Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 13 October 2026
D. Fairaizl
Independent Researcher
11 April 2026

A Standard for the Training and Inference of Machine Learning Models via Avian Parameter Carriers

draft-fairaizl-avian-ml-parameters-00

Abstract

Current machine learning infrastructure relies heavily on high-bandwidth digital interconnects for parameter synchronization, gradient aggregation, and model weight distribution. This document proposes an alternative: Avian Parameter Carriers (APC), building on the physical transport layer established in RFC 1149 and the quality-of-service extensions of RFC 2549.

The authors note that the bandwidth-delay product of a pigeon carrying a 512GB NVMe drive over 5 kilometers remains competitive with certain cloud providers during peak billing hours.

This specification is considered Feature Complete. It addresses scale (Section 14), observability (Appendix A), security (Sections 11 and 11.4), regulatory compliance (Section 9.3), and pigeon hygiene (Sections 3.3, 10.4, and 11.4.5). Future revisions MAY address quantum parameter transport (Section 14.6) and UV-spectrum adversarial plumage design (Section 11.4.2). The pigeons are considered stable.

Status of This Memo

Copyright Notice

Table of Contents

1. Overview and Rationale   3
2. Definitions and Terminology   4
3. Architecture   5
3.1. Training Loop   5
3.2. Batch Size and Payload Constraints   6
3.3. Loss Monitoring and the Loss Log   7
4. Bias-Variance Considerations   8
4.1. Variance   8
4.2. Explainability   8
4.3. Chain-of-Flight vs. Explainability   9
5. Regularization Techniques   9
5.1. Weight Decay   9
5.2. Early Stopping   9
5.3. Dropout and Carrier Shuffling 10
5.4. Data Augmentation 10
6. Distributed Training 10
6.1. Parameter Server Architecture 10
6.2. Federated Learning 11
6.3. Bandwidth-Delay Product 11
7. Transfer Learning 11
8. Quality of Service 11
9. Data Privacy and Gradient Leakage 12
9.1. Regulatory Compliance 12
9.2. Gradient Inversion Attacks 12
9.3. Data Subject Rights and the Predatory Interception Protocol 13
10. Electrostatic Discharge (ESD) Considerations 14
10.1. Threat Model 14
10.2. Mitigation Requirements 14
10.3. Scroll Integrity Verification 15
10.4. Shock-Based Training Instability (SBTI) 16
11. Security Considerations 17
11.1. General Threat Model 17
11.2. Messenger-in-the-Middle (MitM) Attacks 17
11.3. Prompt Injection via Avian Mimicry 18

11.4. Carrier Visual Obfuscation and Adversarial Plumage Patterning 19
11.4.1. Threat: Raptor-as-Classifier 19
11.4.2. Colorimetric Patterning Requirements 19
11.4.3. Steganographic Dual-Use Encoding 20
11.4.4. Null Gradient Signaling via Pattern Absence 21
11.4.5. Carrier Preparation and the Scrub-Before-Reuse Requirement 21
11.5. Security Considerations Summary 23
12. Ethical Considerations 24
12.1. Ethical Treatment of Gradient-Bearing Carriers 24
12.2. Bias in Carrier Selection 24
12.3. Informed Consent for Raptors 25
12.4. Environmental Considerations 25
13. Motivating Example: A Single Training Epoch 26
14. IANA Considerations 27
15. Scalability Considerations and Forward Compatibility 28
15.1. Model Size Scaling 28
15.2. Framework and Format Versioning 29
15.3. Topology Scaling: Beyond the Single Parameter Server . . 30
15.4. Inference Scaling and Serving 31
15.5. New Framework Onboarding 32
15.6. Quantum Computing Compatibility 32
16. Known Limitations 33
17. References 35
17.1. Normative References 35
17.2. Informative References 35
Acknowledgments 36
Smart Leg Band: Hardware Specification and Integration 36
Overview 36
Hardware Components 37
Table 1: Moisture Level Action Matrix 37
Data Ingestion and Telemetry 38
Optical Capture and Colorimetric Decode System 39
Firmware Security 40
Author’s Address 40

1. Overview and Rationale

The machine learning community has grown increasingly dependent on dense GPU clusters, high-speed interconnects, and cloud infrastructure with unpredictable pricing models. RFC 1149 [RFC1149] established that IP datagrams may be transmitted via avian carrier. This specification extends that foundation to accommodate the specific requirements of neural network training, including forward passes, backpropagation, gradient descent, and the existential uncertainty of whether the model is actually learning anything.

Avian Parameter Carriers offer several advantages over conventional infrastructure:

a. Carriers self-provision given adequate grain and water.

b. No software licensing fees. No CUDA version conflicts.

c. Gradient delivery failures are observable and attributable (see Section 4.2).

d. Carriers naturally implement a form of dropout regularization through attrition (see Section 5.3).

2. Definitions and Terminology

The following definitions apply throughout this document.

Carrier (C):
A homing pigeon (Columba livia domestica) serving as the physical transport medium for model parameters. The carrier maintains no persistent state between flights. This is considered a feature.

Parameter Scroll (PS):
A printed or written representation of model weights, gradients, or hyperparameters affixed to the carrier’s leg using archival-grade adhesive. Maximum payload is constrained by leg circumference and carrier tolerance.

Flock (F):
A distributed cluster of Carriers operating in parallel. Flock coherence is not guaranteed. Flock size SHOULD be determined by available loft space and local ordinances.

Loss Surface (LS):
The multidimensional landscape of model error as a function of parameter values. In APC implementations, the loss surface is not directly observable and must be inferred from returned scroll contents and the demeanor of the receiving researcher.

Electrostatic Discharge Event (EDE):
An unplanned transfer of static charge resulting in corruption or destruction of Parameter Scroll contents. Distinguished from a Carrier Attrition Event (CAE) in that the scroll arrives but is no longer legible. Both result in a dropped gradient. Only one requires replacing the hardware.

Carrier Morale (CM):
A scalar value representing a Carrier’s current disposition toward active dispatch duty. Carrier Morale is not directly measurable but may be inferred from observable behavioral indicators including: loft entry latency, scroll acceptance behavior, voluntary approach to the dispatch perch, and general demeanor. High Carrier Morale correlates with low epoch latency and reduced route deviation. Low Carrier Morale, particularly following a Parameter Pop event (Section 10.4), correlates with increased variance and non-standard flight paths.
Carrier Morale is the only hyperparameter in this specification that responds positively to grain, rest, and kind words. It does not respond to learning rate scheduling.

RAID-0 Dove Configuration (RDC):
A striped array of Carriers transporting complementary shards of a single Parameter Scroll. All shards MUST arrive for reconstruction to succeed. The fault tolerance of this configuration is zero. It is named optimistically.

Loft Telemetry Dashboard:
A containerized monitoring interface, deployable via Docker Compose on any OCI-compliant container orchestration platform, providing real-time telemetry from Smart Leg Band-equipped Carriers. Named for the philosophical concept of total surveillance. The irony of applying this to pigeons is acknowledged.

The following requirement keywords are used in this document:

MUST Non-negotiable.

SHOULD Recommended unless pigeons are unavailable.

MAY Use your judgment. The pigeons will use theirs.

MUST NOT Seriously. Do not do this.

SHOULD NOT The authors have tried this. It did not go well.

3. Architecture

3.1. Training Loop

The standard APC training loop proceeds as follows:

a. Initialize model weights. Print on acid-free paper.

b. Divide Parameter Scrolls across available Carriers. Assignment SHOULD be random to prevent positional bias. Carriers MUST NOT be allowed to self-select scrolls based on weight, as this introduces selection bias.

c. Dispatch Carriers to remote training node.

d. Wait. MAY occupy time by checking GPU availability on alternate infrastructure or reviewing current cloud billing status. The authors note that completing either task typically motivates continued investment in APC.

e. Upon return, aggregate received parameters. Carriers that do not return SHOULD be treated as dropped gradients. Apply gradient clipping accordingly.

f. Compute loss. Record on the Loss Log (see Section 3.3). Repeat from (b).

Each complete loop constitutes one training epoch. Epoch duration is a function of flight distance, headwind, and carrier motivation, the last of which is not currently formalizable.

3.2. Batch Size and Payload Constraints

Batch size is physically constrained by Flock size. Larger batches require proportionally more Carriers. The authors observe that very large batch training (batch size > 512) presents significant loft management challenges and is RECOMMENDED only for well-funded institutions.

Mini-batch gradient descent is the preferred approach. Stochastic single-sample updates via individual Carrier are theoretically supported but result in training instability proportional to wind conditions.

Payload capacity MUST be considered when selecting storage media for Parameter Scroll transport. A standard M.2 NVMe drive weighs approximately 7-10 grams. A homing pigeon’s maximum comfortable payload is approximately 75 grams, yielding a theoretical maximum of seven drives per Carrier. Practitioners SHOULD target a conservative 50% utilization to preserve aerodynamic performance and Carrier morale.

For model shards exceeding single-Carrier payload capacity, implementors MAY employ a RAID-0 Dove Configuration (RDC), in which a Parameter Scroll is striped across multiple Carriers dispatched simultaneously. Reconstruction at the receiving node requires all shards to arrive.

The fault tolerance of a RAID-0 Dove Configuration is zero. The authors wish to be unambiguous on this point. If one Carrier in an RDC does not return, the entire shard set MUST be discarded and the epoch repeated. Practitioners considering RDC for production workloads SHOULD reflect carefully on this property and then consider logistic regression instead.

3.3. Loss Monitoring and the Loss Log

Conventional ML infrastructure provides continuous, real-time loss curves updated at each training step. APC implementations do not support this.

Loss MUST be computed at the central loft upon gradient aggregation and recorded manually in the Loss Log (LL), a physical ledger maintained by the loft operator. The Loss Log SHOULD include:

i. Epoch number.

ii. Number of Carriers dispatched.

iii. Number of Carriers returned.

iv. Computed training loss.

v. Computed validation loss, if a validation Flock was dispatched.

vi. Weather conditions at time of dispatch.

vii. Smart Leg Band telemetry summary, if Appendix A hardware is deployed.

viii. Any anomalous Carrier behavior observed during the epoch, including but not limited to: prolonged absence, return without scroll, return with incorrect scroll, or return with scroll in a condition suggesting the Carrier sat on it.

Loss curves SHOULD be plotted by hand on graph paper. Automated plotting via the Loft Telemetry Dashboard is available to implementors who have deployed the hardware described in Appendix A.

A flat or increasing loss over consecutive epochs indicates the model is not converging. Practitioners SHOULD verify the following before concluding the architecture is at fault:

a. Scrolls are being attached prior to dispatch, not after.

b. Gradient direction is correctly indicated (descent, not ascent). The scroll orientation MUST be verified. The authors note this failure mode has occurred.

c. The validation Flock is distinct from the training Flock. Using the same Carriers for both constitutes data leakage and will produce artificially optimistic validation loss. This is not the Carriers' fault.

NaN loss values indicate a numerical instability in the gradient computation, or that the scroll was exposed to rain en route and is no longer legible. The two cases are distinguished by examining the scroll. If the scroll is damp, the loss is indeterminate. Discard the epoch. Allow Carrier to dry before reuse. Smart Leg Band-equipped deployments receive automated moisture alerts prior to Carrier return (see Appendix A, Appendix "Overview"), which is considered a significant operational improvement.

4. Bias-Variance Considerations

4.1. Variance

High variance in model performance will manifest as high variance in carrier return times. A model that has overfit to training conditions will show strong performance in fair weather and catastrophic degradation when a neighboring loft releases recreational pigeons. Practitioners SHOULD regularize accordingly.

4.2. Explainability

A principal advantage of APC over conventional black-box training infrastructure is full operational explainability. When a parameter update fails to arrive, the cause is observable:

i.   Carrier encountered adverse meteorological conditions.

ii.   Carrier was distracted by a local flock (see [BERGEN2001]).

iii. Cage was left open at the dispatch node. This is a human error and MUST be logged in the experiment record.

iv. Carrier did not return. Cause unknown. This is the honest answer and is more useful than a NaN loss.

Practitioners familiar with SHAP values will note that case (iii) is equivalent to a high-attribution input feature with a data pipeline error. The pigeon is blameless in both cases.

4.3. Chain-of-Flight vs. Explainability

Recent literature [BAREZ2025] establishes that chain-of-thought reasoning traces in large language models are not equivalent to genuine interpretability. The authors note that a pigeon's flight path is similarly non-explanatory: the carrier arrives or does not. The route taken is unobserved, non-reproducible, and likely influenced by factors outside the training distribution.

APC implementations MUST NOT treat carrier arrival as evidence that the intended route was followed. Gradient integrity SHOULD be verified upon receipt via checksum. Checksums MUST be printed in a font legible to the receiving researcher, as OCR error rates on leg-mounted scrolls remain non-trivial (see [BERGEN2001]).

Smart Leg Band GPS logging (Appendix A) provides partial flight path reconstruction and is RECOMMENDED for implementations requiring route auditability. The authors note that knowing the route taken does not constitute understanding why. This is also true of transformer attention maps.

5. Regularization Techniques

5.1. Weight Decay

Physical parameter scrolls degrade over repeated use due to moisture, mechanical wear, and carrier enthusiasm. This constitutes a natural form of weight decay. The decay rate is a function of weather and scroll material and is not hyperparameter-tunable in the current specification.

5.2. Early Stopping

Training SHOULD be halted when validation loss fails to improve over a defined number of epochs. In APC implementations, practitioners will typically identify this condition when they run out of grain before the model converges. This is considered a valid stopping criterion.

5.3. Dropout and Carrier Shuffling

Carrier attrition provides a natural analog to dropout regularization. A Carrier that does not return effectively masks the corresponding parameter subset from the current update. The dropout rate is environment-dependent and not directly configurable.

The authors recommend maintaining a 20% reserve Carrier pool to compensate. Practitioners SHOULD NOT attempt to implement structured dropout via deliberate Carrier interception. This is both logistically complex and ethically inadvisable.

A peer review of this document identified a statistical bias inherent to naive dropout-via-attrition: if certain Parameter Scrolls are consistently heavier, or if specific Carriers exhibit reduced motivation for particular routes, those gradient components will be non-uniformly dropped across epochs. This constitutes a systematic bias in which the weights most likely to be dropped are precisely those the Carrier finds most burdensome. The authors acknowledge this is not random dropout.

To mitigate this effect, Carrier-to-Scroll assignment MUST be shuffled between epochs. No Carrier SHOULD be assigned the same scroll position in consecutive epochs. Implementations SHOULD maintain a Carrier Rotation Log to verify compliance. The authors note that the Carriers themselves do not maintain such a log and cannot be relied upon to self-report assignment history.

5.4. Data Augmentation

Carriers operating in varied meteorological conditions implicitly expose the model to augmented training distributions. Rain, wind, and seasonal variation constitute a natural augmentation pipeline. No additional implementation is required. This is one of the few areas in which APC outperforms GPU-based training without qualification.

6. Distributed Training

6.1. Parameter Server Architecture

A central loft serves as the parameter server. Remote training nodes dispatch Carriers to the central loft bearing gradient updates. The central node aggregates received updates and returns updated weights via return Carrier.

Consistency guarantees are eventual. Practitioners accustomed to synchronous gradient aggregation SHOULD adjust expectations.

6.2. Federated Learning

APC is naturally suited to federated learning scenarios in which training data cannot leave the local node. Only gradients are transported, preserving data locality. Privacy guarantees are proportional to the discretion of all parties with access to the Carrier population.

6.3. Bandwidth-Delay Product

The effective bandwidth of an APC link is determined by Carrier payload capacity and round-trip flight time. For a Carrier transporting a 512GB NVMe drive over 5 kilometers at standard homing pigeon airspeed (approximately 80 km/h in favorable conditions), peak throughput substantially exceeds that of a T1 line. This comparison is made seriously and has been previously noted in the networking literature.

Latency remains non-competitive.

7. Transfer Learning

Pre-trained Carriers, having internalized routes to known destinations, exhibit faster convergence on familiar tasks. This is directly analogous to fine-tuning a pre-trained foundation model: the expensive representational work has already been completed; the practitioner need only adapt the terminal behavior.

Catastrophic forgetting has not been observed in Carriers. The authors attribute this to the comparatively limited parameter space of the avian hippocampus and the absence of gradient descent in biological systems.

8. Quality of Service

QoS extensions established in RFC 2549 [RFC2549] apply to APC without modification. Priority Carriers SHOULD be identified via distinct leg-band coloring. The authors note that Carriers do not observe priority markings and will proceed at their own discretion regardless.

Service classes from RFC 2549 (Concorde, First, Business, Coach) map naturally onto model sizes:

Concorde
: Frontier model, trillion-parameter scale. Requires a very large Carrier.

First
: 70B parameter range. Feasible with standard Carrier and favorable winds.

Business
: 7B parameter range. Recommended for most home-hosted implementations.

Coach
: Logistic regression [RUDIN2019]. A single scroll. Arrives reliably. Usually correct. Fully interpretable. Underrated. A sparrow could carry it. This is not an insult to logistic regression.

9. Data Privacy and Gradient Leakage

9.1. Regulatory Compliance

Parameter Scrolls in transit may constitute personal data under applicable privacy regulations if the training dataset contains personally identifiable information and the model has memorized it, as modern neural networks are known to do.

Under the General Data Protection Regulation (GDPR) and similar frameworks, gradient updates can encode and leak information about individual training samples. This property does not change when the gradients are printed on paper and attached to a bird. The medium is novel. The risk is not.

Practitioners MUST assess whether Parameter Scrolls constitute personal data under applicable law prior to dispatch. The authors observe that "we sent it via pigeon" is not currently recognized as a valid data transfer mechanism under Article 46 of the GDPR. This may change. The authors are not optimistic.

9.2. Gradient Inversion Attacks

Research has demonstrated that model gradients can be inverted to reconstruct training data with meaningful fidelity. An adversary intercepting a Parameter Scroll mid-flight could potentially reconstruct training samples from the gradient values encoded therein.

This attack vector is partially mitigated by the following factors inherent to APC:

a. Gradient values are encoded in human-readable decimal notation. Reconstruction requires the adversary to also solve OCR and decode the researcher’s handwriting.

b. The adversary must physically intercept the Carrier. This constrains the attack surface to parties with access to the flight corridor and sufficient reflexes.

c. Partial scroll delivery (see Section 5.3) means the adversary may only obtain a subset of gradient values, reducing reconstruction fidelity.

Practitioners handling sensitive training data SHOULD apply differential privacy mechanisms prior to printing [PHONG2018]. Adding calibrated Gaussian noise to gradient values before encoding them on the scroll is RECOMMENDED. The noise level MUST be sufficient to provide meaningful privacy guarantees and MUST NOT be so large as to render the gradient values indistinguishable from the researcher’s normal handwriting variation.

9.3. Data Subject Rights and the Predatory Interception Protocol

If a data subject exercises their right to erasure under applicable privacy law, the practitioner MUST ensure that gradient information derived from that subject’s data is purged from all Parameter Scrolls, including scrolls currently in transit.

Recalling Carriers mid-flight for the purpose of scroll amendment is not supported in the base specification. Peer review of this document identified this as a critical compliance gap under GDPR Article 17.

To address this gap, implementations requiring full regulatory compliance MAY deploy the Predatory Interception Protocol (PIP), in which a trained raptor (Falco peregrinus or equivalent) is dispatched to intercept and neutralize the Carrier prior to arrival at the destination node, thereby preventing the compromised scroll from entering the gradient aggregation pipeline.

The authors acknowledge that PIP raises significant concerns under Section 5.3 ("ethically inadvisable"). Practitioners operating under EU AI Act high-risk classification SHOULD consult legal counsel before implementing PIP. Practitioners operating under less stringent regulatory regimes MAY treat mid-flight scroll recovery as out of scope and document this decision in their Data Protection Impact Assessment.

The authors offer no opinion on whether a DPIA can legally reference a hawk.

10. Electrostatic Discharge (ESD) Considerations

10.1. Threat Model

Parameter Scrolls printed on standard paper and transported via avian carrier are subject to electrostatic discharge events arising from:

a. Triboelectric charging during flight, particularly at altitude or in low-humidity conditions.

b. Contact with synthetic loft materials, carrier leg-bands fabricated from acrylic or nylon, or researchers wearing polyester.

c. Rapid descent from altitude in dry atmospheric conditions, which may generate charge differentials sufficient to corrupt high-density gradient encodings.

d. The researcher removing the scroll from the Carrier without first grounding themselves. The authors note this is the most common ESD failure mode observed during informal testing and is entirely preventable.

A peer review of this document identified an additional threat not addressed in prior drafts: potential difference upon landing. A Carrier that has spent 45 minutes triboelectrically charging against the atmosphere arrives at the loft with a significant accumulated charge relative to ground. A loft operator who is themselves grounded presents as a preferred discharge path. The authors classify this as a "Parameter Pop" event and note it is unpleasant for all parties.

10.2. Mitigation Requirements

Practitioners MUST observe the following ESD precautions:

a. Loft operators MUST wear ESD-safe wrist straps when handling Parameter Scrolls. This applies during both attachment and retrieval operations. Wrist straps MUST be connected to a ground reference common with the receiving perch (see item (f) below).

b. Scrolls SHOULD be printed on anti-static paper where available. In the absence of anti-static paper, standard paper treated with anti-static spray is acceptable. Treated scrolls MUST be allowed to dry completely before attachment. A damp scroll introduces ambiguity (see Section 3.3, NaN loss handling).

c. The receiving loft MUST maintain relative humidity between 40% and 60%. Humidity below 40% significantly increases ESD risk. Humidity above 60% significantly increases scroll legibility risk. The authors acknowledge this is a narrow operational window.

d. Carriers MUST NOT be dispatched during thunderstorms. This requirement exists for multiple reasons, of which ESD is only the third most important.

e. Smart Leg Band implementations (Appendix A) MUST log triboelectric charge accumulation during flight and set the ESD_RISK flag upon landing if accumulated charge exceeds the defined threshold. The Loft Telemetry Dashboard SHOULD alert loft operators prior to scroll retrieval.

f. The receiving perch MUST be constructed from a dissipative material (surface resistivity between 10^5 and 10^9 ohms per square) and MUST be connected to building ground via a continuous conductor. This provides a controlled, low-energy discharge path for the returning Carrier and eliminates the loft operator as the path of least resistance. The authors consider this the most important addition prompted by peer review.

10.3. Scroll Integrity Verification

Upon receipt, each Parameter Scroll MUST be inspected for signs of ESD damage prior to gradient aggregation. Indicators include:

i. Scorch marks or discoloration inconsistent with normal soiling.

ii. Partial erasure of printed values in a pattern consistent with arc discharge rather than moisture.

iii. Values that parse as valid floating-point numbers but are implausible given the expected gradient magnitude for this architecture.

iv. The Carrier appears to have had a bad day.

Scrolls failing integrity verification MUST be discarded and the corresponding gradient treated as dropped. Under no circumstances SHOULD the practitioner attempt to recover partially legible values by interpolation and incorporate them into the gradient aggregate. The authors have done this. The model did not recover.

10.4. Shock-Based Training Instability (SBTI)

Peer review of this document identified a secondary consequence of the Parameter Pop event (Section 10.1) not addressed in prior drafts: behavioral impact on the Carrier.

A Carrier subjected to an uncontrolled electrostatic discharge event upon landing may exhibit reduced motivation in subsequent epochs. This manifests as increased route deviation, extended dwell time at intermediate locations, and in severe cases, a documented reluctance to re-enter the dispatch loft. The authors classify this as Shock-Based Training Instability (SBTI).

SBTI is operationally significant because it is self-reinforcing. A Carrier that experienced a Parameter Pop in epoch N is more likely to exhibit high variance in epoch N+1, producing the same high-variance loss characteristics as an overfit model, but arising from loft infrastructure failure rather than gradient pathology. The two causes are distinguished by consulting the Loss Log weather notes and the Carrier’s observable disposition. A model that overfit is a modeling problem. A Carrier that was electrocuted is a facilities problem. Both present identically in the loss curve.

Mitigation is addressed by the dissipative perch requirement (Section 10.2). The perch provides a controlled discharge path that eliminates the abrupt high-voltage event and replaces it with a gradual, low-energy equalization. A Carrier that lands on a properly grounded dissipative perch experiences no detectable discharge event and proceeds to the scroll retrieval queue with motivation intact.

The authors note that this is the only section of this document in which the welfare of the Carrier and the integrity of the gradient are addressed by the same physical component. The dissipative perch is therefore both an ESD mitigation and an animal welfare measure, and MUST be treated as mandatory on both grounds.

Carriers exhibiting persistent SBTI symptoms SHOULD be rotated to non-dispatch duties pending behavioral recovery. Forcing an SBTI-affected Carrier to continue active dispatch duty introduces systematic variance into the training loop that cannot be corrected by regularization. The Carrier Rotation Log MUST record SBTI events and the affected Carrier’s return-to-active status.

11. Security Considerations

11.1. General Threat Model

Parameter scrolls in transit are subject to interception, modification, and consumption. The last failure mode is novel to this transport layer and MUST be considered in threat modeling.

Practitioners operating in adversarial environments SHOULD encrypt parameter scrolls. Note that encrypted scrolls require a decryption step at the receiving node, which introduces latency proportional to the legibility of the researcher’s handwriting.

A Carrier that has been compromised MUST NOT be reintegrated into the Flock without full parameter re-initialization. Supply chain security for grain and nesting materials is out of scope for this document.

Model poisoning via corrupted gradient injection is theoretically possible if an adversary gains access to the dispatch loft. Physical perimeter security is RECOMMENDED. A lock on the cage door would also address the failure mode documented in [BERGEN2001] and the data subject erasure gap identified in Section 9.3.

11.2. Messenger-in-the-Middle (MitM) Attacks

A class of attack not addressed in RFC 1149 or RFC 2549 is the Messenger-in-the-Middle attack, in which an adversary intercepts a Carrier mid-route and substitutes a modified or entirely fabricated Parameter Scroll prior to release.

Known attack vectors include:

a. Breadcrumb-based route deviation, in which an adversary lays a trail of grain to redirect the Carrier to an intermediate node outside the intended flight corridor.

b. Unauthorized grain-based congregation points (UGBCP), in which a stationary food source in the flight corridor induces the Carrier to land and rest, during which time scroll substitution may occur.

c. Loft impersonation, in which a fraudulent destination loft is established within the Carrier’s homing radius, exploiting the Carrier’s navigation system rather than intercepting it in flight.

Mitigations include:

a. Scroll signing via printed cryptographic hash. The receiving node MUST verify the hash prior to gradient aggregation. This does not prevent interception but detects substitution.

b. Flight corridor hygiene. Practitioners SHOULD ensure the flight path is free of unauthorized grain sources. This is operationally difficult to enforce at scale.

c. Smart Leg Band GPS logging (Appendix A) provides post-hoc route verification. Scrolls from Carriers whose logged route deviates significantly from the expected path SHOULD be treated as suspect and subjected to enhanced integrity verification.

11.3. Prompt Injection via Avian Mimicry

A novel attack surface is introduced when APC infrastructure operates in geographic proximity to facilities housing psittacine or mimetic avian species, including but not limited to African Grey parrots (Psittacus erithacus), Common Hill Mynas (Gracula religiosa), and Northern Mockingbirds (Mimus polyglottos).

These species are capable of reproducing loft call signals, return confirmation sounds, and in documented cases, human speech patterns used in loft management. An adversary with access to a sufficiently trained mimetic bird could:

a. Trigger premature loft door opening via synthesized return call, allowing unauthorized Carrier ingress or egress.

b. Issue false verbal instructions to loft operators, including scroll handling commands, Carrier assignment directives, or dispatch authorizations.

c. In implementations using voice-activated Smart Leg Band pairing (Appendix A), inject unauthorized configuration commands into the telemetry pipeline.

The authors note that (b) and (c) are functionally equivalent to prompt injection attacks against large language models, in which adversarial input in the environment causes the model to execute unintended instructions. The mechanism differs. The effect does not.

Mitigations SHOULD include:

a. Out-of-band verification of verbal loft instructions via a secondary, non-auditory channel.

b. Perimeter exclusion zones for mimetic species. The authors acknowledge this is difficult to enforce in areas with established wild mockingbird populations.

c. Smart Leg Band voice command authentication SHOULD require a passphrase not reproducible by common mimetic species. Selection of an appropriate passphrase is left as an exercise for the implementor, subject to the constraint that African Grey parrots have demonstrated vocabulary acquisition in excess of one thousand words [PEPPERBERG1999]. Choose accordingly.

11.4. Carrier Visual Obfuscation and Adversarial Plumage Patterning

This section describes a dual-purpose technique combining Carrier visual obfuscation with steganographic gradient encoding. The technique addresses two distinct threats: adversarial human interception (Section 11.2) and opportunistic raptor predation (a threat implicitly present in all APC deployments but not previously formalized in this specification).

11.4.1. Threat: Raptor-as-Classifier

Birds of prey employ a visual classification system refined over approximately 65 million years of supervised learning on a dataset of considerable size. This classifier is highly optimized for the detection of Columba livia domestica in open airspace and represents a meaningful threat to Carrier availability, particularly at altitude.

The authors note that this classifier is also vulnerable to adversarial examples. Research in the human domain has demonstrated that carefully constructed visual perturbations can cause deep neural networks to misclassify objects with high confidence. The same principle applies to biological classifiers, including raptors. A Carrier whose visual appearance has been shifted outside the raptor's training distribution for "pigeon" is less likely to be correctly classified as prey.

This is not a new observation. It is the operating principle of every bird that is not a pigeon. This document formalizes it as a security mitigation.

11.4.2. Colorimetric Patterning Requirements

Carriers MAY be marked with animal-safe, non-toxic, water-soluble dyes in patterns selected to shift their visual classification away from Columba livia domestica. The following requirements apply:

a. All dyes MUST be certified non-toxic and animal-safe. Practitioners MUST verify certification before application. The authors are not responsible for outcomes arising from the use of uncertified dyes. This is not a hypothetical disclaimer.

b. Patterns SHOULD be designed to maximize dissimilarity from the natural appearance of Columba livia domestica as observed from above, which is the primary viewpoint of the raptor classifier. Side and frontal profiles are secondary threat surfaces.

c. Pattern selection SHOULD account for the known visual spectrum of local raptor species, which extends into ultraviolet wavelengths invisible to humans. A pattern that appears disruptive to the human eye MAY still be classified as "pigeon" by a UV-sensitive predator. The authors acknowledge this significantly complicates pattern design and note that no existing tool adequately addresses it. This is future work.

d. Patterns MUST be documented in the Carrier Rotation Log (Section 5.3) to ensure consistent identification of individual Carriers across epochs.

11.4.3. Steganographic Dual-Use Encoding

The colorimetric patterns applied per Section 11.4.2 MAY simultaneously encode gradient metadata using a pre-shared colorimetric key known only to the dispatch and receiving lofts. This constitutes a steganographic encoding layer in which:

a. An uninformed observer, including an adversary who successfully intercepts the Carrier, observes decorative or identification markings.

b. The receiving loft, equipped with the colorimetric key and a calibrated optical capture device (see Appendix A, Appendix "Optical Capture and Colorimetric Decode System"), decodes the pattern prior to physical handling of the Carrier, recovering gradient metadata before the Parameter Scroll is retrieved.

The steganographic layer MUST NOT be used as a substitute for scroll-level encryption (Section 11.1). It is an obfuscation layer, not a cryptographic one. An adversary who obtains the colorimetric key can decode all past and future transmissions. Key rotation SHOULD occur at regular intervals. Key rotation requires repainting the Carrier, which is addressed in Section 11.4.5.

11.4.4. Null Gradient Signaling via Pattern Absence

A significant operational advantage of the steganographic encoding scheme is the ability to distinguish between a Carrier that has returned with an intact scroll and a Carrier that has lost its scroll in transit.

Under conventional APC operation, a Carrier returning without a Parameter Scroll is indistinguishable at a distance from a Carrier returning with one. The loft operator must physically inspect each returning Carrier to determine scroll status, introducing latency and handling risk (see Section 10).

In steganographically-equipped deployments, the absence of expected colorimetric encoding on a returning Carrier constitutes a NULL_GRADIENT signal, interpretable by the optical capture system prior to Carrier landing. This enables:

a. Automated Loss Log annotation indicating scroll loss rather than aggregation failure.

b. Pre-emptive gradient clipping adjustment before the epoch is finalized.

c. Early notification to the loft operator that the Carrier should be directed to the scrub station before reintroduction to the active Flock.

The authors consider this one of the more practically useful contributions of this specification.

11.4.5. Carrier Preparation and the Scrub-Before-Reuse Requirement

The use of colorimetric patterning introduces an operational overhead not present in unencoded APC deployments: a Carrier bearing steganographic markings from a prior epoch MUST be fully scrubbed before reassignment to a new gradient payload.

Failure to scrub results in pattern contamination, in which residual encoding from a prior epoch is interpreted by the receiving loft’s optical system as current metadata, producing corrupted gradient annotations. This is functionally equivalent to gradient poisoning and MUST be treated as such.

Scrubbing requirements:

a. All prior colorimetric markings MUST be fully removed using warm water and a mild, animal-safe detergent. The Carrier MUST be allowed to dry completely before new markings are applied.

b. The Carrier MUST be inspected under both visible and, where equipment is available, ultraviolet light to confirm complete marking removal prior to re-encoding.

c. Scrub events MUST be logged in the Carrier Rotation Log with timestamp, operator identity, and confirmation of clean status.

d. A wet Carrier MUST NOT be dispatched. This requirement appears in multiple sections of this document. The authors note this is not coincidental.

The authors acknowledge that the scrub-before-reuse requirement adds meaningful operational overhead, particularly in high-throughput deployments with rapid epoch cycling. For a 70B parameter model dispatched across 140 Carriers, the time required to scrub, dry, inspect, and re-encode the entire active Flock will in many cases exceed the epoch flight time itself, creating a hard throughput ceiling that no gradient optimization can address.

To mitigate this constraint, implementations with sustained training workloads SHOULD adopt the Dual-Flock Pipeline (DFP), in which the total Carrier population is divided into two sub-flocks of equal size operating in alternating phase:

a. While Sub-Flock Alpha is In-Flight carrying the current epoch’s Parameter Scrolls, Sub-Flock Beta is in the Scrub/Dry/Re-encode pipeline being prepared for the subsequent epoch.

b. Upon Sub-Flock Alpha’s return and gradient aggregation, Sub-Flock Beta is immediately dispatched. Sub-Flock Alpha enters the scrub pipeline.

c. Epoch cadence is therefore limited by max(flight_time, scrub_dwell_time) rather than their sum.

The Dual-Flock Pipeline requires doubling the total Carrier population relative to single-flock operation. Practitioners MUST size the scrub facility accordingly. A scrub facility capable of processing N Carriers per hour MUST be available to support a Dual-Flock Pipeline with sub-flocks of size N.

The authors note that scrub facility throughput is a function of warm water availability, drying capacity, Carrier cooperation, and the number of researchers assigned to scrub duty. Of these, Carrier cooperation is the least configurable and the most consequential. Practitioners SHOULD factor a cooperation variance multiplier of at least 1.3x into scrub dwell time estimates.

This is the second documented case in which pigeon hygiene directly constrains model training throughput. The first was rain.

11.5. Security Considerations Summary

For implementors who have reached this section without reading Sections 11.1 through 11.4, the following summary is provided. The authors note that not reading the preceding sections is itself a security risk.

The APC architecture introduces three primary threat classes not present in conventional ML infrastructure:

Interception:
An adversary may physically intercept a Carrier in transit, obtaining the Parameter Scroll and potentially reconstructing training data via gradient inversion (Section 9.2). Mitigations include scroll encryption (Section 11.1), steganographic encoding (Section 11.4.3), and flight corridor hygiene (Section 11.2). Complete elimination of interception risk requires eliminating the flight corridor, which also eliminates the protocol. This is considered an acceptable tradeoff only if the researcher has access to alternative infrastructure.

Modification:
An adversary who intercepts and re-releases a Carrier bearing a modified or substituted scroll introduces corrupted gradients into the training loop (Section 11.2). This is distinguished from a hardware failure in that the corruption is intentional and may be designed to be undetectable. Scroll signing and GPS route verification provide partial mitigation. Neither is foolproof. Neither is the protocol.

Consumption:
An adversary, or a sufficiently motivated hawk, may consume the Carrier entirely, resulting in permanent loss of the gradient payload, the storage media, and the Carrier. This threat class is unique to APC among all distributed ML transport protocols currently in the literature. There is no cryptographic mitigation for consumption. Adversarial plumage patterning (Section 11.4) is the primary defense. The authors acknowledge this is a sentence that has not previously appeared in an IETF security summary.

12. Ethical Considerations

12.1. Ethical Treatment of Gradient-Bearing Carriers

The Carriers described in this specification are living organisms. Their welfare is not incidental to the operation of the protocol. It is a direct operational dependency. A Carrier whose welfare is neglected exhibits reduced Carrier Morale (Section 2), increased epoch latency, elevated route deviation, and a higher probability of non-return. Ethical treatment of Carriers is therefore both a moral obligation and a performance optimization. The authors note this is one of the few cases in distributed systems where the two are identical.

Practitioners MUST provide:

a. Adequate nutrition. Carrier Morale degrades measurably under caloric deficit. Grain quality SHOULD be appropriate to the Carrier’s operational workload. A Carrier dispatched on a long-haul RAID-0 configuration SHOULD receive proportionally more grain than one dispatched on a single-scroll Coach-class assignment.

b. Adequate rest. Dispatching a Carrier before full recovery from a prior epoch introduces fatigue-related variance into the training loop. Minimum rest intervals SHOULD be established based on round-trip distance and prevailing weather conditions. Rushing the rest interval to meet a training deadline is the avian equivalent of reducing the learning rate warmup period. The consequences are similar.

c. Veterinary access. Carriers exhibiting symptoms of illness, injury, or persistent SBTI (Section 10.4) MUST be removed from active duty and assessed by a qualified avian veterinarian. The Loss Log MUST record the removal and estimated return-to-duty date. A Carrier that is unwell is not a gradient delivery mechanism. It is an unwell bird. Treat it accordingly.

12.2. Bias in Carrier Selection

Selection of Carriers for dispatch assignments MUST NOT introduce systematic bias into the gradient transport process. Known sources of selection bias include:

a. Preferential assignment of lighter scrolls to smaller or lower-Morale Carriers. This produces non-uniform gradient dropout correlated with Carrier physical characteristics, which is not random dropout and cannot be modeled as such.

b. Consistent assignment of high-priority scrolls to a small subset of high-performing Carriers. This creates a single point of failure in which the loss of one or two Carriers disproportionately affects training outcomes. The Carrier Rotation requirement (Section 5.3) partially addresses this but MUST be actively enforced rather than assumed.

c. Unconscious preference for Carriers that have previously returned quickly, without accounting for the possibility that fast return times reflect route familiarity rather than general performance. A Carrier that is fast on a known route MAY be slow or unreliable on a novel route. Generalization cannot be assumed from training performance. The authors note this is also true of ML models.

12.3. Informed Consent for Raptors

The Predatory Interception Protocol (Section 9.3) involves the deployment of a trained raptor to intercept Carriers bearing compromised Parameter Scrolls. The raptor participates in this protocol involuntarily, as raptors cannot provide informed consent to participation in ML compliance workflows.

The authors acknowledge this is an unresolved ethical gap. Deployment of PIP MUST be reviewed by an institutional ethics board, or the nearest available equivalent, before implementation. Documentation of this review MUST be retained.

The authors also note, for completeness, that the Carrier subject to PIP has also not provided informed consent. This is noted without resolution. The GDPR does not currently address avian data subjects. This may change.

12.4. Environmental Considerations

Large-scale APC deployments involve significant numbers of Carriers operating in shared airspace. Practitioners MUST assess the environmental impact of their Carrier fleet on local ecosystems, including but not limited to:

a. Competition with wild pigeon populations for food resources within the flight corridor.

b. Disruption to existing bird of prey territories caused by increased prey density. Practitioners deploying adversarial plumage patterning (Section 11.4) to deter raptor predation should be aware that sustained deterrence may alter local predator foraging behavior in ways that extend beyond the flight corridor.

c. The carbon footprint of printing gradient values on paper at scale. The authors have not computed this figure. The authors suspect it compares favorably to a 1,800-GPU training cluster. The authors acknowledge they may be motivated to believe this.

13. Motivating Example: A Single Training Epoch

The following example illustrates a complete training epoch under the APC protocol using the hardware and procedures defined in this document. All values are representative. Weather conditions are based on Bergen, Norway in April, as this is the only location for which empirical APC performance data exists [BERGEN2001].

Configuration:

Model: Llama-class, 7B parameters (Business class)

Quantization: 4-bit, ~3.5GB effective size

Carriers: 7 (six active, one reserve)

Storage media: 1TB NVMe per Carrier (6TB total, 58% utilized)

Route distance: 5km one-way

Weather: Overcast, 12 deg C, wind NNW at 18km/h

Scrub status: All Carriers freshly scrubbed, dry, re-encoded

```
06:47  Researcher arrives at loft.  Attaches ESD wrist strap.
       Verifies perch ground continuity.  Confirms Carrier
       Morale is adequate.  One Carrier (C-4) appears
       reluctant.  C-4 is assigned the reserve role.

06:52  Scroll Header printed for each of six active Carriers:
       "PyTorch 2.7 / safetensors / 4-bit NF4 / Shard N of 6
       / Llama-7B layer 18-21 / Key v4"
       Scrolls attached.  Colorimetric encoding applied.
       Dispatch cage opened.

06:53  Carriers dispatched.  Wind slightly adverse.
       Loss Log entry: Epoch 47, 6 of 6 dispatched, 06:53.
       Weather: OVC, 12C, NNW 18.  C-4 on reserve.

07:38  C-2 returns first.  ESD_RISK flag: negative.
       Moisture reading: 14% (Optimal).  Colorimetric
       decode: valid, Key v4.  Scroll retrieved.
       Checksum: verified.  C-2 proceeds to re-encode queue.
```

```
07:41  C-5 returns.  ESD_RISK flag: negative.
       Moisture reading: 22% (Humid).  Ink Blur correction
       applied.  Gradient processed with high-variance flag.
       C-5 proceeds to re-encode queue.

07:44  C-1, C-3, C-6 return within 90 seconds of each other.
       All nominal.  Gradients verified.

08:31  C-4 has been in reserve for 98 minutes and is now
       being dispatched to replace C-2 while C-2 undergoes
       scrub.  C-4’s earlier reluctance has resolved
       following grain and rest.  Carrier Morale: restored.

09:15  No further Carriers expected.  C-2 (dispatched 08:31)
       has not yet returned.  This is within expected range
       given adverse wind.

09:47  C-2 returns.  Nominal.

10:00  GRADIENT AGGREGATION:
       Shards received: 6 of 6.
       High-variance flags: 1 (C-5, humidity).
       Dropped gradients: 0.
       Training loss: 1.847 (epoch 47).
       Validation loss: 2.103 (epoch 47).
       Delta from epoch 46: -0.031 training, -0.019 validation.
       Assessment: converging.  Continue training.
```

Loss Log entry closed. All Carriers in scrub/re-encode pipeline. Dual-Flock Beta dispatched at 09:55 carrying epoch 48 scrolls. Epoch 47 total elapsed time: 3h 07m. Model is learning. Slowly. This is expected.

The authors note that the above epoch proceeded without incident. This is not always the case. The Loss Log for epoch 23 contains the entry: "C-3 returned without scroll. Scroll location unknown. C-3 unavailable for comment." This entry has not been resolved.

14. IANA Considerations

This document has no IANA considerations. The authors previously held this position without elaboration.

Following reviewer feedback, the authors have reconsidered and now formally propose the following IANA registries for APC implementations:

Carrier Status Code Registry: A registry of standardized status codes for Carrier disposition, analogous to HTTP status codes. Proposed initial entries:

200 Returned nominal. Gradient accepted.

204 Returned nominal. Scroll missing. No content.

301 Redirected to alternative loft. Cause unknown.

404 Did not return. Carrier not found.

408 Return timeout. Epoch terminated.

418 I am a teapot. (Reserved. Applicability to avian carriers is under investigation.)

500 Internal Carrier error. See veterinarian.

503 Carrier temporarily unavailable. Scrubbing.

Scroll Header Framework Identifier Registry: A registry of valid framework identifier strings for use in the Scroll Header (Section 15.2). New entries MAY be submitted by framework maintainers. The registry MUST include a deprecation date field. Entries for frameworks that have been deprecated SHOULD be retained for historical reference, as practitioners MAY still encounter scrolls bearing deprecated identifiers in long-running loft archives.

The authors note that neither registry requires IANA action at this time, as APC has not achieved sufficient deployment scale to warrant formal registration. The authors express confidence that this will change.

The authors considered requesting a new IP protocol number for APC but concluded that the existing best-effort delivery model of IP adequately captures the operational characteristics of the transport layer.

15. Scalability Considerations and Forward Compatibility

15.1. Model Size Scaling

At the time of this writing, frontier models are estimated at 1.7-1.8 trillion parameters (Section 8, Concorde class). The trajectory of model scaling suggests this figure will continue to increase. The APC architecture must address the physical implications of this trend.

A single 32-bit floating-point parameter requires 4 bytes of storage. A 1.8 trillion parameter model therefore requires approximately 7.2 terabytes in full precision, or approximately 900 gigabytes at 1-bit quantization. At the time of publication, no commercially available leg-band storage medium approaches this capacity.

Practitioners SHOULD apply aggressive quantization prior to scroll encoding. The authors note that 4-bit quantization, now standard practice for local inference, reduces the 1.8T parameter model to approximately 900GB, which is achievable via a multi-Carrier RAID-0 Dove Configuration (Section 3.2) with a fleet of approximately 1,800 Carriers each carrying a 512GB NVMe drive.

The authors acknowledge that a fleet of 1,800 Carriers represents a meaningful operational commitment. Institutions unable to sustain this fleet size SHOULD consider whether they are training the correct model for their resource envelope. This advice applies equally to conventional infrastructure.

As model sizes continue to grow, this specification anticipates two possible evolutionary paths:

a. Advances in storage media density will increase per-Carrier payload capacity, maintaining feasibility at current Carrier-to-parameter ratios.

b. Advances in model compression, quantization, and distillation will reduce effective model size faster than raw parameter counts grow, improving the ratio in the other direction.

The authors consider path (b) more likely and note that it represents a genuine alignment of interests between the ML efficiency research community and the avian transport community. This alignment has not previously been documented.

15.2. Framework and Format Versioning

The ML framework ecosystem evolves rapidly. Parameter Scroll format compatibility across framework versions is a non-trivial operational concern. A scroll serialized under PyTorch 2.x tensor format may not be directly interpretable by a receiving node running a subsequent major version, a different framework entirely, or a researcher who has not updated their deserialization tooling since the model was dispatched.

This specification REQUIRES that all Parameter Scrolls include a Scroll Header (SH) prepended to the gradient payload. The Scroll Header MUST be printed in a standardized, human-readable format and MUST contain:

a. Framework name and major version (e.g., "PyTorch 2.7").

b. Serialization format identifier (e.g., "safetensors", "pickle", "GGUF").

c. Quantization scheme and bit depth.

d. Scroll sequence number and total scroll count, for RAID-0 Dove Configuration reassembly.

e. Colorimetric key version, if steganographic encoding is in use (Section 11.4).

f. A single-line human-readable description of the model architecture sufficient to detect obvious mismatches at the receiving node before gradient aggregation begins. The authors RECOMMEND a format such as: "Llama-class, 70B, MoE 8x9B, layer 42 of 80." A receiving researcher who reads this line and does not recognize the architecture SHOULD NOT proceed with aggregation.

The Scroll Header MUST be verified by the receiving node before the gradient payload is processed. A version mismatch MUST produce a clear error in the Loss Log and MUST NOT result in silent gradient corruption. The authors note that silent gradient corruption on conventional infrastructure is a well-documented and deeply unpleasant failure mode. The scroll header exists precisely to make this class of error loud and attributable.

15.3. Topology Scaling: Beyond the Single Parameter Server

Section 6.1 describes a hub-and-spoke topology with a central parameter server loft. This topology does not scale to the distributed training configurations required by frontier models, which employ pipeline parallelism, tensor parallelism, and expert parallelism across hundreds or thousands of nodes.

This specification defines three extended topologies for large-scale APC deployments:

Ring-Flock Topology:
Carriers travel a circular route between N training nodes, each node updating its assigned parameter shard before dispatching the Carrier to the next node. Gradient aggregation is distributed

across the ring. Latency scales linearly with ring size. The authors note that a Ring-Flock of 64 nodes, each at 5km separation, describes a flight circuit of 320km. At standard homing pigeon cruising speed this represents an epoch duration of approximately 4 hours under favorable conditions. This is slower than conventional ring-AllReduce but requires significantly less InfiniBand cabling.

Hierarchical Flock Topology:
Training nodes are organized into regional clusters, each with a local aggregation loft. Local gradients are aggregated regionally before a subset of Carriers transport the regional aggregate to the global parameter server. This reduces total Carrier count at the cost of introducing aggregation delay at two levels. Practitioners familiar with gradient compression and local SGD will recognize this topology. The pigeon version offers equivalent theoretical properties and superior scenic variety.

Expert-Parallel Flock Topology:
For Mixture-of-Experts architectures, each expert’s parameters are assigned to a dedicated sub-flock. Expert routing decisions are encoded in a separate Routing Scroll dispatched ahead of the parameter Carriers. The Routing Scroll MUST arrive before the parameter Carriers to avoid aggregation at the wrong expert node. In practice, this means the Routing Scroll Carrier MUST be dispatched first and MUST be a faster-than-average Carrier. Selection criteria for the Routing Scroll Carrier are outside the scope of this document but SHOULD include demonstrated navigational reliability and a history of not joining recreational flocks en route.

15.4. Inference Scaling and Serving

This specification has addressed training throughput in detail. Inference serving presents distinct scaling challenges.

A deployed model receiving queries must return predictions within a latency envelope acceptable to the requesting application. For most applications this envelope is measured in milliseconds to seconds. APC inference latency is measured in hours. This gap is not currently bridgeable within the constraints of avian flight physics.

The authors therefore formally RECOMMEND that APC implementations decouple training and inference infrastructure. APC is appropriate for the training loop. Inference SHOULD be served via conventional digital infrastructure after model weights have been transported to the serving node by Carrier and loaded onto appropriate hardware.

The authors note that this hybrid architecture -- avian training transport, digital inference serving -- is consistent with the broader principle that the right tool should be selected for each phase of the ML lifecycle. APC excels at asynchronous, high-payload, low-frequency gradient transport. It does not excel at sub-second token generation. These are not the same problem. Treating them as the same problem is a category error that no amount of grain will resolve.

15.5. New Framework Onboarding

As new ML frameworks emerge, this specification requires only that their serialization formats be registerable in the Scroll Header framework identifier field (Section 15.2). No changes to the physical transport layer are required. The Carrier does not care what framework serialized the weights it is carrying. This is one of the more durable properties of the APC architecture and is considered a design strength.

Framework deprecation SHOULD be handled by sunsetting the corresponding Scroll Header identifier. Scrolls bearing deprecated framework identifiers MUST be flagged at the receiving node. Whether to process them anyway is a decision for the receiving researcher, who by that point presumably knows what they are doing. Or does not, and the Loss Log will reflect this.

15.6. Quantum Computing Compatibility

The emergence of fault-tolerant quantum computing raises the question of whether APC remains viable as a transport layer for quantum ML workloads, or whether quantum methods will supplant the requirement for avian carriers entirely.

The authors address this in two parts.

Part I: Quantum ML Parameter Transport. Quantum ML models represent parameters as quantum states rather than classical floating-point values. Quantum states cannot be copied without disturbance (the No-Cloning Theorem, Wootters and Zurek, 1982 [NOCLONING]). A Parameter Scroll encoding a quantum state would therefore constitute a measurement of that state, collapsing the superposition and destroying the quantum information in the act of printing it.

The authors conclude that APC is fundamentally incompatible with native quantum parameter transport. A pigeon cannot carry a qubit. More precisely: by the time the qubit has been attached to the pigeon’s leg, it is no longer a qubit. It is a classical value, and the quantum advantage has been surrendered to the leg band.

This is not a limitation of the pigeon. It is a limitation of measurement. The pigeon is, once again, blameless.

Part II: Will Quantum Methods Supplant the Pigeon? Quantum computers offer theoretical speedups for specific problem classes, including optimization problems relevant to ML training. However, the question of whether quantum methods will supplant APC conflates two distinct concerns:

a. Whether quantum computers will perform ML training faster than classical computers.

b. Whether, given (a), the parameter transport layer becomes irrelevant.

Regarding (a): quantum advantage for general ML training remains undemonstrated at scale. Current quantum hardware operates at qubit counts and error rates that preclude practical ML workloads. The authors note that "practical quantum ML" has been approximately five years away for approximately fifteen years. Carrier fleets established today are unlikely to be rendered obsolete by quantum hardware before they reach retirement age.

Regarding (b): even assuming quantum ML training achieves practical advantage, the trained model weights must still be transported to inference infrastructure. If those weights are classical (as is likely for deployed models, given the state of quantum memory), the APC architecture remains a viable transport option.

The authors therefore conclude that quantum computing does not supplant the requirement for pigeons. It may, in the long term, change what is printed on their scrolls. The scrolls themselves, and the birds carrying them, remain relevant.

A future revision of this specification MAY address hybrid quantum-classical parameter transport in which classical gradient approximations of quantum circuit parameters are encoded on Parameter Scrolls. The authors consider this a tractable extension. The pigeons have no opinion on quantum mechanics and are not expected to develop one.

16. Known Limitations

The following limitations are acknowledged:

a. Gradient aggregation latency is measured in hours. Real-time inference is not supported in this release.

b. Carriers operate during daylight hours only. Overnight training runs require advance planning and a night-shift loft manager.

c. Carrier availability is subject to seasonal variation, molting cycles, and local predator populations. A high-availability configuration MUST include redundant Carrier capacity. Deployments utilizing the Predatory Interception Protocol (Section 9.3) MUST account for the additional impact on Carrier availability.

d. The current specification does not support backpropagation through time. Or through weather.

e. Loss curves are updated at most once per epoch. This is substantially less frequently than practitioners accustomed to TensorBoard will expect. The authors recommend the Loss Log (Section 3.3) and patience.

f. Recalling a Carrier mid-flight to correct a data privacy error requires implementation of PIP (Section 9.3). Regulatory and ethical review is advised before deployment.

g. ESD wrist strap compatibility with common loft gloves has not been verified. This is future work.

h. The RAID-0 Dove Configuration (Section 3.2) provides no fault tolerance. The authors consider this a feature of accurate nomenclature rather than a limitation of the specification.

i. Passphrase selection for mimetic species resistance (Section 11.3) is an open research problem.

j. Colorimetric pattern design in the ultraviolet spectrum requires equipment and expertise not commonly available in resource-constrained environments. Raptor threat modeling for UV-sensitive predators remains an open problem. Practitioners in hawk-dense geographic regions SHOULD treat this as a priority. Others MAY defer.

k. The scrub-before-reuse requirement (Section 11.4.5) introduces Carrier dwell time as a potential training throughput bottleneck. Pool sizing MUST account for scrub duration. Scrub duration is a function of marking density, water temperature, Carrier cooperation, and the researcher’s patience. Only the first two are configurable.

l. Interpretability is excellent. Performance is adequate. These are not unrelated observations.

17. References

17.1. Normative References

[RFC1149] Waitzman, D., "A Standard for the Transmission of IP Datagrams on Avian Carriers", RFC 1149, April 1990, <https://www.rfc-editor.org/rfc/rfc1149>.

[RFC2549] Waitzman, D., "IP Datagrams on Avian Carriers with Quality of Service", RFC 2549, April 1999, <https://www.rfc-editor.org/rfc/rfc2549>.

17.2. Informative References

[BAREZ2025]
Barez, F. and Y. Bengio, "Chain-of-Thought Is Not Explainability", Oxford Martin AI Governance Initiative, July 2025, <https://aigi.ox.ac.uk/publications/chain-of-thought-is-not-explainability/>.

[BERGEN2001]
Bergen Linux User Group, "Implementation of RFC 1149: Informal Report", April 2001, <https://blug.linux.no/rfc1149/writeup/>.

[NOCLONING]
Wootters, W. K. and W. H. Zurek, "A Single Quantum Cannot Be Cloned", Nature, Vol. 299, pp. 802-803. Cited in Section 14.6 to explain why a pigeon cannot carry a qubit. The authors note that Wootters and Zurek did not anticipate this application of their theorem. The theorem holds regardless., 1982.

[PEPPERBERG1999]
Pepperberg, I. M., "The Alex Studies: Cognitive and Communicative Abilities of Grey Parrots", Harvard University Press. Cited in support of Section 11.3 threat modeling. Alex knew over 100 words and could identify objects by color, shape, and material. He was a legitimate security concern., 1999.

[PHONG2018]
Phong, L. T., "Privacy-Preserving Deep Learning via Additively Homomorphic Encryption", IEEE Transactions on Information Forensics and Security. Cited for gradient inversion attack background. The authors note this threat predates the pigeon., 2018.

[RUDIN2019]
Rudin, C., "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead", Nature Machine Intelligence, January 2019, <https://arxiv.org/abs/1811.10154>.

Acknowledgments

The authors thank the Bergen Linux User Group for their foundational empirical work, without which the latency figures in Section 6.3 would be theoretical rather than measured.

The authors also thank the three Carriers from the 2001 Bergen test whose fate remains undocumented. Their contribution to the literature is noted.

The authors thank the reviewers whose comments prompted the addition of Sections 9 and 10, the Carrier Shuffling requirement in Section 5.3, the RAID-0 Dove Configuration definition, the dissipative perch specification, the Messenger-in-the-Middle section, the prompt injection via avian mimicry threat model, the steganographic plumage encoding scheme, the NULL_GRADIENT signaling protocol, and the scrub-before-reuse hygiene requirement. The quality of this document is directly attributable to their rigor. The scrub-before-reuse requirement in particular represents a genuine operational insight that the authors had not previously considered and are somewhat embarrassed not to have anticipated.

The authors thank Alex the parrot (1976-2007) posthumously. His last words were "You be good. I love you." He did not live to see the threat he represented formalized in an RFC.

This document was submitted late due to pigeons.

Smart Leg Band: Hardware Specification and Integration

Overview

The Smart Leg Band (SLB) is an optional hardware extension to the APC architecture providing in-flight telemetry, automated scroll integrity monitoring, and loft handshake capabilities. Deployment of the SLB is OPTIONAL but RECOMMENDED for implementations where scroll loss rate exceeds 20% or where regulatory compliance requires documented chain-of-custody for Parameter Scrolls.

Hardware Components

The reference implementation specifies the following components:

Microcontroller:
ESP32-C3, selected for compact footprint, integrated Wi-Fi and Bluetooth LE, and deep sleep current draw of approximately 5 microamperes. The ESP32-C3 MUST be coated with a conformal protective resin prior to installation. Pigeons WILL attempt to debug uncoated hardware using their beaks. This is not hypothetical. The authors classify uncoated ESP32 exposure to Carrier grooming behavior as "Physical Bit-Rot" and note that it is entirely preventable. Coat the board.

Moisture Sensor:
A capacitive sensor element modified for paper-contact operation, wrapped circumferentially around the Parameter Scroll beneath the leg band. Provides continuous humidity readings for the scroll surface. Threshold values and corresponding actions are defined in Appendix "Table 1: Moisture Level Action Matrix".

Power Supply:
A single-cell LiPo battery, minimum 100mAh, selected to maintain the Carrier’s payload budget within the constraints of Section 3.2. The ESP32-C3 MUST operate in Deep Sleep between sample intervals to preserve capacity across the full expected flight duration. Wake interval SHOULD be configurable but SHOULD NOT exceed 5 minutes to ensure adequate temporal resolution of moisture events.

Loft Detection:
Either (a) Wi-Fi RSSI-based proximity detection, triggering loft handshake when signal strength from the loft access point exceeds a configurable threshold, or (b) a Hall Effect sensor at the loft trap door, triggering on Carrier ingress. Option (b) is preferred for reliability. Option (a) is preferred when the researcher has already 3D-printed the trap door and does not wish to reprint it.

ESD Monitor:
Capacitive touch pin monitoring for triboelectric charge accumulation. Logs timestamp and estimated charge level for any discharge event detected during flight. Sets ESD_RISK flag in metadata if accumulated charge exceeds operational threshold upon approach to loft.

Table 1: Moisture Level Action Matrix

Moisture (%)	WET_BIT	Action
0 - 20	0	Optimal. Process gradient normally.
21 - 50	0	Humid. Apply Ink Blur correction algorithm prior to OCR.
51 - 80	1	Damp. MAY attempt OCR. Treat resulting gradient as high-variance. Weight accordingly in aggregation.
> 80	1	MUST discard. Treat as NULL gradient. Allow Carrier to towel-dry before reassignment. Do not attempt OCR. The authors have attempted OCR at this moisture level. The results were not gradients.

Data Ingestion and Telemetry

The SLB operates in Store-and-Forward mode. All sensor readings are logged to the ESP32-C3 internal flash during flight. No active transmission occurs in-flight, as Wi-Fi connectivity at operational altitude is not supported and the power budget does not accommodate sustained radio operation.

Upon Carrier entry into the loft (detected per Appendix "Hardware Components", Loft Detection), the ESP32-C3 exits Deep Sleep, connects to the local Wi-Fi network, and transmits the accumulated telemetry log to the Loft Telemetry Dashboard prior to any physical interaction with the Carrier by loft personnel.

The Loft Telemetry Dashboard is a containerized web application deployable via Docker Compose on any OCI-compliant container orchestration platform. It provides:

a. Real-time display of incoming SLB telemetry.

b. Automated WET_BIT alert, notifying loft operators of scroll moisture status before handling.

c. ESD_RISK flag display, prompting operators to verify wrist strap grounding and perch continuity before approaching the Carrier.

d. Automated Loss Log population for moisture-flagged epochs, reducing manual transcription errors.

e. Carrier Rotation Log maintenance, enforcing the shuffling requirement of Section 5.3.

f. MitM route deviation alerts, flagging Carriers whose GPS track deviates from the expected corridor by more than a configurable threshold.

Docker Compose configuration for the Loft Telemetry Dashboard is available in the project repository. The authors note that the repository does not currently exist but express confidence that it will by the time this RFC is published.

Optical Capture and Colorimetric Decode System

Implementations deploying steganographic scroll encoding (Section 11.4) MUST equip the receiving loft with a calibrated optical capture device positioned to image each returning Carrier prior to landing on the receiving perch.

Hardware requirements:

a. Camera resolution MUST be sufficient to resolve the colorimetric encoding pattern at the expected approach distance. Minimum 12 megapixels is RECOMMENDED for standard leg-band pattern densities.

b. Where Section 11.4.2(c) UV threat modeling is in scope, the capture device SHOULD include a UV-capable sensor or a secondary UV-band camera. Standard CMOS sensors are not UV-sensitive and will not detect residual markings in that spectrum.

c. Lighting at the approach corridor MUST be consistent and controlled. Variable ambient lighting introduces colorimetric decoding errors. The authors RECOMMEND a covered approach corridor with fixed artificial lighting calibrated to the colorimetric key’s reference illuminant. The authors note this is a meaningful infrastructure investment and that a well-lit trap door is also acceptable for low-security deployments.

Decode pipeline:

a. On Carrier approach detection (via SLB RSSI or Hall Effect trigger, per Appendix "Hardware Components"), the optical system captures a reference image.

b. The captured image is processed against the current colorimetric key to extract gradient metadata.

c. If no valid encoding is detected, the system sets NULL_GRADIENT status for this Carrier and notifies the loft operator that scroll retrieval will yield no gradient data and that the Carrier should proceed directly to the scrub station.

d. If encoding is detected but fails key validation, the Carrier MUST be quarantined. This indicates either key mismatch (epoch management error) or scroll substitution by an adversary who obtained the colorimetric key but not the full key schedule. Treat as a potential MitM event per Section 11.2.

e. Decoded metadata MUST be fused with SLB telemetry before the Loss Log entry for this epoch is finalized.

The optical capture system operates independently of the SLB and requires no hardware on the Carrier. It is therefore compatible with non-SLB deployments that have opted into colorimetric encoding only. The authors consider this a useful deployment flexibility and note that it also means the Carrier’s ESP32-C3 is not required to know anything about the steganographic layer, which simplifies firmware scope and reduces the surface area available to beak-based debugging.

Firmware Security

SLB firmware MUST be cryptographically signed. Over-the-air firmware updates are supported via the loft Wi-Fi network and MUST be authenticated before installation.

Voice-activated configuration commands, if implemented, MUST be authenticated via passphrase as specified in Section 11.3. The authors reiterate that the passphrase MUST be selected with awareness of local mimetic species populations. An African Grey parrot has a documented vocabulary exceeding one thousand words [PEPPERBERG1999]. Implementors in affected regions are advised to plan accordingly.

Physical access to the ESP32-C3 is prevented by conformal resin coating (Appendix "Hardware Components"). The authors wish to be clear that this coating serves dual purpose: it protects against moisture ingress, and it protects against the Carrier. Both threats are real. Both are addressed by the same countermeasure. This is considered an elegant solution.

Author’s Address

D. Fairaizl
Independent Researcher
Undisclosed Location
Email: a.fairaizl@gmail.com