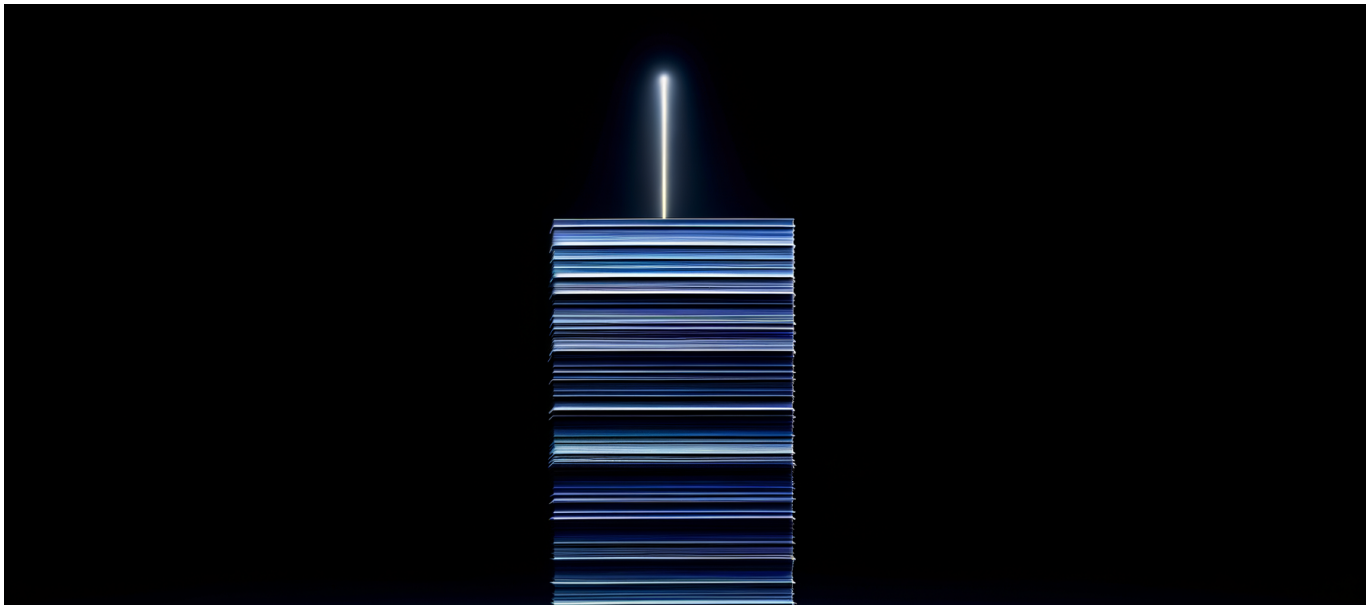

Rocket Software Launches EVA AI Assistant on January 27—Traces Mainframe Issues From Symptoms to Code Lines Using Model Context Protocol

The AI industry spent 2025 building chatbots for email summarization while ignoring the infrastructure processing $3 trillion in daily transactions. Rocket Software just shipped an AI assistant that debugs COBOL running your mortgage payments.

The News: Rocket EVA Brings Agentic AI to Legacy Infrastructure

Rocket Software launched Rocket EVA on January 27, 2026—an enterprise virtual assistant designed for automated operational diagnostics across mainframe, distributed, and cloud environments. The announcement marks the first production

deployment of the Model Context Protocol (MCP) standard for mainframe operations at enterprise scale.

EVA's core capability: tracing operational issues from initial symptoms through system interactions to the exact line of responsible code. Not approximate guidance. Not suggested areas to investigate. The specific line number where the problem originates.

The assistant integrates with Rocket's existing data scanners and third-party MCP servers, creating an interoperable framework for governed system data access. Natural-language queries replace the tribal knowledge currently required to diagnose MIPS efficiency problems, workload placement decisions, and code refactoring opportunities.

SD Times covered the launch as part of January's AI updates, positioning it alongside cloud-native AI announcements—but missing the significance of what Rocket actually built. This isn't another wrapper around GPT-4. It's an agentic system with governed access to infrastructure that predates the internet.

Why This Matters: The $350B Infrastructure Layer Everyone Ignores

Mainframes process 70% of Fortune 500 transactions. They run the SWIFT network. They calculate your credit score. They process 30 billion ATM transactions annually. And until January 27, they had no AI-native diagnostic tooling.

The economics tell the story. A single MIPS (Million Instructions Per Second) costs between $1,500 and $3,000 annually on IBM Z-series hardware. Large enterprises run 50,000 to 100,000 MIPS. When a batch job runs inefficiently—consuming 15% more MIPS than necessary—the monthly bill increases by hundreds of thousands of dollars. The expertise to diagnose these inefficiencies is retiring at 10,000 professionals per year.

EVA targets the gap between institutional knowledge walking out the door and the systems that still need human understanding to operate.

The second-order effects hit three constituencies immediately:

Winners: Enterprises with Legacy Technical Debt

Banks, insurers, and government agencies running COBOL applications written in 1987 now have a path to understanding code that predates their current engineering staff's careers. EVA doesn't require migrating off mainframes—a process that has bankrupted modernization budgets for two decades. It makes existing infrastructure more observable.

Losers: Boutique Mainframe Consulting Firms

The $15 billion mainframe services market relies on scarcity. Fewer than 200,000 professionals worldwide understand z/OS internals at a diagnostic level. When an AI can trace symptoms to code lines, the premium for human expertise compresses. Consulting engagements that previously required two-week investigations become one-day validations of EVA's findings.

Changed: The AI Infrastructure Stack

Model Context Protocol matters more than EVA itself. Anthropic released MCP in late 2024 as a standard for connecting AI systems to data sources. Rocket's implementation proves MCP works for the most complex, most regulated, least cloud-native infrastructure in enterprise computing. If MCP handles mainframes, it handles everything.

Technical Depth: How EVA Actually Works

Understanding EVA requires understanding what it replaces. Traditional mainframe diagnostics follow a human-intensive pattern: an operator notices a batch job running long, opens SMF (System Management Facilities) records, correlates timestamps with resource consumption, identifies the LPAR (Logical Partition) showing anomalies, traces to the specific job step, then reads the responsible COBOL paragraph to find the inefficiency.

This process requires expertise in z/OS, JCL, SMF record formats, COBOL structure, and whatever application logic the code implements. It takes hours. It depends on documentation that may not exist. It fails when the expert who wrote the code retired in 2019.

The MCP Architecture

EVA uses Rocket MCP Server as its governed access layer. IDC's analysis of EVA emphasizes the governance model: MCP doesn't give the AI unrestricted system access. It provides structured, auditable, policy-constrained data retrieval.

The architecture works in three layers:

Query Layer: Natural-language input parsed into diagnostic intent. "Why did the overnight settlement batch run 40% longer than Tuesday?" becomes a structured request for temporal comparison across resource metrics.

Context Layer: Rocket MCP Server retrieves relevant SMF records, job logs, system configurations, and application metadata. The retrieval is governed—the AI can't access data outside its authorization scope, and every query logs for audit.

Analysis Layer: EVA correlates the retrieved context, identifies anomalies, and traces causation chains. The output includes confidence scores for each diagnostic step, creating an explainable path from symptom to root cause.

The code-line-level tracing deserves specific attention. COBOL programs often exceed 50,000 lines. Traditional debugging requires reading sequentially, understanding data division structures, and mentally executing procedure division logic. EVA's approach treats code as contextual data—retrieving the specific paragraphs executed during anomalous time windows and correlating them with resource spikes.

Integration Architecture

EVA's interoperability model differentiates it from closed AI systems. The third-party MCP server support means enterprises can connect EVA to:

- ServiceNow for incident correlation
- Splunk for log aggregation
- Application-specific data sources via custom MCP implementations
- Cloud infrastructure APIs for hybrid workload analysis

This matters because mainframe problems rarely stay on mainframes. A batch job might call a distributed API, which queries a cloud database, which returns slowly

because of a misconfigured connection pool. EVA's multi-environment support traces these chains across infrastructure boundaries.

What The Benchmarks Don't Show

Rocket hasn't published diagnostic accuracy benchmarks—an understandable omission given that mainframe diagnostics lack standardized test sets. The closest proxy: MIPS efficiency analysis, where EVA reportedly identifies optimization opportunities that human analysis misses.

The benchmark gap matters for enterprise adoption. CTOs evaluating EVA need proof that AI-generated diagnostics won't create false confidence in incorrect conclusions. A misdiagnosed root cause wastes more time than no diagnosis at all.

Early deployments will likely use EVA for triage rather than resolution—identifying the investigation vector for human experts rather than replacing human judgment entirely. This is the appropriate trust calibration for v1.0 of any agentic system touching production infrastructure.

The Contrarian Take: What Most Coverage Gets Wrong

The press release narrative frames EVA as "conversational AI for mainframes." This undersells the actual achievement and oversells the interaction model.

Undersold: MCP's Governance Model

The real innovation is governed agentic access to legacy systems. Previous AI-for-mainframe attempts failed because they required either unrestricted system access (unacceptable for regulated industries) or pre-extracted data snapshots (too stale for real-time diagnostics).

MCP solves the access control problem that blocked enterprise AI adoption for five years. Every query is scoped, logged, and policy-constrained. Compliance teams can audit exactly what the AI accessed and why. This isn't a feature—it's the enabler of every other capability.

The tweetable version: EVA matters less than MCP. Rocket just proved

that agentic AI can meet banking regulatory requirements.

Oversold: The Conversational Interface

Natural-language interaction sounds transformative until you realize mainframe operators don't want to type sentences. They want dashboards, alerts, and CLI tools that integrate with existing workflows.

The conversational model serves a specific use case: executives and architects querying systems they don't operate daily. For production operators, EVA's value comes from API integration, automated alerting, and scripted invocation—not chat windows.

Rocket's marketing emphasizes the natural-language interface because it demonstrates capability intuitively. But production deployments will likely bury the chat interface beneath automation layers within six months.

Overhyped: Immediate Productivity Gains

Every AI announcement promises productivity improvements. EVA's actual productivity impact depends on factors Rocket can't control:

- Quality of existing documentation (EVA can only trace what's instrumented)
- Complexity of application interdependencies (some root causes span decades of accumulated logic)
- Willingness of operations teams to trust AI-generated diagnostics (cultural, not technical)

Organizations with well-instrumented, well-documented mainframe environments will see immediate value. Organizations with thirty years of technical debt will need six months of MCP configuration before EVA becomes useful.

Underhyped: The Retirement Crisis Mitigation

The average COBOL programmer is 58 years old. The number of professionals who understand z/OS internals decreases by approximately 10,000 annually through retirement. Enterprises face a fifteen-year window where institutional knowledge disappears faster than it can be transferred.

EVA doesn’t replace this expertise—but it captures its outputs. Every diagnostic session that traces a symptom to a root cause creates an auditable record of how that conclusion was reached. Over time, EVA’s context accumulates the investigative patterns that currently exist only in human memory.

This is the strategic play most coverage misses. Rocket isn’t selling a diagnostic tool. They’re selling institutional knowledge preservation.

Practical Implications: What You Should Actually Do

If You Run Mainframes

Immediate action: Inventory your MCP readiness. EVA’s value depends on what data sources you can connect. Start mapping SMF record availability, job scheduling metadata, and application-level instrumentation.

30-day action: Request an EVA pilot scoped to a single application domain—preferably one with recent performance issues and available historical data. Use the pilot to validate diagnostic accuracy against known root causes.

90-day action: If the pilot validates, develop MCP integration roadmaps for third-party data sources. EVA’s cross-environment tracing only works when all environments expose MCP-compatible interfaces.

If You Build AI Infrastructure

Immediate action: Study MCP implementation patterns. Rocket’s use case proves the protocol scales to enterprise requirements. Your roadmap should include MCP server development for any proprietary data sources you maintain.

Architecture consideration: Governed access layers belong between AI agents and production systems, regardless of infrastructure type. EVA’s model—structured queries, policy constraints, audit logging—applies to Kubernetes clusters as much as mainframes.

The pattern to replicate: Context retrieval with confidence scoring. EVA’s diagnostic outputs include explainability metadata that allows humans to validate AI

reasoning. This pattern should become standard for any AI system touching production infrastructure.

If You’re Evaluating Vendors

EVA’s launch changes the competitive landscape for mainframe tooling. IBM’s AI strategy for Z-series has focused on modernization—moving workloads off mainframes rather than making them more observable. Rocket’s bet: enterprises will run mainframes for another twenty years and need better operational tooling for the duration.

Questions to ask IBM: Where is MCP support in your Z-series AI roadmap? If Rocket ships governed agentic access before IBM, the modernization narrative weakens.

Questions to ask Rocket: What diagnostic accuracy can you demonstrate on my specific application stack? Benchmarks against generic workloads matter less than performance on sixty-year-old COBOL with minimal documentation.

Code and Architecture to Explore

MCP server implementation isn’t Rocket-specific. Anthropic’s MCP specification is open, and reference implementations exist for common data sources. If you’re building agentic systems:

- Review the Rocket EVA announcement for integration pattern details
- Evaluate MCP server development for your internal data sources
- Design governance layers that log queries, enforce scope constraints, and provide audit trails

The architectural principle: AI agents should never have direct database or system access. MCP-style intermediation provides the control plane that enterprise security teams require.

Forward Look: Where This Leads in Six to Twelve Months

Q2 2026: MCP Becomes Enterprise Standard

EVA's launch validates MCP for the most demanding enterprise use case imaginable. By April, expect announcements from ServiceNow, Datadog, and Splunk adding MCP server capabilities. The protocol becomes the standard interface between AI agents and enterprise data sources.

Implication: Teams building agentic systems should prioritize MCP over proprietary integration approaches. The winner-take-all dynamics of protocol standards mean early MCP adoption creates durable advantages.

Q3 2026: Competitive Response from IBM

IBM cannot allow Rocket to own the AI-for-mainframes narrative. Expect Watson-branded diagnostic tooling with MCP compatibility, likely announced at SHARE or Think 2026. IBM's response will emphasize deeper Z-series integration—hardware-level telemetry that Rocket can't access.

Implication: Enterprises evaluating EVA should negotiate contracts that account for competitive alternatives arriving within twelve months. Pricing leverage increases as options multiply.

Q4 2026: Agentic Diagnostics Expand Beyond Mainframes

EVA's architecture applies to any complex system with observable components and traceable causation chains. Expect Rocket or competitors to ship EVA-style diagnostics for:

- SAP environments (similarly complex, similarly dependent on tribal knowledge)
- Oracle database clusters (similar MIPS-style licensing creates optimization pressure)
- Kubernetes clusters (different technology, same observability challenges)

The pattern Rocket proved—governed agentic access, cross-system tracing, code-level root cause identification—generalizes beyond mainframes.

2027: Autonomous Remediation

EVA's current scope is diagnostic, not remedial. The logical extension: not just

identifying the line of code causing problems, but suggesting or implementing fixes. This requires additional governance layers—change approval workflows, rollback capabilities, blast radius containment.

Rocket's roadmap likely includes remediation capabilities, but they won't ship until enterprises trust EVA's diagnostic accuracy. Expect advisory recommendations (EVA suggests code changes) before autonomous remediation (EVA implements code changes).

The Infrastructure AI Gap Closes

For five years, AI investment concentrated on consumer-facing applications while enterprise infrastructure remained static. Chatbots improved; batch processing didn't. The gap reflected AI's training data—abundant internet text, scarce mainframe documentation.

Rocket EVA signals the gap closing. Not because mainframe documentation suddenly became abundant, but because MCP provides a governed path from AI models to system context. The training data limitation becomes irrelevant when AI can query live systems for the context it needs.

The implications extend beyond mainframes. Any infrastructure with observable state and traceable causation chains becomes amenable to EVA-style diagnostic AI. The question shifts from "can AI understand this system?" to "does this system expose MCP-compatible interfaces?"

Enterprises running legacy infrastructure—whether COBOL on z/OS or Perl on Solaris—now have a template for AI integration. Govern the access. Log the queries. Trace the causation. Surface the root cause.

Rocket EVA isn't a product announcement; it's proof that agentic AI works on the infrastructure everyone said was too old, too complex, and too critical to touch—and that changes what's possible for every system built before 2010.