OCTAVE

EBOOK

10 questions that instantly separate quality-trained AI from marketing hype

Introduction:
Cutting through the AI hype

It's not hard for a QMS vendor to claim they have AI now. In fact, the challenge isn't even finding AI-powered solutions. It's finding ones that actually work for quality management.

Generic AI doesn't understand that "CAPA" has specific regulatory meaning or that root cause analysis requires more than asking "why" five times. AI built for quality management, trained on regulatory frameworks, quality terminology and industry workflows, delivers fundamentally different results.

These 10 questions help you separate substance from hype and ask vendors to explain HOW their AI works in ways that reveal depth — or lack of it.

01

Solving your quality challenges

How does AI reduce manual, repetitive tasks while maintaining data accuracy and compliance?

What to look for: Strong solutions provide specific examples, such as “our AI auto-populates CAPA forms using data from linked nonconformance reports” not “we automate documentation.” Ask for evidence that the AI actually improves accuracy. Look for confidence indicators that show “I’m 92% confident this is correct” versus blindly suggesting answers. Audit trails matter. You should see who accepted the AI suggestion, when, and what the original recommendation was.

Be wary of vendors who emphasize speed without explaining how they maintain data quality. “Generate CAPAs 10x faster” sounds great until you’re explaining to an auditor why three AI-generated investigations have incorrect root causes. The best solutions embed compliance checks into the automation and ensure the AI won’t let you skip required fields or suggest actions that violate your procedures.

02

What specific quality management challenges does your AI address, and what measurable outcomes have customers achieved?

What to look for: Avoid vendors who describe AI capabilities without connecting them to business outcomes. “Our AI uses natural language processing” tells you nothing. “Our AI reduced CAPA cycle time from 87 days to 62 days at a Tier 1 medical device manufacturer” tells you something useful.

Ask for customer references or case studies that demonstrate real-world impact. If they can’t provide a name or PDF due to partner agreements, they should at least be able to provide findings such as cycle time reduction percentages, error rate decreases, hours saved per CAPA or audit findings decreases.

Be skeptical of generic promises. If they claim their AI works for medical devices, automotive and food manufacturing equally without backing it up with numbers, they probably don’t have a deep understanding of them. Strong vendors articulate how their AI addresses your industry’s specific challenges, such as 21 CFR Part 820 compliance for medical devices, IATF 16949 requirements for automotive, etc.

03

How does the AI help shift from reactive problem-solving to proactive risk prevention?

What to look for: Many vendors claim predictive capabilities. But do they deliver? Look for solutions that explain how they identify risks before they materialize, not just that they do it.

Ask for specific use cases: “Can your AI spot that we’ve had three supplier deviations with high moisture content in the past six months and flag incoming lots from that supplier for additional testing?” “Does it notice that CAPAs in Building 3 consistently miss deadlines and recommend process improvements?” “Will it alert me if a procedure change affects five other controlled documents?”

Strong solutions help you prevent problems, not just document them faster after they occur. If the demo only shows AI filling out forms, ask what happens after the forms are complete. Proactive value lives in pattern recognition, trend analysis and early warning systems.

04

AI foundation & intelligence

What domain expertise and data informs your AI models, and how does this ensure recommendations are relevant and compliance-ready?

What to look for: This question separates AI built for quality management from generic AI bolted onto QMS platforms.

Generic large language models don't understand that "CAPA" and "corrective action" have specific regulatory meanings, that root cause analysis requires more than asking "why" five times or that "effectiveness verification" isn't optional in FDA-regulated industries. Ask vendors: "What makes your AI recommendations relevant to quality management versus someone asking ChatGPT for help?"

Strong vendors explain how their AI is trained or tuned with quality domain knowledge, including decades of CAPA data, regulatory frameworks (FDA guidance, ISO standards) and industry-specific terminology. They should understand the difference between a CAPA, a deviation and a nonconformance — and generate context-appropriate recommendations accordingly.

05

How does your AI handle situations with uncertainty or multiple valid options, and what transparency do users have into AI reasoning?

What to look for: AI isn't always certain, and that's fine, but you need to know when it's uncertain.

Look for transparency mechanisms, such as confidence scores ("I'm 73% confident this is the root cause—consider additional investigation"), explanations of reasoning ("I'm recommending Engineering based on three similar CAPAs where equipment modifications were required") or flags for human review ("This is an unusual situation—verify my suggestion").

Avoid "black box" AI that provides answers without explaining why. Ask vendors: "What happens when your AI encounters an edge case or a situation it hasn't seen before?" Weak answers deflect or claim their AI just works. Strong answers demonstrate thoughtful design for uncertainty.

As an example, when multiple departments could be responsible for an issue, does the AI pick one arbitrarily, or does it flag the ambiguity and suggest you involve both? The latter respects your judgment. The former assumes it knows better than you do.

06

Integration & user experience

How is AI integrated into daily workflows. Is it embedded in existing processes or accessed through separate tools?

What to look for: There's a critical difference between AI embedded in your workflows versus AI that requires switching to separate dashboards, chatbots or analysis tools.

Embedded AI delivers recommendations at the point of need. While you're filling out a CAPA form, it suggests similar past investigations. While you're assigning an investigation, it recommends the team based on issue type. While you're scheduling effectiveness checks, it calculates appropriate timelines based on problem frequency. You never leave your current screen.

Separate tools create friction. Users must remember to use them, copy data between systems and context-switch repeatedly. Adoption suffers. Ask vendors: "Do I need to leave my current workflow to access AI capabilities?" If the answer is yes, consider the real-world adoption challenges.

The best AI feels invisible. It's just a natural enhancement to existing work, not an additional step you have to remember.

07

What level of control and configurability do users and administrators have over AI features and recommendations?

What to look for: One-size-fits-all AI rarely fits anyone well. Your organization has unique processes, terminology, risk tolerances and compliance requirements.

Strong solutions offer configurability at multiple levels:

- Administrators: Control where AI is deployed, what confidence thresholds trigger recommendations, which teams get which features and how aggressively automation is applied
- Individual users: Accept or reject each suggestion, provide feedback that improves future recommendations and disable AI features they don't find helpful

Be wary of forced automation that removes human judgment. "Our AI automatically assigns all CAPAs" sounds efficient until it routes a critical product failure to the wrong department because your organizational structure doesn't match the vendor's assumptions.

Ask: "Can we tailor this AI to match our definition of quality, not just yours?" and "What happens if the AI recommendation doesn't fit our situation?" Strong answers demonstrate flexibility.

08

Trust, governance & evolution

How does your solution ensure AI recommendations are explainable, traceable and aligned with regulatory standards?

What to look for: In regulated industries, “the AI told me to” isn’t an acceptable audit response.

You need three things:

- **Explainability:** Why did the AI make this recommendation? “Based on 47 similar CAPAs, Engineering resolved this type of equipment failure” is explainable. “The model determined Engineering was appropriate” is not.
- **Traceability:** Can we document the AI’s role in decisions? Audit trails should show what the AI recommended, what the user accepted or modified and when.
- **Regulatory alignment:** Does this meet FDA 21 CFR Part 11, ISO 13485, EU AI Act or other applicable standards?

Strong vendors reference specific frameworks they follow, such as the EU AI Act for high-risk systems, ISO 42001 for AI management and NIST AI Risk Management Framework. They explain how their solution creates compliant documentation.

Vendors who can’t articulate their AI governance approach should send make you pause and think twice. It’s possible they’re AI isn’t trained or tuned for quality management. Regulated industries require AI that understands all of the nuances of each space.

09

How do you balance automation with human expertise and judgment in critical quality decisions?

What to look for: The best AI-powered QMS solutions augment human expertise rather than replace it.

Look for "human-in-the-loop" design where AI handles repetitive tasks (form population, document generation, data compilation) and surfaces insights (trend analysis, risk identification and similar past issues), but humans maintain control over critical decisions (root cause determination, corrective action selection and effectiveness verification).

Be skeptical of vendors who promise to "eliminate human error" through full automation. Quality management requires judgment, context and accountability that AI can't fully replicate. Root cause analysis isn't just filling out a form — it's investigative work that requires understanding your processes, recognizing patterns and exercising judgment about what's a symptom versus what's a system failure.

Strong solutions explain how they preserve user agency. Can recommendations be overridden? Does the system learn from human corrections? Is there clear ownership of decisions? Who's accountable if the AI suggestion is wrong?

The goal is to adopt AI as a trusted advisor not as an automated dictator.

10

What's your approach to evolving AI capabilities, and how do customer needs shape your roadmap?

What to look for: AI technology is evolving rapidly. Your QMS investment should evolve with it.

Ask vendors about their innovation strategy:

- Are they developing AI in-house, partnering with technology leaders or white-labeling third-party tools?
- How do they incorporate customer feedback into development?
- What's their track record of AI innovation, and can they show progression over the past two years?

Look for evidence of customer-driven innovation, such as early adopter programs, design partnerships and customer advisory boards. Companies that develop AI in isolation from actual quality professionals often build features that sound impressive but don't solve real problems.

Be cautious of vendors with vague roadmaps ("we're exploring AI opportunities") or those who can't articulate how they'll keep pace with AI advancements.

Also ask about their partnership ecosystem. Access to best-in-class AI technology often comes through strategic partnerships with AI companies not just internal development. A QMS vendor partnering with a leading AI platform can deliver better AI capabilities than a QMS vendor trying to build everything themselves.

Final question to ask yourself: "Is this vendor treating AI as a strategic differentiator they're investing in, or as a marketing checkbox they're checking off?"

One final comment

Don’t just ask these questions — dig into the answers. Ask for demos that show specific examples. Request customer references you can actually call. Push back on vague responses.

The vendors who get defensive or evasive when you ask tough questions are telling you something important. The vendors who welcome the scrutiny and provide specific, honest answers, including limitations, are the ones worth your time.

Quality management is too important to hand over to AI you don’t understand and can’t trust. These questions help you find AI that enhances your expertise rather than replacing it.

AI-powered QMS evaluation: 10 essential questions

1	**Manual tasks, accuracy & compliance?**	Show me confidence indicators that prevent incorrect AI suggestions.
2	**Specific challenges & measurable outcomes?**	Give me customer results in my industry with metrics and a reference to call.
3	**Proactive vs. reactive?**	Walk me through a real example where AI prevented a risk before it materialized.
4	**What domain expertise informs your AI?**	Explain what makes your AI understand quality management vs. generic AI—does it know the difference between CAPA, deviation and nonconformance?
5	**How does AI handle uncertainty?**	Show me when your AI is uncertain. How does it communicate that and explain its reasoning?
6	**Embedded in workflows or separate tools?**	Show me where AI appears. Do I stay in my current screen or switch to a separate tool?
7	**What control do users and admins have?**	Show me controlling confidence thresholds, deployment locations and how users reject suggestions.
8	**Explainable, traceable & regulatory-aligned?**	Which frameworks do you follow (EU AI Act, ISO 42001, NIST) and show me your audit trail?
9	**Automation vs. human judgment?**	Walk me through a critical decision. Where does AI help vs. where humans maintain control?
10	**How do AI capabilities evolve?**	Are you building in-house or partnering, and give examples of customer-shaped features?

About Octave

Octave is a leader in enterprise software, turning data into decisive action and intelligence into your edge. Our software solves for and simplifies complexity, from the design and build to operations and protection of people, property and assets – for any scope, at any scale. For decades, we've partnered with customers to sharpen performance, elevate efficiency and amplify results. From factory floors to entire cities, our solutions are tuned to scale up what's possible from day one onward.