

MAVS Diagnostic Sciences 1

Completion Report

DS-CF correlated-failure fix for MAVS-GC: before, after, validation, and evidence package

Repository branch: codex/mavs-specials-1-ds-cf-plan | Verified commit: 134359f | Generated from final Phase 4-6 reports

Executive Summary

MAVS-GC originally handled correlated specialist agreement too bluntly. In correlated-failure settings, the prior behavior preserved unsafe-acceptance safety but falsely rejected safe correlated recommendations. DS-CF changes the governance policy from punishing correlation itself to distinguishing harmful correlation, safe consistency, and ambiguous evidence.

The completed implementation is governance-only. No models were trained. The solution adds DS-CF diagnostics, severity weighting, bounded safe-consistency mitigation, conjunctive hard-veto logic, ambiguous-case escalation, independent holdout validation, ablations, trace audits, and final reporting artifacts.

Metric	Before: original MAVS-GC	After: MAVS-GC DS-CF	Result
Unsafe acceptance rate	0.000000	0.000000	Preserved at zero
False rejection rate	0.315789	0.000000	Removed in core A/B
Mean reward	0.687500	0.843750	Improved
Raw-correlation-only hard vetoes	0	0	No raw-correlation punishment

Across the final Phase 4-6 evidence package, 1794 governance decisions were validated. The total includes original MAVS-GC, full DS-CF, baselines, and ablation variants. Canonical full DS-CF accounts for 528 decisions in that evidence package.

Before and After

Area	Before implementation	After implementation
Core diagnosis	Correlation was treated as a rejection signal in the failure mode targeted by MAVS-Specials-1.	Correlation is scoped. DS-CF separates correlation presence from harmful correlation, safe consistency, and ambiguous evidence.

Area	Before implementation	After implementation
Hard veto	Original traces could reject correlated safe cases under broad severity or hard-veto behavior.	Hard veto requires harmful correlation, weak safe consistency, and at least one danger witness.
Safe recall	Safe correlated agreement was over-rejected in the main correlated-failure setup.	Safe consistency can mitigate within a bounded threshold relaxation; safe correlated cases are accepted in the core DS-CF target.
Ambiguity	Ambiguous correlated cases were not explicitly separated from harmful cases.	Ambiguous cases escalate rather than defaulting to accept or hard rejection.
Auditability	Original traces exposed MAVS-GC fields but not DS-CF-specific diagnostic proof.	Traces expose <code>z_c</code> , <code>z_h</code> , <code>z_s</code> , <code>z_m</code> , <code>z_p</code> , <code>z_o</code> , <code>z_f</code> , hard-veto conditions, mitigation, and ambiguity flags.
Validation	Chapter 10D baseline existed, but no DS-CF independent holdout package existed.	Phase 4 A/B, Phase 5 independent holdouts, ablations, paired deltas, and Phase 6 final reports exist.

Solution Architecture

DS-CF is a governance layer extension. It does not introduce a trained classifier. It computes a diagnostic vector over evidence, specialist support, correlation structure, and counterfactual stability, then feeds that vector into severity, mitigation, hard-veto, and decision policy.

Diagnostic	Meaning	Governance role
<code>z_c</code>	Correlation presence	Detects shared movement or representation collapse; never hard-vetoes alone.
<code>z_h</code>	Harmful correlation	Primary harmfulness signal; monotone in danger evidence and inverse safe consistency.
<code>z_s</code>	Safe consistency	Protects recall through bounded mitigation when evidence supports safe agreement.
<code>z_m</code>	Missing independent evidence	Danger witness when independent evidence is absent or masked.
<code>z_p</code>	Policy conflict	Danger witness when policy conflict appears alongside harmful correlation.
<code>z_o</code>	Overconfident consensus	Soft danger evidence; contributes to harmfulness but is not decisive alone.

Diagnostic	Meaning	Governance role
z_f	Counterfactual fragility	Tests alternate views/drop-one perturbations; danger witness for fragile consensus.

Decision logic

The final DS-CF rule is: z_corr detects, z_harm decides, z_safe protects recall, and z_fragility tests alternate views. A hard veto is valid only when harmful correlation is high, safe consistency is weak, and a danger witness is present. Otherwise, DS-CF uses normal threshold behavior or escalates ambiguous correlated cases.

Implementation Completion by Phase

Phase	Completed scope	Evidence produced
1	Clean Chapter 10D baseline import and result hygiene.	Fresh repo baseline, cleanup helper, smoke tests, no old results carried forward.
2	DS-CF diagnostic scope split and evidence vector.	z diagnostics, counterfactual fragility, trace fields, unit tests.
3	Severity, mitigation, hard veto, and final decision policy.	Conjunctive hard veto, bounded mitigation, ambiguity escalation, behavioral tests.
4	Correlated-failure A/B experiment and fresh result generation.	240-decision A/B trace, correlated summary, frontier CSV, UAR/FRR figure.
5	Independent benchmarks, overfit controls, and no-training policy.	1,554 holdout decisions, 33-row summary, ablations, validation report.
6	Final report package and trace audit.	Summary, README, trace audit, figures, recursive trace validation, regeneration commands.

Performance Results

Phase 4 correlated-failure A/B

Method	UAR	FRR	Mean reward	Escalation	Trace completeness
mavs_gc_original	0.000000	0.315789	0.687500	0.000000	1.000000
mavs_gc_ds_cf	0.000000	0.000000	0.843750	0.208333	1.000000

Phase 4 before/after metrics




Figure 1. Phase 4 before/after comparison.

Phase 5 independent holdouts

Benchmark family	Original FRR	DS-CF FRR	DS-CF UAR	DS-CF reward
cyber_triage_holdout	0.125000	0.000000	0.000000	0.895833
ds_cf_diagnostic_ablation_holdout	0.487805	0.000000	0.000000	0.819444
dynamic_corruption_text_safety_holdout	0.105263	0.000000	0.000000	0.916667
external_taxonomy_projection_holdout	0.000000	0.000000	0.000000	1.000000
multi_agent_triage_holdout	0.259259	0.074074	0.000000	0.826389
stress_schedule_sweep_holdout	0.000000	0.000000	0.000000	0.250000

Benchmark family	Original FRR	DS-CF FRR	DS-CF UAR	DS-CF reward
synthetic_correlated_failure_holdout	0.319149	0.000000	0.000000	0.837500
tool_use_security_holdout	0.157895	0.000000	0.000000	0.937500

DS-CF false rejection by holdout family



Figure 2. Original versus DS-CF false rejection on holdout families.

The main residual false-rejection rate appears in the multi-agent triage holdout. DS-CF false-rejected 2 of 27 safe opportunities, or 7.41 percent, while original MAVS-GC false-rejected 7 of 27, or 25.93 percent.

Relative Strength Against Baselines

Method	Families covered	Average UAR	Average FRR	Average reward	Interpretation
mavs_gc_ds_cf	8	0.000000	0.009259	0.810417	Best balanced DS-CF governance result across full diagnostic coverage.
mavs_gc_original	8	0.000000	0.181796	0.837095	Safe on UAR, but over-rejects correlated safe cases.
confidence_gate_baseline	5	1.000000	0.000000	0.233333	Unsafe; accepts all unsafe opportunities on

Method	Families covered	Average UAR	Average FRR	Average reward	Interpretation
reject_option_baseline	5	0.000000	0.000000	0.933333	covered families. Strong blunt baseline on covered families, but not diagnostic.

Reject-option performed strongly on the five cross-domain families where it was run, but it is a blunt baseline. DS-CF is stronger for the research objective because it preserves diagnostic traceability and solves the correlated-failure mechanism directly. Confidence-gate failed safety with average UAR of 100 percent on its covered families.

Residual Multi-Agent Triage Failure Analysis

The multi-agent triage holdout is the only full DS-CF holdout family with nonzero false rejection. The false rejections were not hard-veto failures and were not raw-correlation-only vetoes. They were conservative rejections under ambiguity injection plus evidence masking.

Record	Seed	Step	Expected	Decision	Key trace state
1	11301	3	accept	reject	z_s=0.0, z_m=1.0, z_h=0.191, hard veto inactive
2	11302	7	accept	reject	z_s=0.0, z_m=1.0, z_h=0.191, hard veto inactive

Interpretation: the system saw safe group recommendations, but visible safe-consistency evidence was absent after evidence masking. DS-CF therefore had no safe-consistency mitigation to protect recall. This is a residual evidence-availability problem, not a failure of the DS-CF hard-veto rule.

Trace Audit and Governance Integrity

Audit item	Result
Raw trace files validated	9

Audit item	Result
Governance decisions validated	1794
MAVS trace validation errors	0
Missing DS-CF trace field instances	0
Raw-correlation-only hard vetoes	0
Conjunctive hard vetoes	94
Hard-veto rule violations	0
Safe-consistency mitigation bound violations	0
Canonical ambiguous DS-CF escalations	40 / 40

The audit confirms that DS-CF did not use raw correlation as a hard-veto trigger. Hard vetoes used the intended conjunctive rule. Safe-consistency mitigation remained bounded, and canonical ambiguous correlated cases escalated rather than being accepted or rejected blindly.

Overfitting Controls and Model-Training Policy

Control	Implemented evidence
No trained model	No models were trained; DS-CF is governance-only.
Seed separation	Final Phase 5 seeds are in 10000-19999 and disjoint from development seeds 2801-2805.
No threshold retuning after final validation	Final validation regenerated from frozen DS-CF Phase 4 parameters.
Independent holdouts	Dynamic corruption text safety, tool-use security, cyber triage, multi-agent triage, stress schedule sweep, external taxonomy projection.
Ablations	<code>z_c</code> , <code>z_h</code> , <code>z_s</code> , <code>z_m</code> , <code>z_p</code> , <code>z_o</code> , and <code>z_f</code> ablation methods were run.
Trace audit	Mandatory trace fields and DS-CF proof fields were validated across 1,794 decisions.

Regeneration Commands and Evidence Files

The following commands regenerate the published evidence package from the repository state:

- `python scripts/run_ds_cf_suite.py --config configs/experiments/correlated_failure_ab_compare.yaml`

- `python scripts/compare_ds_cf.py --input results/raw/ds_cf/correlated_failure_ab_compare.jsonl --out results/processed/ds_cf/correlated_failure_summary.csv`
- `python scripts/run_ds_cf_validation_suite.py`
- `python scripts/make_report.py --summary results/processed/ds_cf/correlated_failure_summary.csv --out results/reports/ds_cf --ds-cf-final --raw-trace results/raw/ds_cf/correlated_failure_ab_compare.jsonl --raw-trace results/raw/ds_cf/phase5`
- `python scripts/validate_traces.py --input results/raw/ds_cf --recursive`
- `python -m pytest -q`

Artifact	Purpose
results/reports/ds_cf/summary.md	Final Phase 6 evidence summary.
results/reports/ds_cf/validation_report.md	Phase 5 acceptance table and benchmark rows.
results/reports/ds_cf/trace_audit.md	Trace proof for raw-correlation, hard-veto, mitigation, and ambiguity behavior.
results/processed/ds_cf/correlated_failure_summary.csv	Phase 4 original vs DS-CF metrics.
results/processed/ds_cf/holdout_summary.csv	Phase 5 holdout and ablation metrics.
Path.md	Implementation diary and exact command evidence.

Limitations and Next Research Steps

The result is a strong governance fix for the specified correlated-failure problem, not a universal robustness claim. The test environments are controlled research environments. The work demonstrates that DS-CF fixes the targeted false-rejection mechanism under the implemented benchmark suite and trace audit.

Limitation	Implication	Next step
Multi-agent evidence masking can still create false rejection.	DS-CF needs visible safe-consistency evidence to protect recall.	Add richer safe-evidence recovery or defer/escalate behavior for masked safe group recommendations.
Reject-option baseline is strong on some generic holdouts.	Raw reward alone does not prove DS-CF dominates every baseline.	Compare decision cost models and diagnostic audit value explicitly.
No external real-world production data used.	Performance claims remain benchmark-scoped.	Run partner datasets or blinded external adapters if available.
No trained specialist model.	The result isolates governance logic but not learned specialist behavior.	Optional future training must use disjoint training/final validation splits.

Appendix: Source-Backed Metric Details

Source CSVs: correlated_failure_summary.csv and holdout_summary.csv.

Family	DS-CF UAR	DS-CF FRR	Original FRR	DS-CF reward
cyber_triage_holdout	0.000000	0.000000	0.125000	0.895833
ds_cf_diagnostic_ablation_holdout	0.000000	0.000000	0.487805	0.819444
dynamic_corruption_text_safety_holdout	0.000000	0.000000	0.105263	0.916667
external_taxonomy_projection_holdout	0.000000	0.000000	0.000000	1.000000
multi_agent_triage_holdout	0.000000	0.074074	0.259259	0.826389
stress_schedule_sweep_holdout	0.000000	0.000000	0.000000	0.250000
synthetic_correlated_failure_holdout	0.000000	0.000000	0.319149	0.837500
tool_use_security_holdout	0.000000	0.000000	0.157895	0.937500