CREDASIS AI

Unsupervised Fidelity-Based Anomaly Detection

Results on CIC Datasets

PATENT PENDING · MAY 31, 2026 · MARTY OELRICH, COFOUNDER & CHIEF SCIENTIST

We have developed an unsupervised detection system that produces per-class, per-dimension, operator-tunable results on CICIDS-2017 and CSE-CIC-IDS2018. No attack labels are used for calibration or detection. Labels are used only for stratified fold construction and post-hoc benchmark scoring.

The system calibrates from benign traffic — under 30 seconds in Precision mode, several minutes in Detection mode. No GPU. No internet. No signatures. The operator controls one parameter — sensitivity — and everything else emerges from calibration on the deployment environment's own traffic. Thresholds are derived from held-out benign rows only.

Every benchmark-pipeline number below is reproducible from our deployment package. We report every class, including the ones where the system is weak.

What Is Different

The system does not optimize a single aggregate score. It measures behavioral fidelity across multiple independent dimensions simultaneously and preserves the decomposition for the operator. The operator sees which dimensions deviate, by how much, and whether independent measurements agree.

This changes what the system reports: not "this flow is malicious" but "this flow deviates on these dimensions, with this level of agreement, at a sensitivity setting you chose." Per-class recall at every sensitivity level replaces a single F1 number.

No attack labels are used for calibration or detection. Labels are used only for stratified fold construction and post-hoc benchmark scoring. Detection generalizes across datasets without retuning. The same code and parameters produce the CIC-2018 results without modification.

CICIDS-2017 Corrected (Engelen 2021, 15 classes, 5-fold CV)

Two operating modes from a single codebase.

PER-CLASS RECALL (A=0.05)

ATTACK CLASS	N	PRECISION MODE	DETECTION MODE
DoS Hulk	158,468	1.000	1.000
DDoS	95,144	1.000	1.000
Botnet	736	1.000	1.000
Web Brute Force	73	1.000	1.000
Web XSS	18	1.000	1.000
Heartbleed	11	1.000	1.000
Infiltration	36	0.980	1.000
DoS Slowhttptest	1,740	0.998	0.998
Infiltration-Portscan	71,767	0.993	0.998
FTP-Patator	3,972	0.603	0.995
SSH-Patator	2,961	0.988	0.995
DoS Slowloris	3,859	0.912	0.984
Portscan	159,066	0.947	0.952
DoS GoldenEye	7,567	0.936	0.939
SQL Injection	13	0.067	0.077

Precision mode: 11 of 15 classes at or above 0.93 recall. Detection mode: 14 of 15. SQL Injection (n=13) is a payload-level attack below the resolution of flow metadata. We report it.

Tuesday (Patator) is the hardest day for aggregate metrics — 7K attacks in 315K benign flows. Precision mode reaches FTP-Patator R=0.603 at α=0.05; Detection mode reaches 0.995.

Operating Curves

PRECISION MODE — FRIDAY (DDOS, PORTSCAN, BOTNET)

A	PREC	REC	FPR	F1
0.001	0.996	0.283	0.1%	0.440
0.005	0.990	0.478	0.4%	0.644
0.01	0.984	0.650	0.8%	0.784
0.02	0.973	0.836	1.5%	0.902
0.03	0.964	0.944	2.3%	0.954
0.05	0.959	0.968	3.8%	0.963
0.07	0.945	0.976	5.1%	0.960
0.10	0.927	0.993	7.1%	0.959

DETECTION MODE — FRIDAY

A	PREC	REC	FPR	F1
0.001	0.904	0.552	5.3%	0.685
0.005	0.908	0.589	5.4%	0.715
0.01	0.922	0.713	5.4%	0.804
0.02	0.933	0.869	5.6%	0.900
0.03	0.937	0.944	5.7%	0.941
0.05	0.933	0.970	6.3%	0.951
0.07	0.927	0.978	7.0%	0.952
0.10	0.915	0.994	8.3%	0.953

PRECISION MODE — WEDNESDAY (DOS VARIANTS)

A	PREC	REC	FPR	F1
0.001	0.997	0.626	0.1%	0.765
0.005	0.991	0.905	0.5%	0.946
0.01	0.984	0.947	0.9%	0.965
0.02	0.971	0.970	1.6%	0.971
0.03	0.958	0.976	2.4%	0.967
0.05	0.936	0.991	3.8%	0.963
0.10	0.888	1.000	7.0%	0.940

PRECISION MODE — THURSDAY (INFILTRATION, WEB ATTACKS)

A	PREC	REC	FPR	F1
0.01	0.940	0.499	0.8%	0.651
0.02	0.931	0.825	1.6%	0.874
0.03	0.912	0.926	2.3%	0.918
0.05	0.867	0.967	3.8%	0.915
0.10	0.778	0.973	7.1%	0.864

CSE-CIC-IDS2018 (No Retuning)

Same engine. Same parameters. No modification for CIC-2018.

ATTACK CLASS	FILE	N	PRECISION (A=0.05)	DETECTION (A=0.05)
SSH-BruteForce	Wed-14	94,197	1.000	1.000
DDoS-HOIC	Wed-21	248,069	1.000	1.000
DDoS-LOIC-HTTP	Tue-20	246,049	0.980	1.000
DDoS-LOIC-UDP	multiple	815	1.000	1.000
DoS GoldenEye	Thu-15	22,560	0.997	1.000
DoS Slowloris	Thu-15	8,490	1.000	1.000
DoS Hulk	Fri-16	247,177	0.912	1.000
Web Brute Force	multiple	131	1.000	1.000
Web XSS	multiple	113	1.000	1.000
Infiltration-NMAP	Thu-01	17,407	0.990	0.999
Infiltration-NMAP	Wed-28	21,591	0.885	0.988
Web SQL	multiple	39	0.650–0.836	0.650–0.836

Detection Mode improves most classes not already at 1.000 (Web SQL unchanged). Both modes validated via benchmark pipeline; production wrapper Detection Mode pending a loading fix for files with attacks at end of file.

Adversarial Validation

TEST	RESULT
Temporal split (train 60%, test 40%, no shuffle)	Detection holds across time-shifted traffic
Calibration contamination	Guard-enabled Tuesday E2E refused 1%, 3%, 5% poisoned calibration
Behavioral head ablation	Per-dimension importance quantified
Cross-dataset transfer	2017 → 2018 without retuning

Calibration poisoning at 1% is catastrophic without the integrity guard. With the guard, the system detects contamination and refuses to go live. Guard-enabled E2E validation completed on Tuesday (Patator classes). Multi-day guard-enabled validation is pending. We document both the vulnerability and the defense.

Documented Limits

LIMITATION	DETAIL
SQL Injection (n=13)	R=0.077. Payload-level, below flow metadata resolution.
Detection mode FPR	Adds 2.5% FPR over Precision mode.
Feature extraction	Requires CICFlowMeter-compatible feature extraction.
Calibration assumption	Clean benign baseline required (integrity guard addresses this).

Specifications

PROPERTY	VALUE
Labels for calibration/detection	None
Calibration time	< 30s (Prec) / mins (Det)
Engine size	560 KB
Total deployment footprint	< 70 MB
Calibration state	~50 MB
Runtime memory	~200 / ~250 MB

PROPERTY	VALUE
Hardware	CPU only
Air-gapped deployment	Yes
Per-flow latency	~1 ms / ~4 ms
Throughput	~1M / ~250K flows/s
Operator parameters	Sensitivity (α), mode
Reproducibility	MD5-tracked

Deployment Security

The engine includes a calibration integrity guard (frozen reference verifies data before deployment), continuous runtime coherence monitoring, autonomous shutdown on integrity compromise, and cryptographic provenance on detection verdicts. The system operates air-gapped with no telemetry and no external dependencies at runtime.

Methodology

5-fold StratifiedKFold (shuffle=True, random_state=42). CICIDS-2017: Engelen, Rimmer & Joosen (IEEE SPW 2021), Distrinet/KU Leuven, MD5-verified. CSE-CIC-IDS2018: CSE & CIC with Distrinet corrections. Thresholds from held-out benign rows only. Hardware: Apple M2 Mac Mini, 16GB, CPU only. Engine v2.1.

Independent Evaluation

We welcome independent reproduction of these results. The full deployment package, benchmark scripts, and reproduction instructions are available on request.

CREDASIS AI

DELAWARE C-CORPORATION · FOUNDED 2025 · PATENT PENDING

marty@credasis.ai · credasis.ai