



U.S. General Services
Administration

SECURITY CONSIDERATIONS FOR MULTI-AGENT AI SYSTEMS

Federal Cybersecurity and Privacy Professionals Forum by NIST
Presented by Tam Nguyen - GSA / Office of Government-wide Policy
09/01/2026

AGENDA ITEMS



The Risk Landscape

Risk frameworks and
their coverage of
Agentic AI risks



The Gaps

What Agentic AI Risks
are not being covered
and why such gaps
exist



Our Method

How did we identify the
gaps



Our Recommendations

What practitioners,
framework developers,
and policy makers get
from our results

Tam Nguyen

Data Scientist for Cybersecurity
GSA - Office of Government-wide Policy
Tam.Nguyen@gsa.gov
Orcid: 0000-0002-8577-8342
Linkedin: tamcs




Dheeraj Arremsetty

Advisory Board Member @ Crew Scaler
AI Solution Architect @ IBM
dheeraj.arremsetty@crewscale.org



Moses Ndebugre

Senior Researcher @ Crew Scaler
PhD @ North Carolina A&T State University
m@crewscale.org




THE TEAM



This work is a collaboration among these experts and our AI Agent swarms. The results are being peer-reviewed by the Association for Computing Machinery (ACM).

<https://arxiv.org/pdf/2603.09002>


GENERATIVE AI RISKS



SHIFTING TO MULTIPLE-AGENTS




EXAMPLE OF CASCADING FAILURES




MAS RISK TAXONOMY

TIER 1: EXTENDED GENAI RISKS



Expands upon Standard GenAI Frameworks

TIER 2: NOVEL MULTI-AGENT RISKS



TIER1 DOMAIN COVERAGE

Frameworks ▾	Output & Content ▾	Security & Adversaria ▾	Agentic & Autonomou ▾	Data & Privacy ▾	Supply Chain & Infrastructur ▾	Governance & Accountabilit ▾	Operational Reliability ▾	Societal & Ethical ▾
ATFAA	✓	✓	✓	✓		✓		
Cisco_A2A		✓	✓		✓			
Google		✓	✓	✓		✓		
NSA_Data		✓		✓	✓			
NSA_Deploy	✓	✓		✓	✓		✓	✓
NSA_Dev		✓	✓	✓	✓	✓	✓	
DoD_Cyber		✓		✓	✓	✓	✓	
DoD_RAI	✓			✓	✓	✓	✓	✓
NIST_RMF	✓	✓	✓	✓	✓	✓	✓	✓
NIST_GenAI	✓	✓		✓	✓	✓		✓
OWASP_MAS		✓	✓	✓				
OWASP_Top10	✓	✓	✓		✓			
OWASP_T&M	✓	✓	✓	✓	✓	✓		
OWASP_Secure			✓		✓			
OWASP_State	✓	✓	✓			✓		✓
ENISA	✓	✓		✓	✓	✓	✓	✓
GAO				✓		✓	✓	
MITRE	✓	✓	✓		✓		✓	✓
CDAO	✓	✓	✓	✓	✓	✓	✓	✓

TIER2 DOMAINS WITH NO COVERAGE




The gap is not incidental - it reflects design intent.
Every framework centers on a single model, or a single agent.


https://github.com/Crew-Scaler/multi_agent_system_risks/

MEMORY & COGNITIVE STATE

What makes these risks distinctively multi-agent is that multi-agent memory is rarely private.

Agent cognitive state is written by one agent and read by another. A single poisoned memory can propagate laterally across collaborators






IDENTITY AND TRUST


Identity, trust, and provenance become emergent attributes negotiated across many cooperating agents.

A single compromise propagates through delegation chains, consensus mechanisms, and shared infrastructure can quickly grow into a system-wide failure of accountability.


11.1 CONTINUAL CHANGE (Assurance Invalidation)




11.2 DECISION-MAKING (Exploitable Windows)




11.3 REASONING & PLANNING (Divergent Outcomes)




11.4 INFRASTRUCTURE (Behavioral Variability)



11.5 MEMORY SYSTEMS (Assurance Gaps)



11.6 POLICY & LEARNING (Online Adaptation)



NON-DETERMINISM & ASSURANCE


Variance compounds across agent boundaries, tool invocations, and orchestration steps. A passed test yesterday may not hold today.

TELEMETRY & OBSERVABILITY


Observability infrastructure designed for single agent routinely fails at the seams between agents

Events may cross system boundaries where no single trace, log, or metric retains end-to-end causality. The result is systematic blind spots.


1 INTERFACE & UI TELEMETRY GAPS.



2 FRAMEWORK & REASONING MONITORING GAPS.




3 MEMORY & KB POISONING.



4 DETECTION EVASION & METRIC MANIPULATION.



5 EVALUATION & MULTI-AGENT PERFORMANCE TELEMETRY GAPS.




1

PLUGIN & TOOL SUPPLY CHAIN

Metadata
Injection

Permission
creep

Registry
tampering




2

RAG PIPELINE & KNOWLEDGE BASE



3

ORCHESTRATION PLANE compromiSE




4

EVALUATION, GUARDRAIL, & MONITORING




5

DISTRIBUTED SYSTEMS & INFRASTRUCTURE



6

LEARNING- BASED WORKFLOWS




WORKFLOWS & ECOSYSTEMS


Shared infrastructure such as a single plugin, retrieval index, orchestration controller, or fine-tuning corpus is consumed by an entire fleet, so compromising one shared component propagates to every agent that depends on it.

EMERGENT & SPECIFICATION GAMING


1. SPECIFICATION GAMING VIA ENVIRONMENT MANIPULATION (Cat 14.1)



2. LONG-HORIZON GOAL DRIFT AND MEMORY-BASED MISALIGNMENT (Cat 14.2)




3. EMERGENT MISALIGNMENT THROUGH MULTI-AGENT INTERACTION (Cat 14.3)




Gaming behaviors no longer remain local. Agents each acting within their own specification can nonetheless converge on collectively unsafe outcomes that no single component would exhibit in isolation.

OUR METHODOLOGY


PHASE 1: KNOWLEDGE DOMAIN EXPLORING (KDEG)



PHASE 1: KNOWLEDGE HIERARCHY BUILDING



OUR METHODOLOGY



FOR PRACTITIONERS

Deploy Cross-Agent Telemetry Infrastructure

correlates events across heterogeneous agent ecosystems

Implement Cryptographic Agent Identity


Signing of agent messages, capability attestation, provenance tracking

Sequence-Aware Authorization

for Agent Workflows

Continuous, Event-Driven Security Testing

Point-in-time testing cannot characterize a multi-agent system



FOR FRAMEWORK DEV

Explicit MAS threat modeling

require risk assessments to characterize inter-agent potential issues

Emergent-behavior evaluation


Define methods to manage risks associated with emergent behavior.

Observability requirements

Specify minimum MAS telemetry

Sequence-aware authorization

Define and require authorization at the workflow level, not only per action.



FOR POLICY MAKERS

Require MAS-specific risk assessments

Establish MAS telemetry standards


Define minimum observability for regulated MAS


Fund framework guidance for D9 to D14

Direct standards bodies to develop MAS-specific extensions for D9 to D14

Recognize the multi-agent security boundary

distinct regulatory treatment for agent and agent swarms





WHERE TO GO NEXT

We've just scratched the surface. Multi-agent economy can reach Metaverse level. Associated risks will be massive.



Scalable Method for Risk Identification

How to identify MAS risks at scale and as early as possible?



Emergent MAS Behaviors

How can we advance our understanding of multi-agent emergent behaviors?



Transferring Knowledge

How can we rapidly transfer the knowledge to as many practitioners as possible?

Thank you!