elixir | Cloud

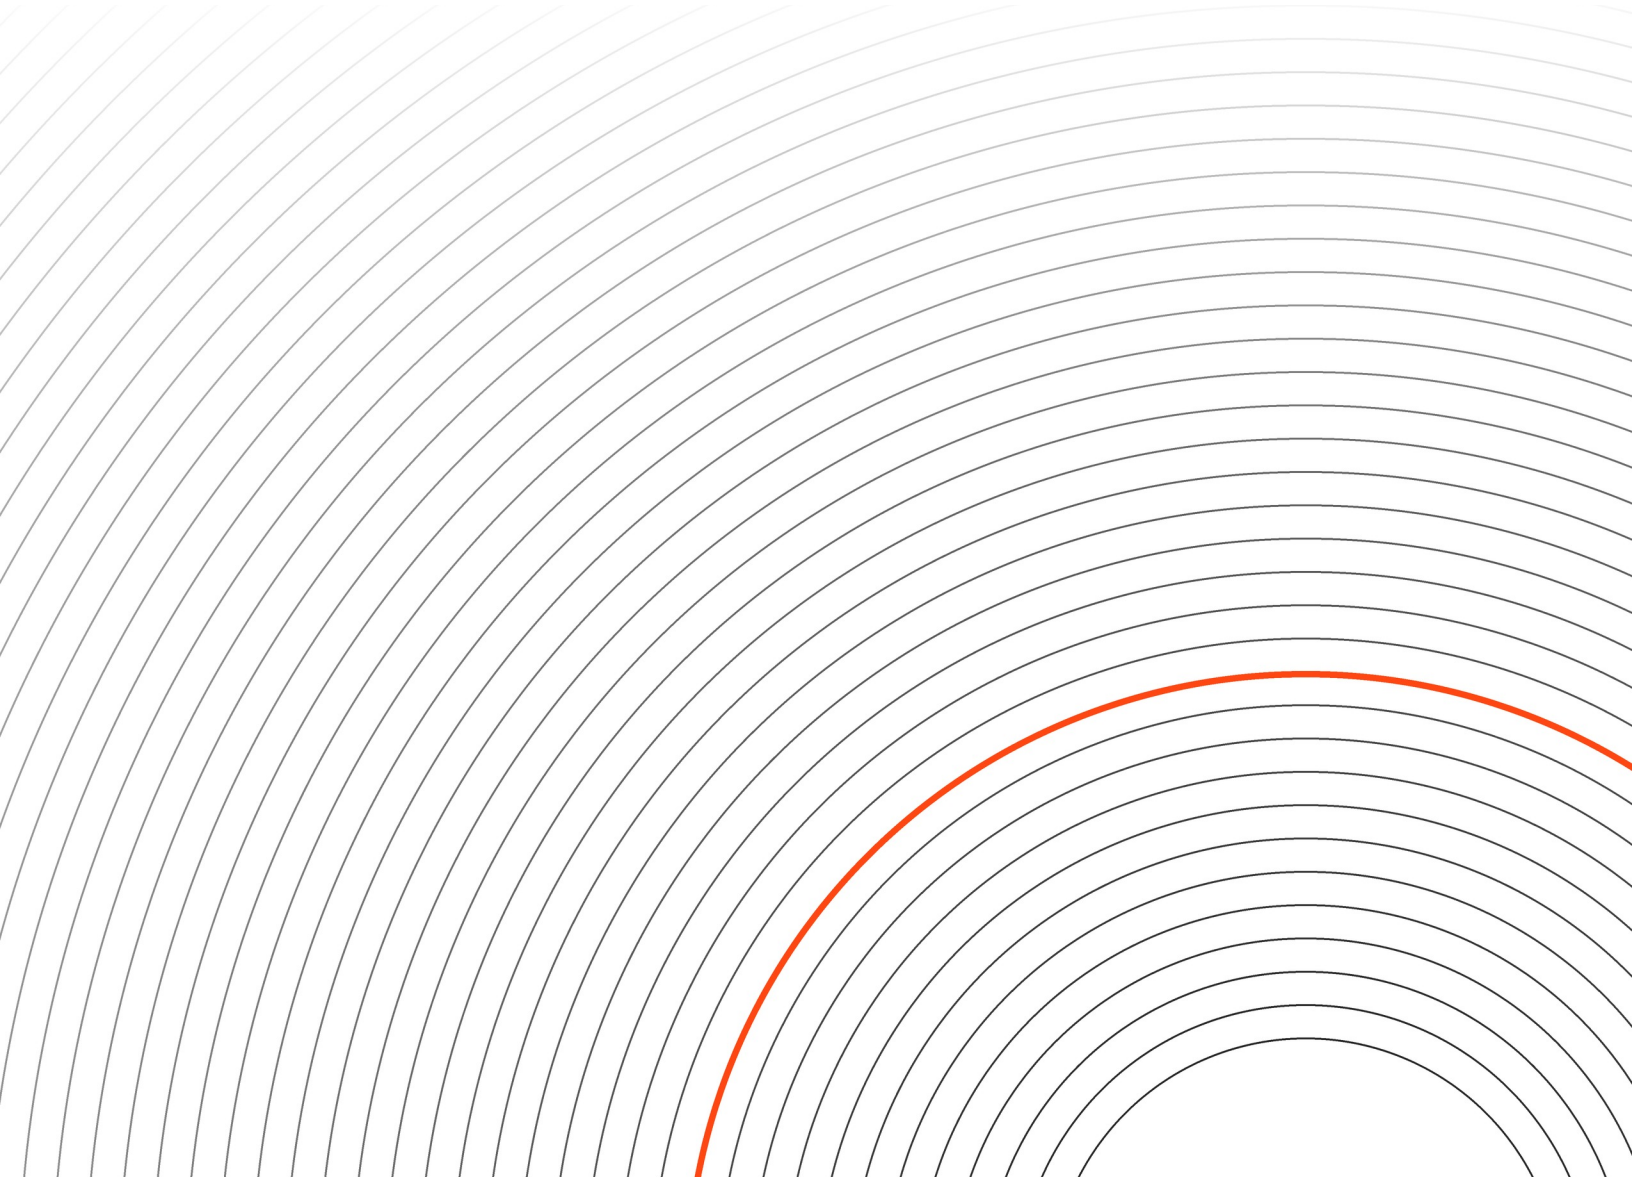
WEBINAR INSIGHTS

Prove It: What AI Can Actually Deliver in Regulated Customer Communications

An expert conversation with QKS Group and Elixir on separating AI hype from production reality in banking, insurance, and healthcare communications

Libby Koehn, Elixir | Saurabh Raj and Amandeep Khanuja, QKS Group | Hart Johnson, Elixir

November 2026



elixir | Cloud

Contents

00 About This Conversation 3

01 From Pilot to Production 5

02 What Makes Risk Officers Hesitate 7

03 Hype vs. Reality 10

04 AI-Generated Text vs. a Governed Library 12

05 Through the Auditor's Eyes 14

06 Fragmentation and the Audit Trail 16

07 Beyond Drafting: The Full Lifecycle 17

08 Prove It: Measuring the Impact of AI 19

09 Five Questions to Ask Your CCM Vendor 21

This asset is drawn from the live webinar held November 5, 2026. Remarks have been edited and condensed for length and clarity.

About This Conversation

Boards are demanding progressive AI roadmaps. Risk committees are demanding proof. In regulated customer communications, the gap between those two demands is where most AI initiatives stall.

In banking, insurance, and healthcare, you can't plug a public large language model into customer communications and hope for the best. On November 5, 2026, Elixir Technologies hosted a fireside chat with two principal analysts from QKS Group to separate what AI can actually deliver from what is still hype. No slides, no sales presentation: just a candid conversation across four areas.

- Where AI adoption in regulated communications stands today
- What is hype and what is real
- What regulators and auditors expect to see
- What buyers need across the full communications lifecycle, not just drafting

The Panel

Libby Koehn

Chief Product Officer, Elixir Technologies

Leads product strategy for Elixir's cloud-native customer communications platform, built for regulated industries like insurance and healthcare. Led the conversation.

Saurabh Raj

Principal Analyst, QKS Group

Researches the customer communications management market: trends, vendor strategies, and enterprise buying priorities.

Amandeep Khanuja

Principal Analyst and Associate Director, QKS Group

Specializes in enterprise communication, customer experience, and AI. Advises enterprises on AI adoption, product strategy, and competitive positioning.

Hart Johnson

Executive Managing Director, DevSecOps, Elixir Technologies

Security and compliance leader responsible for Elixir's HITRUST audits and chief architect of Elixir's AI solutions.

Hosted by Melissa Yearta, Director of Demand Generation, Elixir Technologies.

At a Glance

Five takeaways from the panel

01 **The model is the easy part.**

Pilots rarely stall on AI capability. They stall on ownership, content quality, and evidence: who approves the output, where the data went, and how to prove later what was sent and why.

02 **Risk officers fear the unexplainable.**

The three barriers are data exposure, content risk, and explainability. Being unable to reconstruct a communication worries them more than a single error.

03 **"Agentic" is often just a new label.**

The measurable gains today are in template analysis, migration, quality assurance, and assisted drafting and translation, with a human in the loop.

04 **Deterministic at the core, AI around it.**

Regulated wording belongs in a governed library with approved content and business rules. AI adds value around that core, and the release decision stays with accountable people.

05 **Baseline before you pilot.**

Measure speed, quality, cost, and customer outcomes before AI enters the process, and report a small set of measures the same way every month.

The through-line: **the AI never stands in front of the regulator. The business does. Every AI decision in regulated communications has to be one you can explain, prove, and own.**

01 From Pilot to Production

Where AI adoption actually stands

LIBBY KOEHN ASKS

Looking across the CCM market, where does AI adoption in regulated communications actually stand, and where are enterprises getting stuck between pilot and production?

SAURABH RAJ · QKS GROUP

AI interest is everywhere. Almost every buyer we speak with has an AI agenda, and many are already running pilots. But AI that runs inside a live, regulated communication workflow is still very rare. In this year's evaluation, only a small minority of vendors could show AI in production with measured customer results. For most of the market, AI is still a pilot, a beta, or a roadmap item.

That gap matters most in the industries CCM depends on. Banking and financial services make up roughly a third of the CCM market, and healthcare is the fastest-growing segment we have seen in a couple of years. The largest and fastest-growing parts of this market are also among the most heavily regulated. We see a similar pattern in Asia Pacific, where banks and insurers are active on AI and digital transformation. Buyers are often ahead of the market.

So where do enterprises get stuck? In our view, it is rarely the AI itself. A good model can draft a reasonable letter. The problem is everything around it. Someone has to approve the output, someone has to know where the customer data went, and someone has to be able to prove later exactly what was sent and why. If a pilot can't answer those three questions, it doesn't get past the risk committee, no matter how good the demo looked.

~1/3

of the CCM market is banking and financial services

67%

of communication owners name integrated operations a top priority

45%+

of communication owners name workflow redesign a top priority

Source: QKS Group research, as cited by Saurabh Raj

Those are foundation issues rather than AI features, and they show where the real work sits. In the boardroom, the conversation is about which AI to buy, while the operations team is still trying to change a letter without raising an IT ticket.

AMANDEEP KHANUJA · QKS GROUP

Working with operations teams this year, I see three gaps.

01 Ownership

Very often an innovation or digital team runs the pilot, and it goes well in a controlled setting. Then it lands with compliance and operations, and none of those teams were in the room when it was designed. The pilot works in the demo, but nobody owns it once it has to run in production every day.

02 Content

Many organizations have thousands of templates built over decades. We met with a bank that had 11,000 templates and a goal to reduce them. After mergers and platform changes, content is duplicated, slightly different from one version to the next, or simply out of date. If you point AI at a messy content estate, it doesn't clean it up for you. It just produces more messy output.

03 Evidence

A risk committee wants to see how the output is checked, how it is logged, and who signs off. A good demo doesn't answer those questions. A documented control process does.

For data and AI leaders, this pattern will sound familiar, because it's the same with any model in production. The model is usually the easy part. Ownership, data quality, and monitoring are where the real effort goes.

LIBBY KOEHN · ELIXIR

Moving from pilot to production, you can do a lot of different things with AI, but the question is, should you? It has to scale, it has to return value, and it has to actually optimize the experience. If it doesn't do those things, it's going to fall flat.

There's also a genuine, justifiable fear. In regulated communications, if AI hallucinates, we've got a problem. A hallucination in a marketing message might be embarrassing. In a Truth in Lending statement, it could be a multimillion-dollar compliance error.

AMANDEEP KHANUJA · QKS GROUP

It's easy to adopt AI in customer service and CX, where the level of risk is low. In regulated industries, even a margin of error of 0.001% matters. So the entire mindset shifts from adopting AI to controlling and governing it first. Every conversation Saurabh and I have starts there. If you can control it and govern it, and then move it up the ladder, you'll get the real value you're looking for in a safe environment.

SAURABH RAJ · QKS GROUP

Everybody wants to go from pilot to proof. Nobody wants to be stuck in between after investing heavily, especially with what tokens cost at enterprise scale. Token economics are at stake, so you need that mindset of control from the start.

02 Governance and Risk

What makes risk officers hesitate

LIBBY KOEHN ASKS

Whether the rules are HIPAA, CMS marketing guidelines, TILA disclosures, or UDAAP, what specific risks make risk officers hesitate when AI is proposed in the communication lifecycle?

AMANDEEP KHANUJA · QKS GROUP

When we speak with enterprises, the hesitation is less about what AI can do and more about what happens when something goes wrong. That concern comes down to three interconnected risks, each building on the last.

01 Data exposure

A customer letter is full of personal data, and in healthcare much of it is protected health information under HIPAA. The first thing a risk officer asks is where that data goes when AI touches it: whether it leaves the organization, whether anything is stored along the way, and whether it could ever be used to train something else.

02 Content risk

This is where regulated communications differ from marketing. A TILA disclosure has to show the right terms in the right place. Medicare communications have to follow CMS rules on wording and format. UDAAP asks a simple question that has to be answered at scale: could the message mislead or cause harm? If AI changes one sentence the wrong way, that can become a regulatory finding. In a high-volume program, it becomes a remediation exercise across thousands of customers, with letters to reissue and sometimes money to return.

03 **Explainability**

Months or even years later, a regulator or court may ask what a particular customer received and why. If AI is involved anywhere in that chain, you need to show what it did, who checked it, and which version actually went out.

> ***Risk officers worry less about AI being wrong once than about not being able to explain it afterwards.***
>
> **AMANDEEP KHANUJA, QKS GROUP**

SAURABH RAJ · QKS GROUP

A CAUTIONARY TALE

In 2024, an airline's website chatbot told a customer they could claim a bereavement fare discount after booking, which wasn't the airline's policy. When the customer took the case to a Canadian tribunal, the airline argued the chatbot was a separate legal entity. The tribunal rejected that argument and held the company responsible for everything on its website, including what the chatbot said.

LIBBY KOEHN · ELIXIR

The pillars you identify, data exposure, content accuracy, and reconstruction and auditability, are exactly what we see. It's precisely why generic AI models fail in regulated customer communications. You can't have a black box dealing with CMS, TILA, or HIPAA guidelines.

This is why Elixir built its AI framework the way it did. We anchor everything on AWS Bedrock so it is provably private. We run in a zero-retention environment. We don't train on customer data in any form, and we never scrape the public web. That removes the data exposure risk.

We solve content accuracy and reconstruction differently, too, because we don't guess on compliance. It's too high risk. We run a simple prompt structure against validated, highly curated, deterministic knowledge repositories, such as CMS guidelines, TILA guidelines, KYC, or HIPAA Safe Harbor. Combined with our native AI-facilitated de-identification capabilities, risk officers get control of what's happening: complete data privacy and the traceability they need. The historical record is kept in a governed workflow in their own system, so they can audit and know what happened.

HART JOHNSON · ELIXIR

If there's a problem in a communication, the AI is not going in front of a judge, a jury, or a committee. The business has to go in front of them and explain it.

On data exposure: we are provably private. We don't store any of the content on our side, which reduces your blast radius. From a cybersecurity standpoint, we keep seeing new vulnerabilities published with the maximum severity score of 10. Being provably private means we don't just say we're not storing the content; you can verify it with our tools and see how they run and what they do. That makes a risk officer's evaluation a lot easier, because there's no new data store sitting out there to worry about, especially given the sensitive content people are putting into AI chatbots today.

On accuracy: with 41 years of history, we know this industry and these regulated markets inside and out. We're building curated content that AI models can rely on, and we're using the right AI models for the job. The industry tends to throw the biggest model at everything. We ask what the right underlying model is, how we'll train it, what content it will draw on, and how it can connect to the customer's knowledge base and content, so they get branded, personalized communications at scale, safely, with a human in the loop.

SAURABH RAJ · QKS GROUP

Using the right model is the key to a good AI story. It's not just about picking a good model; you have to train it well, and every high-end model consumes a lot of tokens. Amandeep and I think of it as a priority pyramid: low, middle, and top priority work, with the model matched to the job. Using a top-tier model to translate a single message doesn't make sense. And the people using these models need that AI knowledge, too.

AMANDEEP KHANUJA · QKS GROUP

What I believe will shape the future of CCM is what we call **guided AI**. AI capabilities will keep advancing, and the level of automation available to enterprises will increase significantly. But the more capable AI becomes, the more important it is for enterprises to define what they want AI to achieve, what it is authorized to do, and where the boundaries lie. You establish the direction, the business rules, and the level of authority, and then allow AI to operate intelligently within those boundaries. You should be the one guiding. Don't let AI guide you.

GUIDED AI

The real measure of AI maturity in CCM is how much responsibility an enterprise can confidently entrust to AI while maintaining control of the communication when it reaches its customers.

03 Hype vs. Reality

Where AI delivers measurable value

LIBBY KOEHN ASKS

Your research looks at where AI in CCM is real and where it's still hype. What are the clearest examples of hype, and where is AI delivering measurable value?

SAURABH RAJ · QKS GROUP

THE HYPE

"Agentic" everywhere. Most vendors in our evaluation now use the word, in RFPs and brochures alike. Look closely and much of it is rule-based workflow with a new label. Workflow automation is useful and creates operational efficiency, but it isn't autonomy.

Demos presented as production. A demo shows the best case: clean data and a friendly script. Production means millions of documents, messy data, and an audit at the end. In our evaluation, a number of features described as available were really in beta or still on the roadmap.

Fully autonomous communications. Regulators have a firm grip on this market. We have seen regulators fine investment advisers for misleading claims about their use of AI.

THE REALITY

Template analysis. Reading thousands of templates and spotting duplicated content, the repetitive work teams have done by hand again and again.

Migration. Moving legacy templates to a new platform, work that used to depend on scripts.

Quality assurance. Automated checks that compare each output against the approved version, with a person who approves it.

Assisted drafting, translation, and plain-language rewriting. Always with a human in the loop.

These are where adoption is moving fastest and where we see production results.

AMANDEEP KHANUJA · QKS GROUP

Here is a simple test anyone can use to evaluate vendors.

1. **Ask what happens when the AI is wrong.** A mature answer explains how errors are detected, how they're logged, and who takes responsibility. An answer that simply emphasizes accuracy is weaker.

2. **Ask what the AI is actually allowed to change.** Can it modify disclosure wording, financial figures, or eligibility conditions?
3. **Ask for evidence of production adoption.** If a capability is genuinely delivering value, there should be customers who can explain where it's being used.

LIBBY KOEHN · ELIXIR

The market doesn't need any more smoke and mirrors or glorified rules-based workflows disguised as AI. It needs practical solutions that cut operational friction and keep human review layers in place, so the value is measurable and grounded in reality.

Migration is a great example, and it's part of why we created Elixir Migrate (formerly Catalyst). It takes any compiled output, like a PDF. We don't need the original system, just a data feed, and we can migrate the template into an editable design, including the data mapping and business logic. That can significantly reduce the implementation burden.

Quality assurance is another. Because we use curated knowledge repositories, AI can check whether content meets HIPAA Safe Harbor or TILA guidelines and show where to look. The AI gives a suggestion, and the human decides whether to accept or reject it. The human is always in control. Translation is powerful, too, because you need to reach your audience in the language they need and in the right accessibility format, such as braille. At Elixir, we look at the entire communication lifecycle and only build AI where it unlocks value. And everything is human in the loop.

HART JOHNSON · ELIXIR

Elixir Compare (formerly AetherDiff) shows what building AI where it matters looks like. Checking whether the pixels are in the right spot against the approved output isn't a job that needs AI. Reading the content to confirm the word choice hasn't changed and still matches the customer's approved baseline is a perfect job for AI. So is improving accessibility on published output, like generating alt text for an image.

We don't slap AI on everything just to have AI. We fit it in where it makes sense and does the job well, so we can answer the question: how do we know it's accurate? Again, the AI isn't going in front of a regulator to say, "Oops, my bad." The humans have to answer for that.

LIBBY KOEHN · ELIXIR

Human in the loop gives you governance and control. But if you're using AI, it also needs a return on investment and has to improve efficiency and productivity. If it doesn't do that, it's hype.

04 Governed by Design

AI-generated text vs. a governed content library

LIBBY KOEHN ASKS

How AI interacts with the customer's core system is where compliance is won or lost. Why does it matter whether AI writes customer-facing text itself or works alongside a governed library of approved content, with rules and human approval controls?

AMANDEEP KHANUJA · QKS GROUP

Ultimately the question is how much control an organization keeps over what reaches its customers. The two approaches carry very different levels of risk, even when the output looks similar on the page.

AI WRITES THE TEXT

Probabilistic

Ask the same question twice and you may get two different answers, both perfectly reasonable. That's fine for brainstorming or a first draft. It's a real problem for a disclosure that has to say exactly the same thing every time it's sent.

GOVERNED LIBRARY

Deterministic

Regulated wording is written once, approved by compliance, and stored as controlled content. Business rules decide which approved piece goes to which customer, based on things like the product or their state. The same input always produces the same output.

Because the same input always produces the same output, you can explain exactly why a customer received a particular paragraph. That's what auditors mean when they talk about deterministic communications. Think of an annual notice to Medicare members: the benefits, the costs, and the required statements all have to be exact and match the plan's approved benefits.

So the real question isn't AI or no AI. It's where AI sits in the process. AI is very useful around the governed core: drafting a new paragraph for compliance to review, simplifying complex language, suggesting a translation, or flagging a sentence that looks risky. The release decision stays with the rules and with the people accountable for them.

LIBBY KOEHN · ELIXIR

In our space, unsupervised is a non-starter. Regulated communications are too risky and the fines are too big. The goal shouldn't be to replace the human or bypass the system of record. It should be to supercharge them while keeping deterministic guardrails in place.

A writing assistant shouldn't have free rein. It can be a great copilot and offer good suggestions to refine the language, and curated knowledge repositories can flag when content doesn't meet TILA guidelines, KYC requirements, or the CMS framework. But the human has to stay in control, to say: I've signed off on this, I've accepted it, and I take responsibility that it's ready for publication.

> ***The real question isn't AI or no AI. It's where AI sits in the process.***
>
> **AMANDEEP KHANUJA, QKS GROUP**

05 Through the Auditor's Eyes

What regulators expect to see

LIBBY KOEHN ASKS

When an organization walks a regulator or auditor through an AI-assisted communication process, how much weight do the AI provider's non-training and non-retention commitments carry, and what else would they expect to see?

SAURABH RAJ · QKS GROUP

They carry a lot of weight, but they're not the whole story. Data handling is usually the first objection a risk officer raises, and for good reason. The good news is that the market has moved. Most enterprise AI platforms now commit not to train their models on customer data, and some go further and commit not to store prompts or responses at all. That removes a major barrier. A year or two ago, many pilots stopped right at this point.

It's still worth reading the details. Ask what is logged, what is cached, and for how long. It also helps to separate two layers that often get mixed up:

LAYER 1

The AI provider

How it treats your data: training, and whether prompts and responses are stored.

LAYER 2

Your CCM platform

What it keeps, such as drafts, logs, and outputs, and where they are stored.

Both need clear answers. Ask where each commitment comes from and how long it lasts under the contract. And don't forget data residency. Customers always ask about it, yet it's often missed when buyers and vendors talk.

HART JOHNSON · ELIXIR

It comes down to transparency, because it touches everything: what are we storing, what are we training on, what are we doing with the data. At Elixir, we share how we do this: what we use your data for, what we don't, and even the metadata we store to understand how our product is being used. That's one of the core tenets of our AI strategy. Everything is provable, not only the privacy. We'll show you what we're doing, how we're doing it, and why.

That also answers the residency question, because we can show the entire lifecycle of something as simple as a prompt: where it goes, what happens to it when it comes back, and the full pathway for that data. My advice for everybody is to be transparent, not just in contractual language but in your public-facing materials. We want customers to feel comfortable.

LIBBY KOEHN · ELIXIR

Provably private and zero retention is our architecture. We don't store anything. Information is stored in the customer's environment, because they can choose to keep their prompts, their history, and the AI's recommendations and suggestions in local storage wherever they want: an S3 bucket, an in-house content management system, or their data lake. We have no record of it.

You don't want a vendor storing or persisting your prompts, even for 30 days. Now you have to explain to the auditor why it was there at all, especially if something went wrong.

SAURABH RAJ · QKS GROUP

Commitments are where it starts. How you align them with the architecture is what matters, and explaining it through the architecture carries real weight.

ASK YOUR VENDOR

What is logged? What is cached, and for how long? Where does the data reside? And can you **show** the full pathway of a prompt, not just describe it in a contract?

06 One Connected Record

Fragmentation and the audit trail

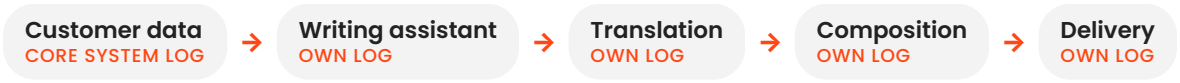
LIBBY KOEHN ASKS

When AI tools are added piece by piece to separate systems, what happens to the audit trail, and what does a traceable record of an AI-assisted communication need to contain?

AMANDEEP KHANUJA · QKS GROUP

One of the biggest risks with AI adoption is that enterprises deploy AI faster than they build the accountability around it. When AI operates across disconnected systems, it becomes increasingly hard to reconstruct how a communication was produced. Put simply, the audit trail breaks into pieces, and it usually happens gradually rather than all at once.

Think of a fairly typical setup:



Each may keep its own log in its own format, so nobody holds the full story of a single letter. When a complaint comes in and someone asks what a customer received, which version it was, and why the wording changed last month, the team has to pull logs from four or five systems.

48% of communication owners told QKS they are adding channels faster than they are connecting the content behind them. Every new tool added without a connected record makes the gap a little wider.

Source: QKS Group survey, as cited by Amandeep Khanuja

HART JOHNSON · ELIXIR

The fragmentation and tool sprawl we're seeing across every industry right now is massive. The way to solve it is to avoid completely disconnected systems and make it easy to integrate different tools and technology. But it ultimately comes down to understanding the full workflow. If you understand how you'll use every piece of the puzzle to build a communication, you can build your logging strategy around it.

Our tools can be used standalone if that's all you need, for example a digital signature. They also connect into the full communications lifecycle. For audit logging, you need a way to connect these systems elegantly, and that starts with strategy.

LIBBY KOEHN · ELIXIR

Our aim isn't to scatter disconnected point solutions. We're building AI capabilities that sit within a single communication workflow, so there's a logical path. Everything is API-enabled, so a capability can run as an isolated instance, but the model runs with zero retention and the audit record is kept locally in the customer's environment. That gives you a traceable stack. Zero retention and provable privacy give us a governance model, with a human in the loop on everything.

07 Beyond Drafting

Evaluating the full communications lifecycle

LIBBY KOEHN ASKS

Generating compliant text is only one step. Over the next 12 to 24 months, which parts of the lifecycle beyond drafting should buyers evaluate, such as output quality assurance, accessibility, translation, print and mail preparation, or bundling and kitting?

Data capture → **Drafting** → Review and approval → Quality assurance →
Translation and accessibility → Print, mail, and digital delivery → Bundling and kitting →
Archive and retrieval

SAURABH RAJ · QKS GROUP

With generative AI, you can write a letter in minutes. But if review, translation, and print preparation still take weeks, the customer sees no difference at all.

The lifecycle also starts upstream, with data. In our survey, roughly 61 to 63 percent of communication owners said the data captured before a communication matters just as much as the communication itself.

Quality assurance matters because every change to a template can break something else, whether that's a logo or a table that spills over. Automated comparison against the approved version needs to catch those problems before they reach the customer.

Accessibility rules are tightening. The European Accessibility Act has applied since June 2025 and covers consumer banking services, among others. In the US, organizations that receive HHS funding have to make their websites, apps, and online documents accessible.

EDITOR'S NOTE

HHS's Section 504 rule requires recipients of HHS federal financial assistance to meet WCAG 2.1 Level AA for web content and mobile apps. In 2026, HHS extended the compliance dates to May 11, 2027 for organizations with 15 or more employees and May 11, 2028 for smaller organizations.

Delivery is changing a lot, too. More customers want documents in apps, portals, and messaging channels, and each of those needs the approved content and a record, just like a printed letter. And at the end of the lifecycle, regulatory communications must be kept and retrieved for years, in the exact version the customer received.

LIBBY KOEHN · ELIXIR

We've done some interesting things here that aren't AI. We built collaborative tooling into our system to speed up review and approval. Then Elixir Compare, one of our AI tools, does both a pixel-level and a contextual comparison, so you can see the differences from one version to another.

Add collaborative tooling and built-in notifications, and you can draft a communication with an AI-assisted approach, with a human approving the suggestions, finalize the template, and immediately kick off a workflow to legal, compliance, or other reviewers to sign off. We can't publish until a human gives formal approval. All versions and content are stored in the customer's environment, because our CCM also takes a zero-retention approach. We're a transitory data processor, so the content stays under the organization's control.

Post-processing matters as well. Many communications aren't a single template. They include contracts, other documents, and attachments. Bundling and kitting, watermarking, encryption, and making sure it's all highly secure become really important. When that work serves the communication workflow, it unlocks real ROI.

08 Prove It
Measuring the impact of AI

LIBBY KOEHN ASKS

To justify an enterprise budget, your investment needs proof. How should buyers measure the impact of AI in their communications, and what needs to be in place before a pilot starts so the results hold up?

AMANDEEP KHANUJA · QKS GROUP

We see a growing gap between enterprises investing in AI and those that can demonstrate business value from it. The difference often comes down to whether they defined what success looks like before introducing AI into the process.

Start with a baseline. It sounds obvious, but many teams skip it and then struggle to prove anything later. If you don't know how long a template change takes today, you can't show that AI made it faster. Measure your current process for a few weeks before the pilot starts.

Then measure outcomes rather than activity. Words generated or prompts used tell you about usage, not value. We suggest four groups of measures that together cover the whole picture:

01 Speed
How long the work takes, such as a template change, compared with your baseline.

02 Quality
How output is checked and how many problems are caught before customers see them.

03 Cost
What it costs to produce and run communications.

04 Customer
Do complaints go down? Do calls asking what a letter means go down? Do more customers act on time, such as paying a bill or renewing a policy?

It also helps to pick a small number of these, perhaps four or five, and report them the same way every month.

SAURABH RAJ · QKS GROUP

Behind AI, it's all data. So don't just measure the impact of AI; measure the impact of data. As legacy communications modernize, data is fragmented, sometimes stored in local systems. Understand which data isn't good for your AI and which data you need to gather first. When we run advisory calls and ask which data a team is working from and where their CRMs and ERPs are, many are still unsure.

HART JOHNSON · ELIXIR

That up-front data strategy is so important when you roll out any AI or any tool. If you don't know your data, and if that data isn't clean, your outputs aren't going to be clean either.

LIBBY KOEHN · ELIXIR

We need to baseline what we're doing today so we can prove the ROI. We need to clean up our data and make sure we have good integration into it. An AI strategy doesn't help much when the batch publication process still runs on manual downloads and data extracted and manipulated in a CSV file. Sanitize the data, enrich the experience, and that's where you start seeing high-value return on investment.

09 Your Vendor Checklist

Five questions for one hour with a CCM vendor

LIBBY KOEHN ASKS

If a buyer has one hour with a CCM vendor to discuss AI, what should they ask?

SAURABH RAJ · QKS GROUP

1. **What is live today?** Which AI features are in production with customers, and which are in beta?
2. **Where does our data go?** Which model provider is behind the AI, and what is stored or cached?
3. **What can AI change, and what is locked?** Disclosures, numbers, and required wording matter most here.
4. **How is AI involvement recorded?** Who changed what, and how AI was used. If an AI-assisted Spanish translation has an issue with terminology or spelling, the AI's role should be in the record.
5. **What happens when AI is wrong?** How errors are caught, who is told about them, and how they get fixed.

AMANDEEP KHANUJA · QKS GROUP

ONE MORE, FOR YOUR OWN TEAM

Before the meeting, agree internally on what problem you are solving and how you'll measure it. Then go into the vendor conversation.

LIBBY KOEHN · ELIXIR

These are exactly the hard lines buyers need to draw. If a vendor stumbles on data privacy, migration, or deterministic compliance questions, they're selling an AI wrapper, not a regulated communications platform.

Look for proof of privacy and proof of the data lifecycle. Look for AI that natively tackles legacy templates and migration. Look for compliance proof. It's not all hype, but AI has to unlock ROI, optimize the experience and the workflow, and make things faster and more efficient, while keeping the record where you can own it and manage it.

Don't buy the demo. Ask for the proof.

Provable privacy, deterministic compliance, and a human in the loop at every step.



If there's a problem in a communication, the AI is not going in front of a judge, a jury, or a committee. ***The business has to go in front of them and explain it.***

Hart Johnson

Executive Managing Director, DevSecOps, Elixir Technologies

Keep the conversation going

Elixir's AI is built on the same principles the panel described: provably private, zero retention, deterministic knowledge, and a human in the loop. See the architecture behind it in our white paper.

Read: Provably Private AI as Infrastructure →

ABOUT ELIXIR TECHNOLOGIES

For 41 years, Elixir Technologies has helped organizations in regulated industries create, manage, and deliver customer communications. Its cloud-native customer communications platform serves insurance, healthcare, banking, and other regulated markets, with provably private, zero-retention AI built in where it unlocks value.

ABOUT QKS GROUP

QKS Group is a global advisory and research firm. Its analysts study technology markets, including customer communications management, to help enterprises select technology and help providers identify growth opportunities.

elixir | Cloud