f_{temporal}(d) = e^{-\lambda(t_{now} - t_{created})}$$

where λ is a domain-tuned constant that prioritizes recent anomalies.

(2) Confidence Score:

$$f_{confidence}(d) = \frac{1}{n}\sum_{i=1}^{n} P_i(v)$$

Derived from Bayesian posterior probability in the Validation Layer (Tier 2), updated iteratively via expert review.

(3) Access Frequency:

$$f_{access}(d) = \frac{log(1 + freq)}{log(1 + freq_{max})}$$

Log-normalization prevents excessive weight concentration on frequently accessed datasets.

(4) PIR Correlation:

$$f_{PIR}(d) = \frac{\mathbf{V}_d \cdot \mathbf{V}_{PIR}}{\|\mathbf{V}_d\| \, \|\mathbf{V}_{PIR}\|}$$

Cosine similarity between data metadata vectors and predefined PIR vectors.

The weight vector $W = \{w_1, w_2, w_3, w_4\}$ is dynamically optimized using a Deep Q-Network (DQN) agent. The system state (S) includes current tier occupancy and processing latency, while the reward function (R) is defined as:

$$R = \alpha \cdot Throughput - \beta \cdot StorageCost + \gamma \cdot PIRHitRate$$

This incentivizes the agent to maximize analytical throughput and PIR-related detection accuracy while minimizing operational costs.

Crucially, the migration between tiers is governed by two thresholds, τ (Promotion) and ρ (Purge), aligned with Intelligence Community (IC) analytic standards:

$\tau = 0.75$, set at the median of the 'High Confidence' range (0.70–0.80), promotes data to Tier 3;

$\rho = 0.30$, set at the lower bound of the 'Low Confidence' tier, triggers purging to mitigate system overload.

Data within [0.30, 0.75) remains in Tier 2 for continued expert evaluation.

Fig. 1. Proposed intelligent data tiering architecture for nuclear activity detection platforms

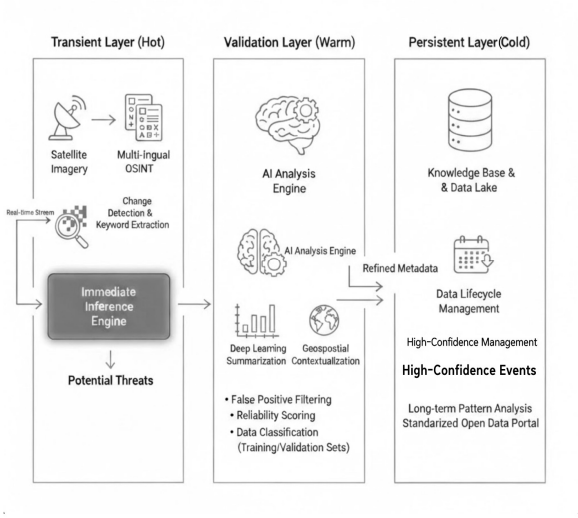

3. Expected Results

Based on empirical benchmarks from analogous CTI frameworks and AIT systems, the proposed architecture is projected to achieve the following performance gains:

- Processing Latency: A 60–85% reduction for high-priority analytical workflows by prioritizing Transient Layer data for immediate inference.
- Storage Costs: A 40–50% reduction by strategically maintaining only 5–10% of total data in Hot Storage.
- Analytical Throughput: A 100x increase (500–50,000 items/day) enabled by the automated Validation Layer pipeline.
- False Positives: A 70–80% reduction through multi-stage Bayesian validation and expert-aligned scoring.
- Scalability: The framework demonstrates linear scalability, as the Persistent Layer (85–90% of total volume) utilizes cost-effective archival technologies. The RL-based migration policy ensures adaptive optimization without manual reconfiguration.

4. Conclusions

This study proposes an intelligent data tiering framework designed to address the critical challenge of information overload in next-generation nuclear activity detection platforms. By dynamically managing the data lifecycle based on temporal relevance, Bayesian confidence, and intelligence priorities, the proposed architecture enables a more balanced distribution of computational resources.

The framework is expected to yield significant operational improvements by alleviating the bottlenecks associated with uniform storage strategies. Preliminary analysis suggests that this approach can substantially reduce access latency and storage costs while enhancing system throughput through prioritized processing paths. Furthermore, the system establishes a data-centric foundation for advanced AI-driven detection, where validated intelligence serves as high-quality training sets and real-time streams enable proactive threat identification.

Future work will focus on: (1) empirical validation through pilot deployment in operational environments to quantify performance gains, (2) refinement of scoring functions and thresholds based on real-world workload analysis, (3) integration of federated learning for secure multi-agency data sharing, and (4) development of Explainable AI (XAI) methods for automated tiering decisions to ensure regulatory and intelligence transparency.

REFERENCES

[1] N. Sun, M. Ding, J. Jiang, W. Xu, X. Mo, Y. Tai, and J. Zhang, Cyber Threat Intelligence Mining for Proactive Cybersecurity Defense: A Survey and New Perspectives, IEEE Communications Surveys & Tutorials, Vol.25, No.3, pp.1748-1774, 2023.

[2] T. D. Wagner, K. Mahbub, E. Palomar, and A. E. Abdallah, Cyber threat intelligence sharing: Survey and research directions, Computers & Security, Vol.87, 2019, Art. no. 101589.

[3] L. Pang, A. Alazzawe, M. Ray, K. Kant, and J. Swift, Adaptive intelligent tiering for modern storage systems, Performance Evaluation, Vol.160, 2023, Art. no. 102332.