

# TEE-Protected Drone Inspection for Critical Infrastructure: Securing Edge Analytics and Trust at Scale

Yongzhi Wang

California State Polytechnic University Pomona, Pomona, CA 91768, USA

## ABSTRACT

Operators of critical infrastructure are required to inspect distributed assets on a regular basis to ensure safety and reliability. Traditionally, inspections have been performed by field workers. While expert technicians are indispensable, manual patrols expose personnel to electrical, physical, and environmental hazards, which may constrain inspection frequency, coverage, and data quality. In recent years, many organizations have started to utilize drone technology to enhance safety, improve efficiency, and reduce operational costs. Unmanned aerial vehicles (UAVs) are a promising productivity amplifier: a pilot can dispatch a drone to nearby functional locations, collect data using various sensors, run pre-processing analytics to identify assets and assess their conditions, then send the data back for additional analysis. Yet the edge-centric workflow that makes drones efficient also expands the attack surface. A compromised UAV or malicious insider can exfiltrate or manipulate observations and processing results, undermining trust of the downstream analytics. This paper proposes a trusted execution environment (TEE)-based data protection framework for UAV-driven infrastructure inspection. It enforces end-to-end confidentiality and integrity across the entire lifecycle of the inspection data. We designed the framework and security protocols of this system and studied different attack scenario to understand the security of the proposed design. Our studies show that, under the threats from external attacks and malicious insiders, the framework enforces security for the full cycle of the inspection data, across acquisition, transport, and processing.

**Keywords:** UAV, Trusted execution environment, Edge computing, Cyber security, Critical infrastructure, Inspection

## INTRODUCTION

Operators of critical infrastructure are required to inspect distributed assets regularly to ensure safety and reliability. These assets span a wide range of domains, including bridges and roads, rail corridors, pipelines, power transmission lines, solar farms, wind turbines, telecom towers, ports, and large industrial facilities. Traditionally, inspections have been performed by field workers. While expert technicians remain indispensable, manual patrols expose personnel to electrical, physical, and environmental hazards, which often constrain inspection frequency, coverage, and data quality.

In recent years, many organizations have begun to leverage unmanned aerial vehicles (UAVs) to enhance safety, efficiency, and cost-effectiveness.

Drones provide a promising productivity amplifier: a UAV operator can dispatch a drone to one or more nearby functional sites, collect high-resolution RGB, thermal, or LiDAR data, apply analytics to identify assets and assess their conditions, then relay the data back for gatekeeping, desktop analysis, and master-data updates. For example, UAVs are now widely used by utilities to inspect power lines—replacing dangerous manual inspections or expensive helicopter surveys—while maintaining or improving data quality and coverage (Rymer & Moore, 2020; Mendu & Mbuli, 2025). In particular, studies note that drones equipped with visual, thermal or LiDAR sensors can detect damaged wires, overheating components, and structural weaknesses—enabling safer, faster, and more frequent inspections across broad networks (Wang, Gao, Xu & Li, 2021; Xing, Cioffi, Hidalgo-Carrió & Scaramuzza, 2023). Such inspection workflows deliver improved safety, reduced human exposure to electrical hazards, more frequent and data-rich inspection, and lower operational cost.

However, the features that make UAVs advantageous, remote operation, wireless data transmission, and onboard sensing, also expand the system's attack surface. UAV security research has documented a wide range of potential cyber-physical threats, including communication-link vulnerabilities, firmware or software exploitation, GPS spoofing, data interception, and sensor attacks (Mekdad et al., 2021; Ceviz, Sen & Sadioglu, 2023; Pratama et al., 2023). Beyond external attackers, there is a serious risk posed by malicious or negligent insiders, individuals such as drone pilots or inspectors, or field workers, who have legitimate access to UAV-collected data. Such insiders could tamper with, delete, or misreport inspection data, exfiltrate sensitive imagery or geolocation metadata, or misuse the drone for unauthorized surveillance or reconnaissance (Ceviz et al., 2023; Mekdad et al., 2021).

Because of these emerging threats, deploying UAV for inspection and surveillance of critical infrastructure should not be treated only as a technological upgrade, but also a shift to a new risk profile. In this paper, we analyzed the security threats inherent in UAV-based infrastructure inspection and introduced a TEE-based data protection framework that safeguards its entire data lifecycle, from collection, to transmission, to final analysis. **To the best of our knowledge, this is the first work to identify and address risks arising not only from external adversaries but also from malicious or negligent insiders operating in the field.** Through detailed attack-scenario case studies, we demonstrated that the proposed framework provides comprehensive protection for data confidentiality, integrity, and provenance across the full inspection workflow.

## SECURITY REQUIREMENTS AND THREATS ANALYSIS

### Security Requirements

When organizations deploy UAVs for infrastructure inspection (e.g., power-line, bridge, pipeline, or large industrial-site inspections), the inspection data typically flows through a life cycle of data collection, data transmission, and data processing (or analytics). For infrastructure inspection, the ultimate

security goal is to securely collect the data from the UAV and transferred to the data processing systems without jeopardizing data **confidentiality** and **integrity**. Since data processing is performed in an enterprise environment, which typically are protected through traditional security measures, the risk appears in the phase of data collection and data transmission. It requires strict confidentiality to prevent unauthorized access to sensitive data and asset information, as well as strong integrity to ensure the data remains accurate and unaltered from collection to analysis. If either property is compromised at any stage, during data collection, transmission, or ground-station processing, the resulting data may be misleading or exposed, undermining trust in inspection results and potentially jeopardizing critical infrastructure.

### Work Flow of UAV-Driven Infrastructure Inspection

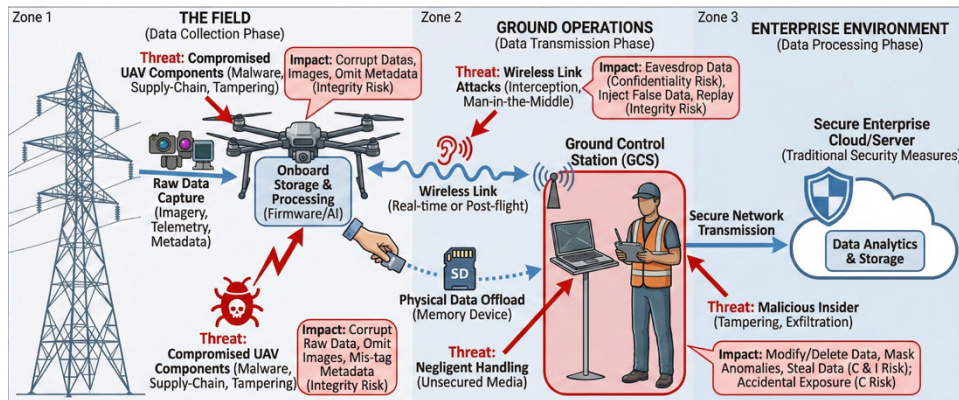
The work flow of the UAV-driven infrastructure inspection can be found in Fig. 1. Data collection and transmission is two critical phases of the workflow. We elaborate each phase below.

- **Data collection:** A drone (UAV) equipped with sensors, including high-resolution cameras (RGB), thermal imagers, LiDAR or other sensors, flies over or along the infrastructure. A trained field worker or UAV pilot controls (or supervises) the flight, performs pre-flight planning, initiates the flight, and manages the data capture. During flight, the UAV records imagery, sensor data, telemetry (GPS, IMU), metadata (time, location, flight parameters), possibly video, thermal readings, 3D scans, etc. This is the “raw data” stage. Some UAV has advanced computing or AI function that can pre-process data so that to reduce the size of the stored data or transmit the pre-processed data to the ground in real-time. For example, an autonomous drone for inspection can perform AI inference based on the sensor data and send the data to the ground only when the data needs additional processing (e.g., detection of a defect on an electric cable).
- **Data transmission:** After or during the flight, the collected data must be transferred from the UAV (onboard storage, memory, or data link) to a ground control station (GCS), or uploaded to storage for later analysis. This may involve wireless links (radio, Wi-Fi, proprietary drone-to-ground communication), or physical data offload via memory cards / storage devices. It also includes the procedure of transmitting data securely from ground control station to the enterprise environment for future processing.

### Security Threats

The critical parties involved in the inspection are the UAV and the field worker (or pilot). The component on the UAV performing data collection and transmission include sensors, flight-control systems, hardware, firmware, and software performing the data collection, storage and processing. If any component on the UAV was compromised, e.g., via malware, supply-chain attack, hardware tampering, the UAV might record incorrect or misleading data, including omitting images of damaged components, capturing corrupted

sensor outputs, mis-tagging GPS metadata, failing to record critical sections, etc. An attacker could also disable logging, corrupt raw data, or cause the UAV to crash, destroying both the UAV and its data without trace. The wireless link used to transmit or offload data from UAV to ground station (or to a network) can be intercepted, eavesdropped, or manipulated if not properly secured. Without proper encryption and authentication, an adversary might intercept raw or pre-processing data. They may inject false data, replay previous data, or corrupt data in transit, breaking confidentiality and integrity. The field worker (or pilot) is responsible for planning and executing flight, collecting data, and operating data transmission and offload. Because the field worker (or pilot) controls the UAV, initiates flights, and often handles data offload, they have legitimate privileges. A malicious insider could intentionally tamper with or delete collected data (e.g., omit images showing infrastructure defects), or modify metadata (e.g., geo-location or timestamps) to mask anomalies. They might even exfiltrate raw imagery / sensor logs / metadata to unauthorized parties (competitors, adversaries, foreign actors). In addition, through negligence or poor handling (unsecured storage, weak encryption, unsecured physical devices), they might inadvertently expose sensitive inspection data. Because insiders operate within authorized domains, standard perimeter defences may not detect or prevent such misuse.



**Figure 1:** Security threats in critical infrastructure inspection<sup>1</sup>.

## THREAT MODEL

In our threat model, we explicitly define the trust boundaries and adversarial capabilities relevant to UAV-based infrastructure inspection. We assume that the enterprise environment, including backend servers, key-management services, and the remote-attestation service, is fully trusted, physically protected, and correctly implements standard security controls. In contrast, we assume that all components operating in the field, the UAV, the field

<sup>1</sup> The figure is generated with the assistance of Gemini 3.

worker (or pilot), and the ground control station (GCS), may be untrusted or partially compromised. The only trusted computing component on the UAV is the hardware-rooted Trusted Execution Environment (TEE), which provides isolated enclaves and supports remote attestation. Everything outside the enclave, including flight-control software, sensors, storage media, firmware, and communication modules, is considered vulnerable to adversarial manipulation.

## PRELIMINARIES

Trusted Execution Environments (TEEs) provide hardware-enforced isolation that allows sensitive code and data to run securely even when the surrounding operating system or applications may be compromised. At the core of a TEE is an enclave—a protected memory region that prevents unauthorized access or tampering by external software, including privileged system processes. TEEs also support remote attestation, a mechanism that enables a remote verifier to confirm that the correct code is running inside a genuine enclave before provisioning secrets or sensitive workloads. Modern TEEs are increasingly deployed on edge and IoT devices to protect on-device analytics, machine-learning inference, and sensor processing. For example, Keystone is an open-source RISC-V-based TEE framework that enables customizable enclave implementations tailored to resource-constrained edge platforms (Lee et al., 2020). Research continues to explore advanced TEE designs, such as multi-enclave architectures, to support secure distributed processing across isolated components. Wang and Habib (2025) implemented a multi-enclave split-inference framework that protects data confidentiality for on-device machine learning, demonstrating how TEEs can safeguard sensitive data and computation on the edge setting, which is ideal for UAV-driven infrastructure inspection.

## TEE-BASED DATA PROTECTION FRAMEWORK

To offer full cycle data protection for UAV-based infrastructure inspection, we propose a TEE-based data protection framework. We isolate data collection and pre-processing inside the hardware rooted enclave on the UAV, enforce data integrity and UAV flight behaviour through autoencoder-based anomaly detection, enforce data confidentiality and integrity through encryption and digital signatures. The security of the three components were all supported by TEE, delivering full cycle provable data security. The system design consists of four pillars.

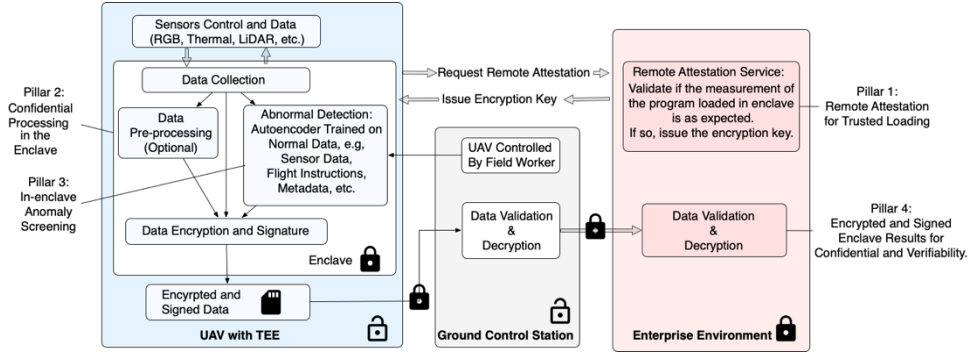
The architecture can be found in Fig. 2.

- **Pillar 1: Remote attestation for trusted loading:** Before a flight, the *remote attestation service* in the enterprise environment performs remote attestation against the UAV, validating the measurement (cryptographic hash) of the program loaded in the enclave. If the measurement matches with the expected value, the program loaded in the enclave is proven to be untampered. Then, the remote attestation service established an

encrypted channel with the enclave, and negotiated a *mission key*,  $K_M$  to the enclave, which is a symmetric encryption key used for encrypting collected or pre-processed data. The enclave also generates a pair of *signing key* ( $K_{pri}$ ,  $K_{pub}$ ), which is a pair of asymmetric key for signing its released data. It releases the  $K_{pub}$  to the enterprise environment when the remote attestation is passed. The remote attestation protected the integrity of the program loaded in the enclave. Any tampering of the program will be detected.

- Pillar 2: Confidential processing in the enclave: The security feature of TEE ensures that the computations in the enclave is protected from eavesdropping and tampering from any parties outside of the enclave. The computation includes, sensor data collection, feature extraction and data pre-processing, as well as the data encryption and signature signing. We introduced an abnormal detection module to the enclave, which can detect manipulations outside of the enclave, such as flight plan change from malicious field worker (or pilot), sensor data tampering from external attackers, etc. This module is also deployed inside the enclave, making it free from tampering and eavesdropping, therefore protecting its correctness.
- Pillar 3: In enclave anomaly detection: We deployed an autoencoder based anomaly detector inside the enclave. The model is expected to detect various adversarial behaviours, such as sensor data tampering from attackers, flight plan deviation by the malicious field worker (or pilot), etc. An autoencoder is an unsupervised neural network model. During training on normal, high-quality data, the encoder learns the essential patterns and structures needed for accurate reconstruction, while the decoder learns to reproduce these patterns from the latent representation. This makes autoencoders particularly useful for anomaly detection: when the model encounters data that deviates from the normal patterns it has learned, such as corrupted, tampered, or unusual sensor readings, it typically struggles to reconstruct it accurately, resulting in a higher reconstruction error. By measuring this error, systems can flag inputs that likely contain anomalies. The model can be trained with the data of previous flight history, including sensor and image data, flight plan (expected route), flight trajectory (actual route), and other metadata. When adversaries tampered the sensor data or changed the flight route, the autoencoder will generate a high reconstruction error, indicating a potential attack. Flags and reconstruction scores will be written to the flight output data and send it to ground control station and enterprise environment for further analysis.
- Pillar 4: Enclave encrypted and signed results for confidential and verifiable provenance: Processed outputs will be encrypted through the missing key  $K_M$  and will be signed with enclave's private key  $K_{pri}$ , an asymmetric key bound to the attested enclave instance. The enterprise environment will be able to validate the authenticity of the result using enclave's public key  $K_{pub}$ . Part of the result can also be validated by the ground control station. For example, during the remote attestation,

the enclave and the remote attestation service can negotiate fine-grained encryption keys, so that some keys can be shared to the ground control station to be used for data validation and decryption. This design ensures results transmitted intact and confidential and establishing a non-repudiable chain of custody from edge to the enterprise environment.



**Figure 2:** TEE-based data protection framework for UAV-driven infrastructure inspection.

## SECURITY ANALYSIS

In this section, we analyze several representative adversarial scenarios and explain how the proposed TEE-based data protection framework prevents, detects, or mitigates each attack. Broadly, the threats fall into two categories: **external adversaries** targeting the UAV or its communication channels, and **malicious insiders**, such as a compromised or untrusted field worker (or pilot). Across all scenarios, the four pillars of our design collectively ensure confidentiality, integrity, and verifiable provenance throughout the data lifecycle.

### External Attacks

**Eavesdropping or tampering with sensor data.** An external attacker may attempt to intercept sensor data during collection or inject manipulated readings into the UAV’s sensor stream. In our system, all sensor data flows directly into the attested enclave, whose memory is cryptographically isolated by hardware. Since the autoencoder-based anomaly detector also executes entirely inside the enclave, any externally induced perturbation, such as tampered imagery, spoofed telemetry, or modified flight-state metadata, will produce a higher reconstruction error, flagging the anomalous input. Because the attacker cannot break into the enclave or modify its execution, it is difficult for him to forge sensor data that both bypasses the anomaly detector and remains consistent with historical flight patterns. As a result, our framework prevents undetected tampering and ensures the integrity of raw sensor inputs.

**Eavesdropping or tampering with inspection results in storage or during communication.** Attackers may target data stored on the UAV (e.g., in local storage) or intercept results transmitted over wireless links between UAV and GCS, or between GCS and the enterprise backend. Pillar 1 guarantees

authenticity through remote attestation, while Pillar 4 ensures that all processed outputs are encrypted using the mission key  $K_M$  and digitally signed using the enclave's private key  $K_{pri}$ . Because the private signing key never leaves the enclave, adversaries cannot forge a valid signature or alter results without detection. Encryption protects confidentiality even if the wireless channel is fully compromised, and tampering with ciphertext will be immediately detected during signature verification at the GCS or enterprise environment. Thus, neither passive eavesdropping nor active modification of the inspection results can succeed.

### Malicious Insiders

**Leaking flight results.** A malicious field worker (or pilot) may try to extract raw sensor data or processed results. However, all sensitive data inside the UAV is encrypted using the mission key  $K_M$ , which is provisioned only to the enclave after successful remote attestation. The field worker (or pilot) does not possess  $K_M$  or the enclave's private key, and therefore cannot decrypt or meaningfully leak raw data or results. Even if they exfiltrate encrypted files, the data remains unintelligible without the corresponding keys held by the enterprise backend.

**Tampering with inspection results.** Because results are signed with the enclave's private key, any modification by the field worker (or pilot), whether by editing files, injecting forged outputs, or corrupting stored results, will be detected during signature verification. The adversary cannot construct valid falsified outputs without access to  $K_{pri}$ . Additionally, Pillar 3's anomaly detection identifies unexpected flight behaviour or sensor inconsistencies that a malicious worker may attempt to induce.

**Deviating from the flight plan.** A malicious operator might attempt to steer the UAV away from its authorized route to hide critical infrastructure defects or avoid areas of interest. The in-enclave autoencoder continuously learns and validates expected flight trajectories, comparing real-time telemetry with historical and mission-specific patterns. Deviations from the planned route trigger high reconstruction errors, resulting in anomaly flags embedded into the signed output. Since the anomaly detector operates inside the enclave and cannot be bypassed or tampered with, the field worker (or pilot) cannot conceal route deviations. These flags are transmitted to the enterprise environment alongside the encrypted and signed results, enabling downstream verification and investigation.

### CONCLUSION

In this paper, we examined the unique threat landscape of UAV-driven inspection and conducted a detailed security analysis of the end-to-end data workflow. To address the identified risks, we proposed a TEE-based data protection framework that protects the full data lifecycle. Our analysis demonstrates that the framework effectively mitigates attacks even when field-side components and communication channels are untrusted, ensuring reliable and tamper-resistant inspection results.

## REFERENCES

- Ceviz, O., Sen, S., & Sadioglu, P. (2023). A survey of security in UAVs and FANETs: Issues, threats, analysis of attacks, and solutions. *arXiv*. <https://doi.org/10.48550/arXiv.2306.14281>
- Dukowitz, Z. (2025, August 13). Powerline inspection drones: An in-depth guide. UAVCoach.
- Harty, P. (2025, August 28). The role of autonomous drones in long-range transmission line inspections. *POWER Magazine*.
- Lee, D., Kohlbrenner, D., Shinde, S., Asanović, K., & Song, D. (2020, April). Keystone: An open framework for architecting trusted execution environments. In *Proceedings of the Fifteenth European Conference on Computer Systems* (pp. 1–16).
- Mekdad, Y., Aris, A., Babun, L., El Fergougui, A., Conti, M., & Lazzeretti, R. (2021). A survey on security and privacy issues of UAVs. *arXiv*. <https://doi.org/10.48550/arXiv.2109.14442>
- Mendu, B., & Mbuli, N. (2025). State-of-the-art review on the application of unmanned aerial vehicles (UAVs) in power line inspections: Current innovations, trends, and future prospects. *Drones*, 9(4), 265. <https://doi.org/10.3390/drones9040265>
- NASA. (2020). A review of unmanned aerial vehicle technology in power line inspection. NASA Technical Reports Server.
- Niyonsaba, S., Konate, K., & Soidridine, M. M. (2023). A survey on cybersecurity in unmanned aerial vehicles: Cyberattacks, defense techniques and future research directions. *IJCNA*.
- Omolara, A. E., Alawida, M., & Abiodun, O. I. (2023). Drone cybersecurity issues, solutions, trend insights and future perspectives: A survey. *Neural Computing and Applications*, 35, 23063–23101. <https://doi.org/10.1007/s00521-023-08857-7>
- Pratama, D., Moon, J., Laksmono, A. M. A., Yun, D., Muhammad, I., Jeong, B., & Ji, J. (2023). Behind the wings: The case of reverse engineering and drone hijacking in DJI Enhanced Wi-Fi protocol. *arXiv*. <https://arxiv.org/abs/2309.05913>
- Rymer, N., & Moore, A. J. (2020). A review of unmanned aerial vehicle technology in power line inspection. NASA Technical Reports Server.
- The Drone Life. (2025). The best drones for power line inspections 2025. *The Drone Life*.
- Wang, Y., & Habib, A. (2025, May). Protect Data Confidentiality for on-Device Machine Learning Through Split Inference. In *2025 10th International Conference on Fog and Mobile Edge Computing (FMEC)* (pp. 290–297). IEEE.
- Wang, Z., Gao, Q., Xu, J., & Li, D. (2021). A review of UAV power line inspection. In *Advances in Guidance, Navigation and Control (Lecture Notes in Electrical Engineering, Vol. 644)*, 3147–3159. Springer.
- Xing, J., Cioffi, G., Hidalgo-Carrió, J., & Scaramuzza, D. (2023). Autonomous power line inspection with drones via perception-aware MPC. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. *arXiv*. <https://arxiv.org/abs/2304.00959>