Persistent

Practitioner Point Of View

From MFA to CFA

Bringing Trust to Age of AI Agents

In enterprise security, few concepts have been as foundational as Multi-Factor Authentication. MFA was elegant in its logic: Prove who you are through something you know, something you have and something you are. Identity rooted in human biology and possession. But something fundamental has shifted.

Dilip Panjwani
Cybersecurity Practice Head

Multi-Factor Authentication was designed for humans. As autonomous agents flood enterprise environments, we need Context Factor Authentication, a new trust paradigm built for machines that act on our behalf.

Across boardrooms and data centers, AI agents are being handed the keys. They are booking meetings, querying databases, writing code, triggering payments and making decisions — autonomously, at machine speed, across distributed cloud environments. They do not have passwords. They do not carry phones. They do not have fingerprints. Yet they need to be trusted.

For two decades, we secured the human in the loop. The loop has changed. Today, the agent IS the loop — and we have no established doctrine for trusting it.

This is the defining identity challenge of the agentic AI era: How do you authenticate an entity that isn't human?

Problem with Applying Human Trust Logic to Machines

MFA, at its core, is a point-in-time verification mechanism. At the moment of login, we challenge the human: Prove three things, get a session token and proceed. The session carries the trust forward. It's a handshake at the door.

AI agents break every assumption baked into this model.

Three Broken Assumptions

Assumption 1

Identity is singular

An agent isn't one entity, it's a chain. A prompt hits an orchestrator, which spawns sub-agents, which call tools, which invoke APIs. Which link in that chain do you authenticate? All of them? None authentically?

Assumption 2

Sessions are bounded

Human sessions expire. Agent workflows run for hours, branch into parallel tasks and hand off context across systems continuously. The session model does not map.

Assumption 3

Authentication is the gate, not the journey

With humans, we verify entry. With agents, the risk isn't at the door; it's inside the house, over the duration of a long, complex, multi-step task.

I've spent nearly two decades watching enterprise security evolve from perimeter firewalls to zero trust, from static passwords to biometric MFA, from on-premise SOCs to cloud-native SIEM. Each shift forced a rethink of what 'trust' means and where it lives. The agentic era demands the same rethink, but faster and at higher stakes.

Introducing Context Factor Authentication (CFA)

What if, instead of challenging an agent to prove identity at a fixed moment, we continuously evaluate the full context of its operations, across the who, what, where, why and how of every action it takes?

That's the premise of Context Factor Authentication (CFA), a trust framework built not for login events but for autonomous action sequences. CFA does not ask 'Are you who you say you are?' It asks the harder, more operationally relevant question: Does everything you're doing right now make sense, given everything we know about your sanctioned purpose?

MFA vs. CFA — Structural Comparison

MFA — Legacy Paradigm

Point-in-time verification
Designed for a human logging into a system. Grants a time-bounded session. Binary outcome: Access or deny. Identity anchored to human attributes.

Trust Model: Static after authentication
Once the handshake succeeds, the session carries forward. Privilege is granted at entry.

Revocation: Session timeout or manual kill
Passive expiry or reactive termination after breach is detected.

CFA — Agentic Paradigm

Continuous, multi-dimensional trust evaluation
Designed for an autonomous agent executing a task chain. Trust is dynamic and revocable at any step. Identity is contextual, defined by purpose, lineage, behavior, environment and intent.

Trust Model: Dynamic, action-by-action
Trust is a confidence score, not a binary state. Recalculated at each decision point in the agent's workflow.

Revocation: Autonomous, mid-task suspension
If the context score falls below threshold, the agent is paused, quarantined or escalated to human review before damage occurs.

Six Context Factors

CFA does not replace existing identity infrastructure; it augments it with a richer signal set purpose-built for non-human actors.

Here are six core context factors that together constitute an agent's trust profile at any given moment.

Agent Lineage

Who spawned this agent? Through what orchestration chain? Verified provenance of the agent's origin and delegation path.

Sanctioned Scope

What was this agent explicitly authorized to do? Actions outside the declared task boundary trigger trust degradation.

Operational Environment

Which network, cloud tenant, API endpoint and data classification tier is the agent operating in? Does it match its declared purpose?

Temporal Coherence

Is the agent acting within expected time windows? Urgency signals or off-hours operations outside norms are anomaly indicators.

Behavioral Baseline

How does this agent's current action pattern compare to its established operating baseline? Drift is a red flag.

Intent Alignment

Can the agent's observed actions be traced back to a legitimate, documented business intent? Orphaned actions have no trust anchor.

These six factors combine into what can be called the Agent Trust Score — a dynamic, continuously recalculated confidence rating that governs whether an agent may proceed, must pause for human review or must be suspended outright.

We do not need to trust the agent. We need to trust the context. The agent is only as trustworthy as the sum of its verifiable surroundings.

Why This Isn't Just a Technical Problem

Leaders who frame agent security as purely an IT or DevSecOps concern are making a governance mistake. The question of agent trust is a board-level risk question because the consequences of getting it wrong land at the board level.

Consider what happens when a compromised or misconfigured agent operates undetected inside an enterprise. It does not trigger a failed login alert. It does not set off a traditional SIEM rule. It behaves exactly like a legitimate workflow because it IS a legitimate workflow, hijacked. The attack surface is not the authentication layer. It is the business logic layer.

This is prompt injection at enterprise scale. It is supply chain compromise via AI orchestration. It is privilege escalation disguised as automation. And it is happening right now, in organizations that have deployed agentic workflows without asking: What is our trust doctrine for these entities?

Agent Attack Vectors and CFA Mitigations

Attack Vector	How It Exploits MFA Gap	CFA Mitigation	Risk
Prompt Injection	Malicious instructions embedded in data the agent ingests redirect its actions outside sanctioned scope.	Sanctioned Scope + Intent Alignment factors detect deviation from declared task.	**High**
Agent Impersonation	A rogue agent presents valid API credentials harvested from a legitimate orchestration chain.	Agent Lineage factor validates full provenance chain, not just endpoint credentials.	**High**
Privilege Drift	Agent accumulates permissions across task iterations beyond its original authorization.	Sanctioned Scope + Behavioral Baseline detect creeping privilege over time.	**High**
Shadow Workflows	Undeclared agent pipelines operate outside governance visibility.	Operational Environment factor flags agents in unregistered infrastructure.	**Medium**
After-Hours Ex-filtration	Compromised agent initiates data access during off-peak windows to avoid detection.	Temporal Coherence factor flags anomalous timing patterns.	**Medium**

Agent Trust Architecture

Implementing CFA is not about buying a new product. It is about building a new capability layer into how an organization deploys and governs AI agents. From an architectural standpoint, five elements are foundational:

Agent Identity Registry

Every AI agent in the enterprise must have a declared identity — a unique, verifiable identifier that records its sanctioned purpose, authorized data access tiers, allowed API endpoints and parent orchestration system. Think of it as a passport for agents: issued, scoped and auditable.

Context Policy Engine

A real-time evaluation layer — analogous to a Zero-Trust policy engine — that scores each agent action against the six CFA dimensions. This sits between the agent and the resources it seeks to access, functioning as a continuous, inline trust broker rather than a perimeter gate.

Immutable Action Ledger

Every action taken by every agent must be cryptographically logged in an append-only audit trail. When things go wrong — and they will — forensic reconstruction of agent behavior across multi-step workflows is non-negotiable. This is the incident response foundation.

Human-in-the-Loop Escalation Triggers

CFA is not about removing humans from the equation. It is about inserting them at the right moments. Configurable trust-score thresholds should trigger mandatory human review before high-consequence actions, such as data deletion, financial transactions or privilege escalation, before proceeding.

Agent Behavior Analytics (ABA)

Just as UEBA established baselines for human behavior, we need ABA — Machine Learning models trained on legitimate agent operation patterns, continuously monitoring for anomalies that CFA's deterministic rules might miss.

The CISO'S Immediate Action List

Audit now

What cannot be seen, cannot be governed. Catalogue every AI agent or agentic workflow deployed in an environment. Shadow AI deployments are your highest risk.

Demand Agent Cards

For every agent in deployment or procurement, require vendors to produce an Agent Identity Card: Declared scope, data access, API permissions, human oversight model.

Extend SIEM and SOAR for agent telemetry

The existing detection stack almost certainly is not instrumented for agent-originated events. This is a gap enterprises need to close before the first agent-enabled incident.

Define trust thresholds

Work with business stakeholders to define which agent actions are auto-approved, which require human confirmation and which are categorically prohibited — regardless of context score.

What Standards Bodies Are Not Telling You

The National Institute of Standards and Technology (NIST), the International Organization for Standardization (ISO) and major regulatory frameworks are beginning to acknowledge the existence of AI agents.

NIST's AI Risk Management Framework, the EU AI Act and emerging guidance from financial regulators such as the Reserve Bank of India (RBI) and the Financial Conduct Authority (FCA) all gesture at accountability and auditability for AI systems. But none of them have yet produced workable operational guidance on agent authentication specifically.

This is the gap that practitioners must fill through doctrine, not wait for regulators to close it through mandate. The organizations that develop robust internal CFA frameworks today will be ahead of both the threat and compliance curves when formal standards arrive — and they will arrive.

For CISOs advising boards: Frame this as a first-mover governance opportunity. The organizations that establish agent trust doctrine proactively become the benchmark against which regulators will draft the eventual standards. That is strategic positioning, not just risk management.

Bigger Picture: Trust as Infrastructure

The agent is not the user. It does not log in. It acts. And every action it takes is either sanctioned or suspect. Our job as security leaders is to ensure we always know which one.

This is an inflection point in enterprise computing. The shift from human-operated workflows to agent-operated workflows is as significant as the shift from on-premises to the cloud. And just as cloud forced us to rebuild security from the ground up, abandoning perimeter thinking in favor of an identity-centric, zero-trust architecture, the agent era forces a similar reconstruction of what trust means.

MFA was a brilliant solution to the problem of human authentication. It moved us from single-factor fragility to multi-layered resilience. CFA must do the same for the agentic layer — moving us from point-in-time credential checks to continuous, context-aware trust evaluation.

The question for every CISO today is not whether to deploy AI agents. That decision is already being made — by the business units, vendors and competitors. The question is whether enterprises will build the trust infrastructure before or after the first major agent-enabled breach forces them to.

The Doctrine We Need, Before the Incident We Don't

Context Factor Authentication is not a product category yet. It is a security philosophy that the practitioner community needs to develop, debate and operationalize — before regulators mandate it and attackers exploit the gap. If you are building agentic workflows or advising organizations that are, the trust question cannot wait for a vendor to answer it. It starts with a doctrine. It starts now.

About the Author

Dilip is a seasoned cybersecurity leader with over two decades of experience serving as CISO across major BFSI institutions, fintech organizations and global IT enterprises.

Across these roles, he has been responsible for securing high value financial systems, regulatory driven environments and large scale digital transformation programs. His work spans enterprise security architecture, threat management, data protection, cloud security and building cyber resilience for organizations operating in some of the most demanding risk landscapes.

Today, Dilip leads the cybersecurity practice and delivery function at Persistent Systems, where he partners with CIOs, CISOs, CTOs and digital leaders to elevate their cybersecurity maturity, strengthen governance and build future ready security programs. His unique blend of practitioner depth and advisory experience gives him a 360 degree view of the challenges enterprises face as they navigate emerging threats like cyber risk management, quantum risk, AI driven attacks and regulatory complexity.

Dilip is a trusted voice in the cybersecurity community and a leader who continues to shape how organizations think about security, resilience and digital trust.

About Persistent

Persistent Systems (BSE: 533179 and NSE: PERSISTENT) is a global services and solutions company delivering AI-led, platform-driven Digital Engineering and Enterprise Modernization to businesses across industries. With over 27,500 employees located in 21 countries, the Company is committed to innovation and client success. Persistent offers a comprehensive suite of services, including software engineering, product development, data and analytics, CX transformation, cloud computing, and intelligent automation. The Company is part of the MSCI India Index and is included in key indices of the National Stock Exchange of India, including the Nifty Midcap 50, Nifty IT, and Nifty MidCap Liquid 15, as well as several on the BSE such as the S&P BSE 100 and S&P BSE SENSEX Next 50. Persistent is also a constituent of the Dow Jones Best-in-Class World Index. The Company has achieved carbon neutrality, reinforcing its commitment to sustainability and responsible business practices. Persistent has also been named one of America's Greatest Workplaces for Inclusion & Diversity 2025 by Newsweek and Plant A Insights Group. As a participant of the United Nations Global Compact, the Company is committed to aligning strategies and operations with universal principles on human rights, labor, environment, and anti-corruption, as well as take actions that advance societal goals. With 468% growth in brand value since 2020, Persistent is the fastest-growing IT services brand in 'Brand Finance India 100' 2025 Report.

USA
Persistent Systems, Inc.
1731 Technology Drive Suite 700
San Jose, CA 95110
Tel: +1 (408) 216 7010
Fax: +1 (408) 451 9177
Email: info@persistent.com

India
Persistent Systems Limited
Bhageerath, 402
Senapati Bapat Road
Pune 411016
Tel: +91 (20) 6703 0000
Fax: +91 (20) 6703 0008

Persistent
Re(AI)magining the World