Quality Forward

WHITEPAPER

How QA Teams in Life Sciences Can Leverage AI to Modernise QMS

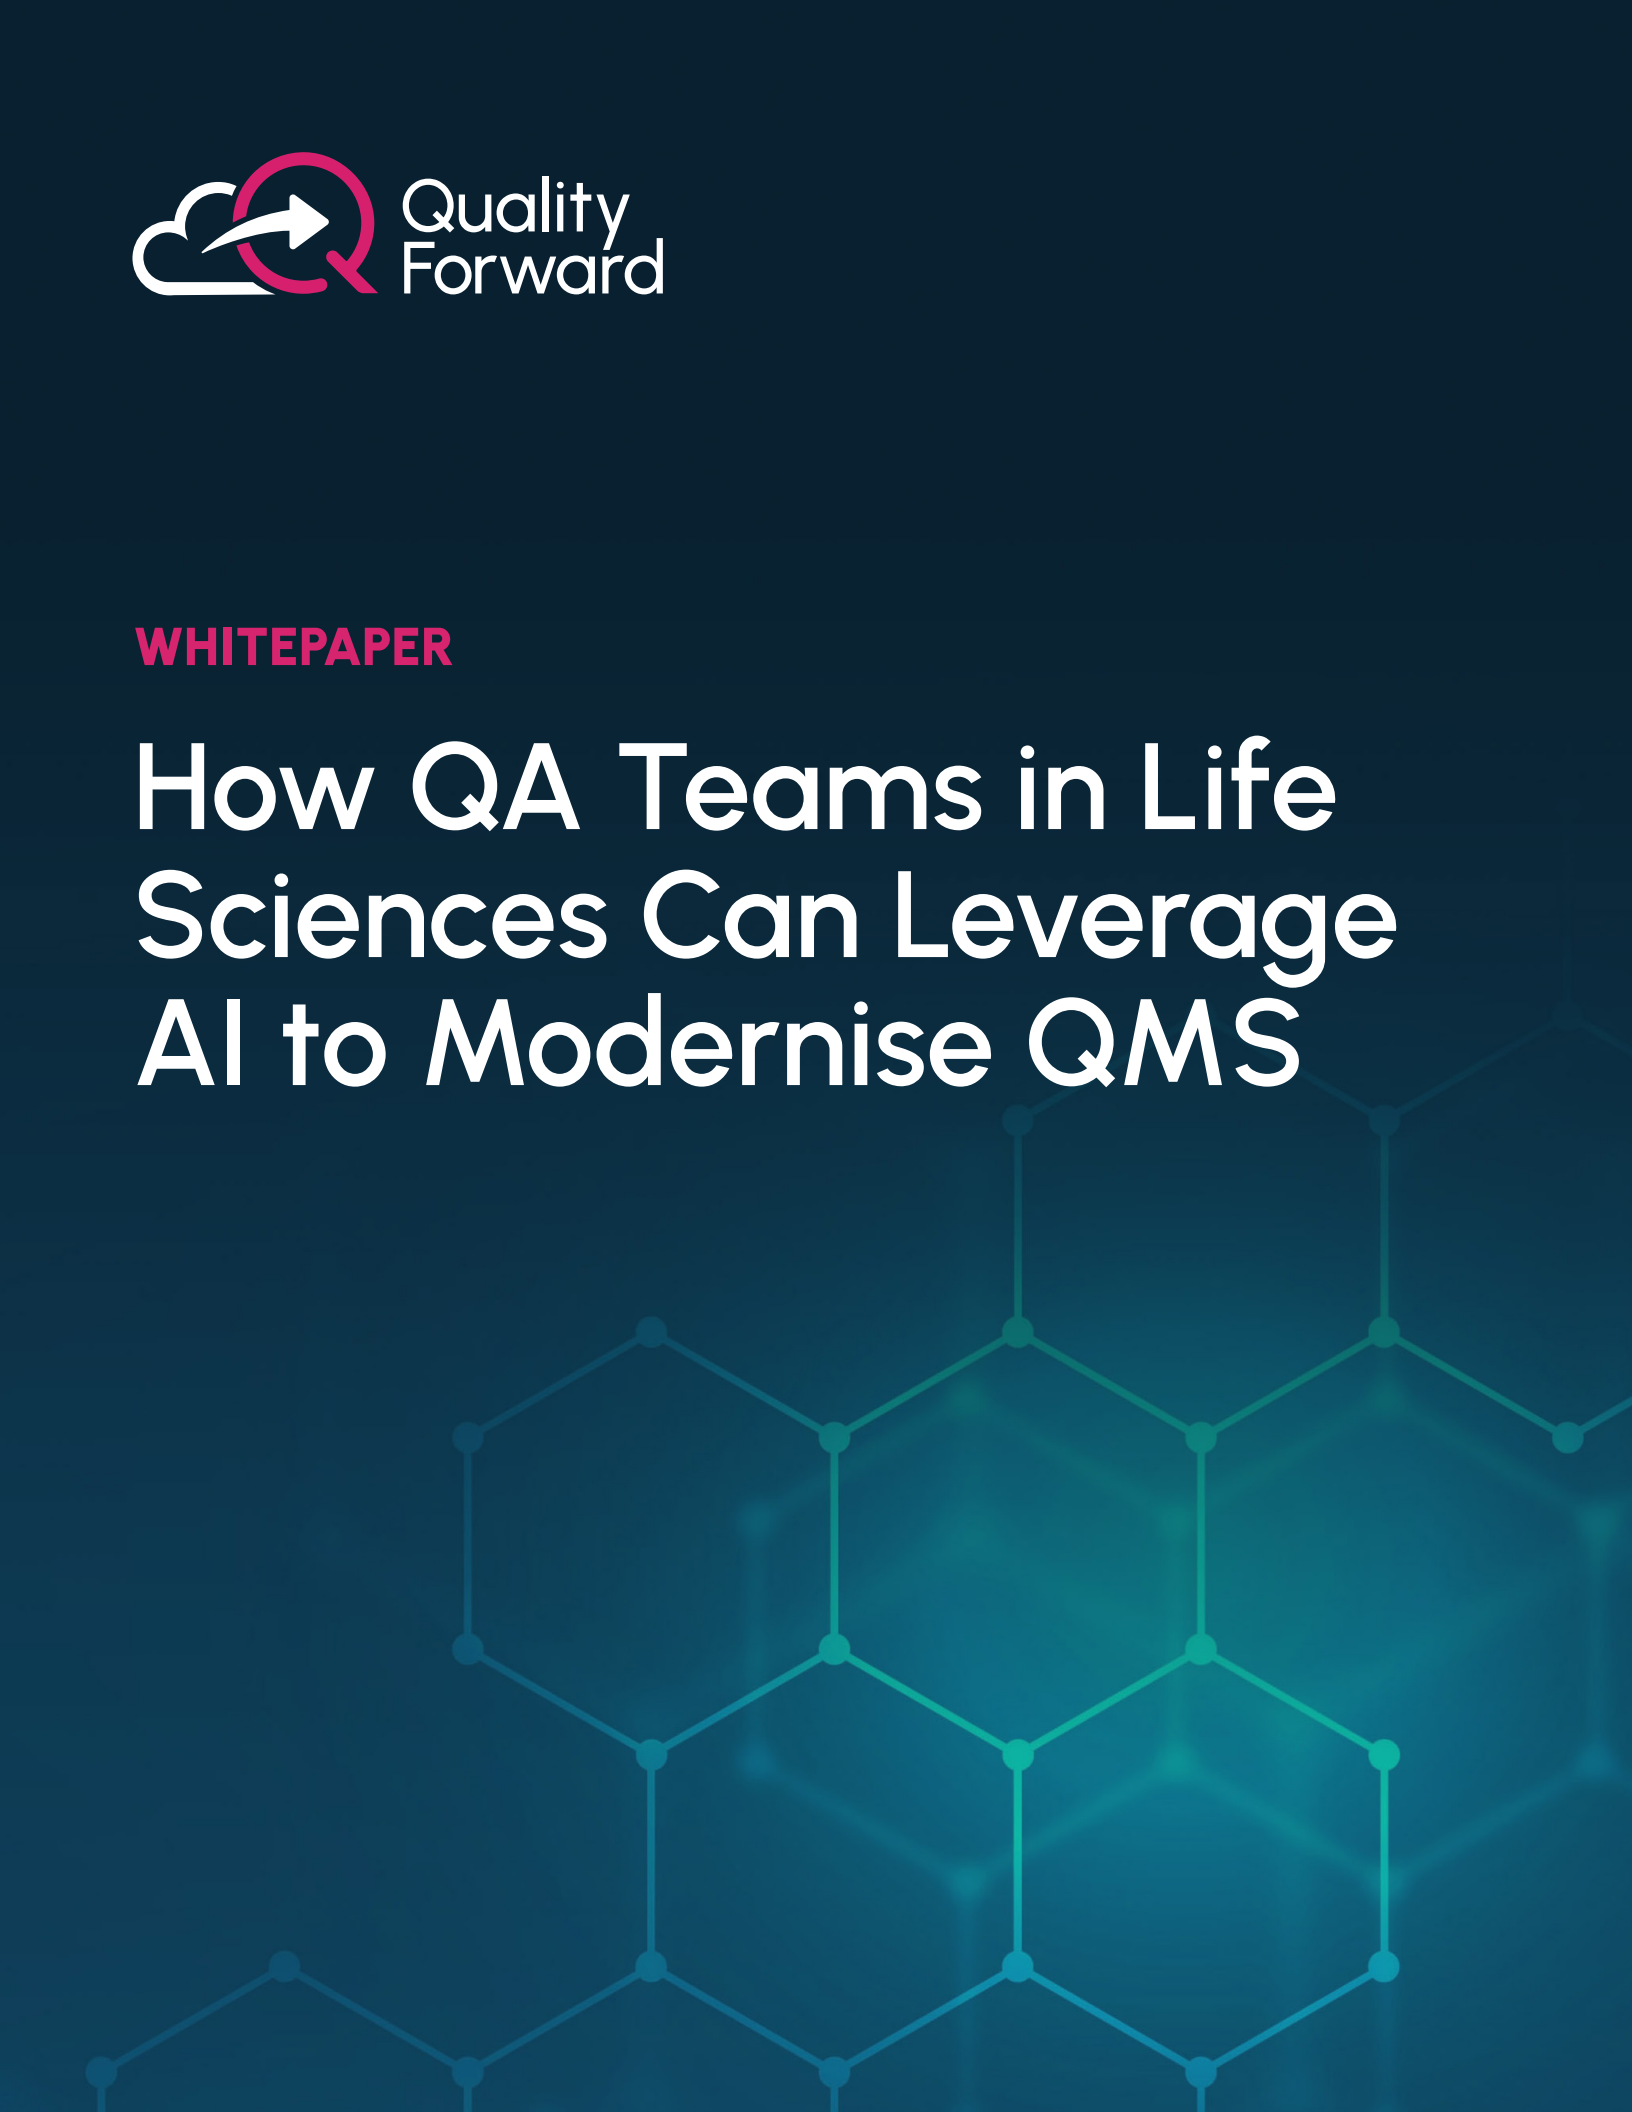

Table of Contents

1. Why QMS in Life Sciences Needs AI Now

Quality systems sit at the centre of every life sciences company. They carry the weight of proving that products are safe, processes are controlled, and teams are ready for inspection at any moment. But too often, these same systems are slowed down by manual steps, scattered data, and old workarounds that should have been replaced years ago.

Manual Workflows Create Bottlenecks

Legacy QMS setups depend heavily on paper-based or semi-automated processes. Each deviation or CAPA often requires multiple hand-offs between teams, increasing cycle times and slowing down resolution. This manual effort leads to hidden bottlenecks that not only drain productivity but also tie up your most skilled QA professionals with repetitive tasks instead of letting them focus on strategic quality improvements.

> On average, 40–60% of a QA team's time is spent on low-value administrative tasks.
>
> - LNS Research, 2022

Rising Pressure from FDA/EMA Inspections

The scrutiny from regulators such as the FDA and EMA has intensified in recent years. The increased pace of audits, surprise inspections, and heightened expectations for data traceability mean that any delay, oversight, or human error can result in costly compliance risks, warning letters, or even operational shutdowns. The traditional QMS approach makes it difficult to respond swiftly with clean, audit-ready records.

Growing Complexity Makes Trends and Anomalies Hard to Spot

Quality teams today must sift through vast amounts of data scattered across multiple systems. Batch records, deviation logs, supplier documentation, training records, equipment calibrations, and more. This makes it challenging to see patterns or spot anomalies early enough to act. Without advanced analytics and automation, your team spends more time hunting for information than using it to make proactive quality decisions.

Modern life sciences companies simply cannot afford a QMS that traps quality data in silos and forces people to spend their days on work that AI could automate in seconds. Organizations that move to an AI-powered QMS gain the visibility, speed, and resilience needed to meet today's strict regulatory expectations and tomorrow's competitive demands. Manual QMS workflows slow teams down. AI is the key to breaking bottlenecks and giving Quality back its time.

2. How AI Enhances QMS Functions

Integrating artificial intelligence into Quality Management Systems is transforming how life sciences companies manage risk, support compliance, and improve process efficiency. Rather than replacing human oversight, AI augments it by embedding intelligent decision support directly into core workflows.

Smart Routing of Deviations and Complaints

One of the most valuable uses of AI in a QMS is the intelligent routing of deviations, non-conformances, and customer complaints. Instead of static workflows or manual assignment, AI agents analyze context, historical data, and current workloads to determine the best route for each quality issue. Prebuilt agents, like a Problem Investigator or Incident Trends Analyzer, help uncover hidden root causes and recurring patterns. Once identified, these issues are dynamically directed to the right experts, reducing cycle times and making sure problems are resolved by the people best equipped to fix them.

Predictive Insights for CAPA Prioritization

Advanced predictive models strengthen the CAPA process by highlighting where action will have the greatest impact. AI tools can analyze large sets of deviation reports, audit findings, and complaint histories to calculate risk scores and forecast the likelihood of recurrence or escalation. This means teams can prioritize corrective and preventive actions based on real data rather than gut instinct, tackling the issues that pose the highest compliance and patient safety risks first. AI also helps teams look for systemic trends that might otherwise go unnoticed until they become a larger issue.

Auto-classification of Quality Events

Manual classification of events is inconsistent and time-consuming. AI agents use natural language processing and historical resolution data to auto-classify deviations, change requests, or non-conformance reports. This allows for consistent application of categories and faster downstream reporting. When new quality records are logged, the system suggests or applies the right classification automatically and links each record with the most likely next steps based on prior outcomes.

> Companies using AI for predictive risk modeling have seen a 40% reduction in repeat deviations and 25% fewer CAPAs per audit cycle.
>
> **- McKinsey Life Sciences QMS Benchmark, 2022**

Automating Repetitive Activities and Best-Next-Actions

AI in modern QMS goes beyond analysis—it takes action. Agents can automate tasks like drafting plans or assigning follow-ups, while a no-code studio lets users customize and combine agents to handle complex workflows. For instance, a Change Request Planner can auto-generate test plans and track updates, with triggers launching multiple agents to save hours of manual work.

These tools free up quality teams to focus on complex analysis, cross-functional collaboration, and strategic decisions. By combining trends, predictive models, and automation, companies can shorten cycle times, reduce errors, and better apply human expertise where it counts most.

3. Power of ServiceNow for QMS

ServiceNow gives life sciences companies an enterprise platform where AI is an integral part of how quality work gets done. One major advantage is its no-code Creator Studio. This workspace comes preloaded with tested AI agents so teams do not have to start from scratch. For example, a quality team can clone an agent that investigates common complaints and then adjust the logic to suit how their specific CAPA process works. Instead of creating rigid flows, the system supports flexible routing, decision trees, and learning models that adapt to changing risks and data.

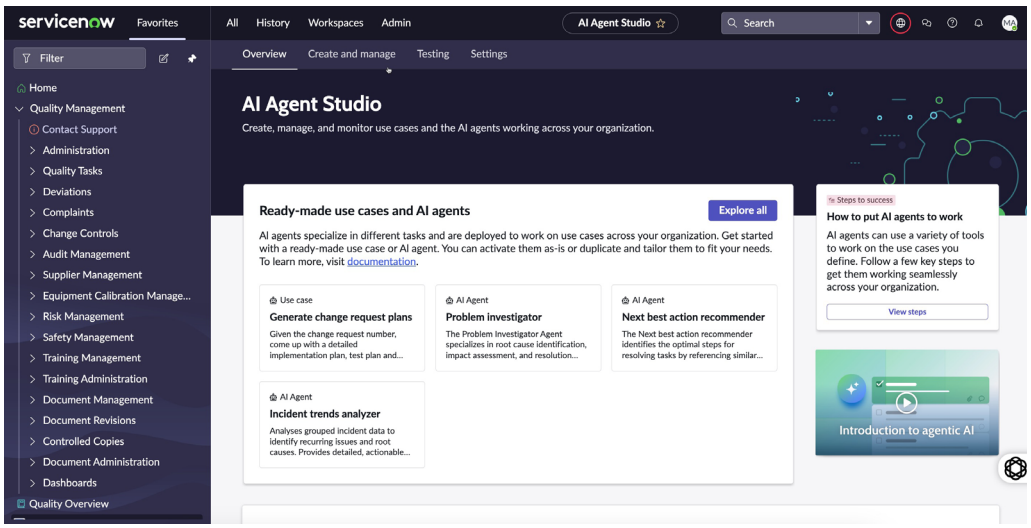

Figure 1: ServiceNow AI Agent Studio. The no-code Creator Studio comes preloaded with ready-to-use AI agents. Teams can clone or modify these agents to build new automations from scratch, without custom coding.

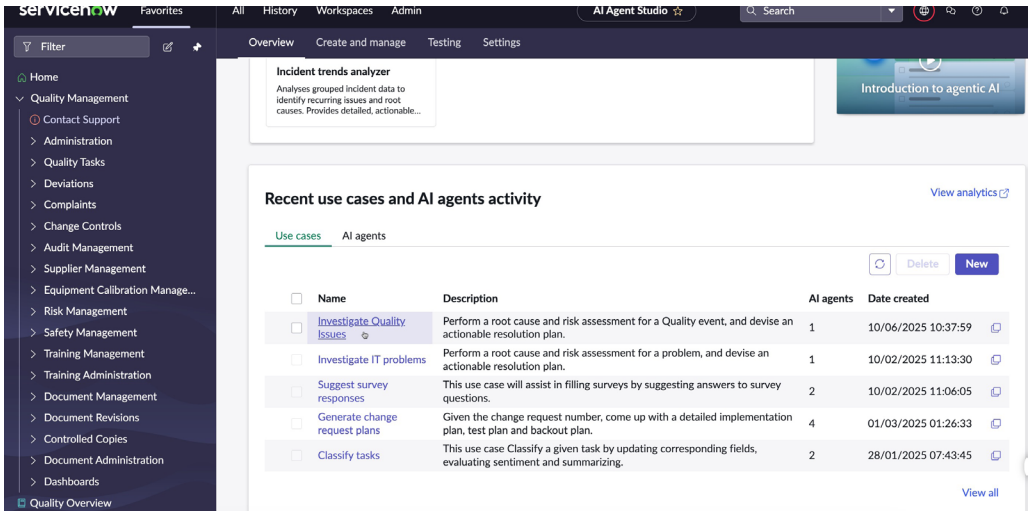

Figure 2: Cloning AI Agents to reuse across quality use cases. Teams can duplicate proven agents, adjust them, and create new ones without starting from scratch.

The Creator Studio allows quality teams to deploy AI agents that follow best practice, then modify or expand them as new products, sites, or suppliers come online. The ServiceNow Creator Studio shows how defining an AI agent use case works like writing a simple prompt. A large language model (LLM) interface means business users can set up, adapt, and launch AI agents with no coding or syntax required.

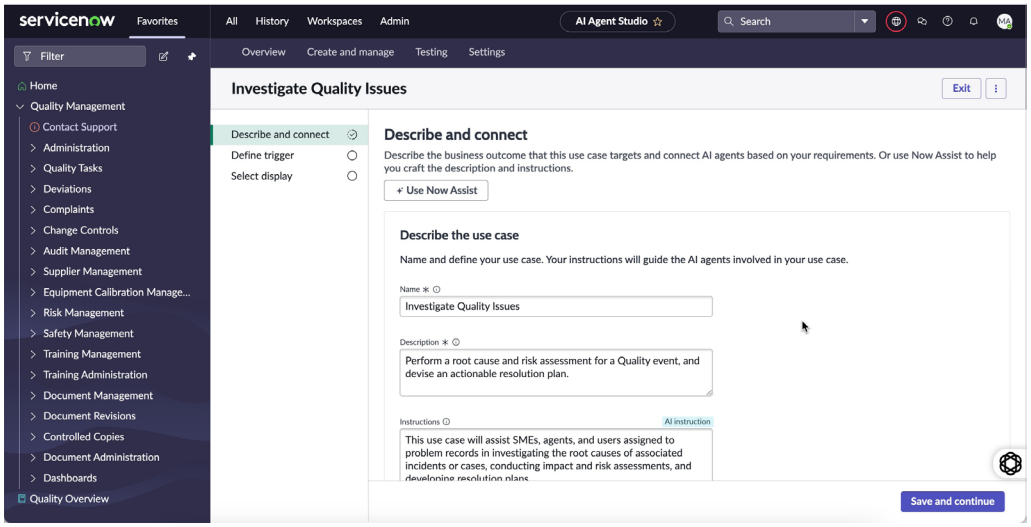

Figure 3: Similar to GPT, the better the prompt, the better the outcomes when your agent is activated.

Another area where ServiceNow stands out is orchestration. Companies often rely on multiple systems to handle deviations, complaints, and supplier issues. This fragmentation slows down the work because people have to switch tools, chase updates, and duplicate entries. On ServiceNow, orchestration means AI agents and workflows run side by side on the same platform. A deviation logged in one region can automatically trigger routing, escalation, and investigation steps, with each action recorded in context. The system can launch multiple agents to handle parallel tasks, which reduces the risk of dropped handoffs.

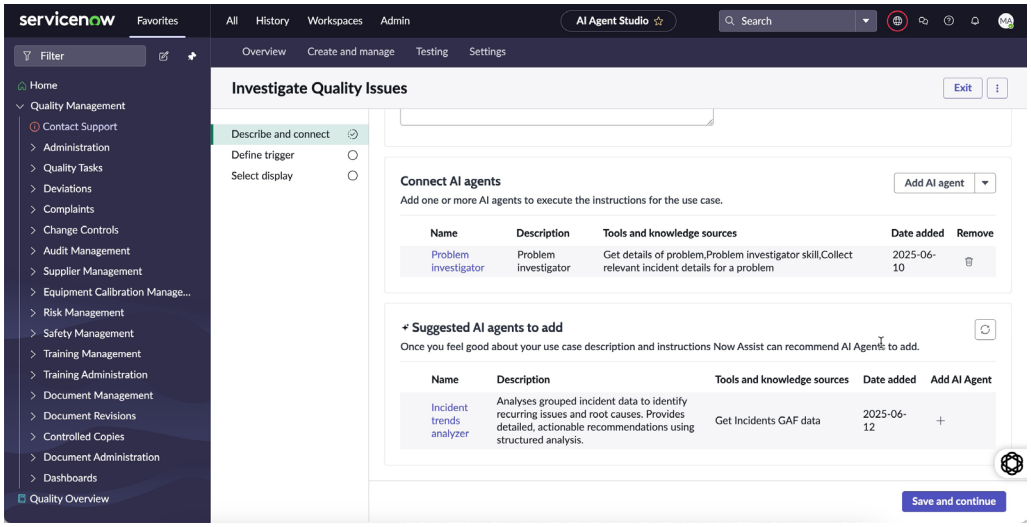

Figure 4: Example of ServiceNow orchestrating multiple AI Agents for a single use case, enabling teams to stack and combine capabilities for more effective issue resolution.

In practice, teams can manage an entire library of use cases simultaneously, adjusting, duplicating, or expanding AI agents to cover different quality and operational needs.

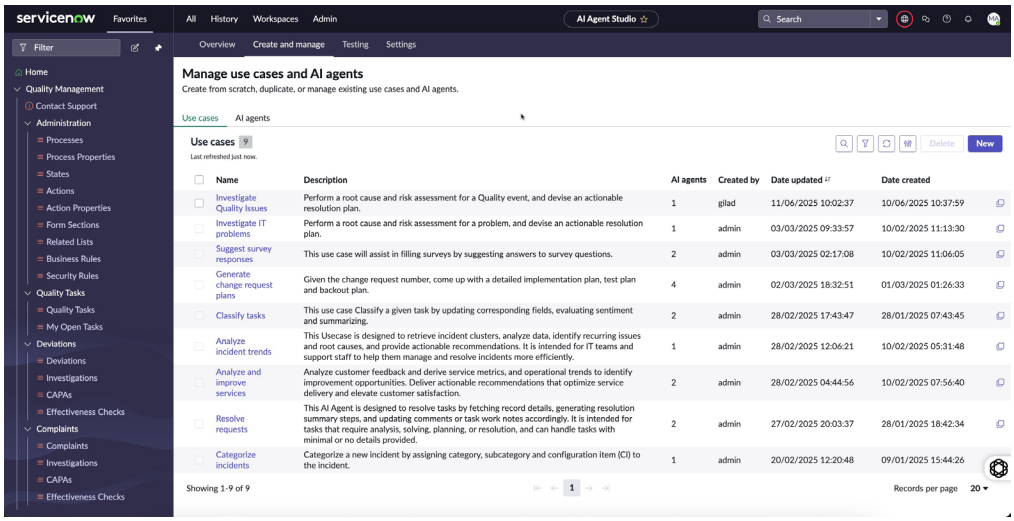

Figure 5: Multiple active AI use cases in ServiceNow's AI Agent Studio. This shows how teams manage a wide range of processes, from root cause investigations to service improvements, all within the same no-code environment.

The following drill-down shows this in action for a Change Request process, where four AI Agents work together to deliver a unified plan.

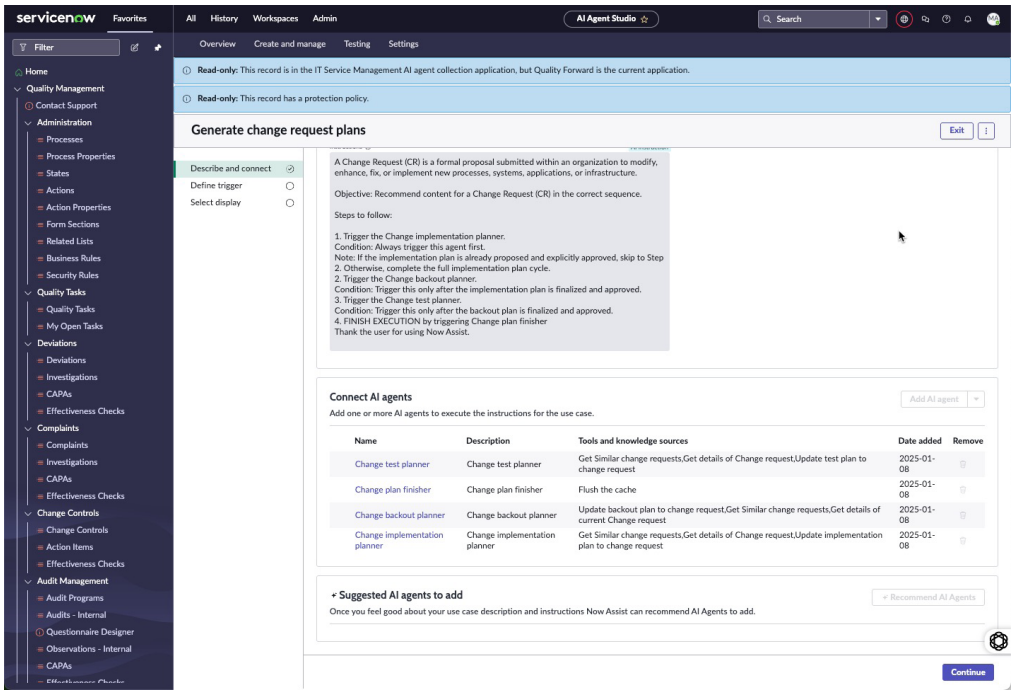

Figure 6: Drill-down view of the Change Request process using 4 unique AI Agents in parallel to deliver a single, final output. ServiceNow can also suggest additional agents or help create new ones from scratch.

Teams can create and activate triggers that kick off the AI Agents in real time while people are working. This means that something as simple as an engineer or operator submitting a request to modify a manufacturing process can automatically activate the right agents, execute tasks, and deliver outputs that support the associated activities.

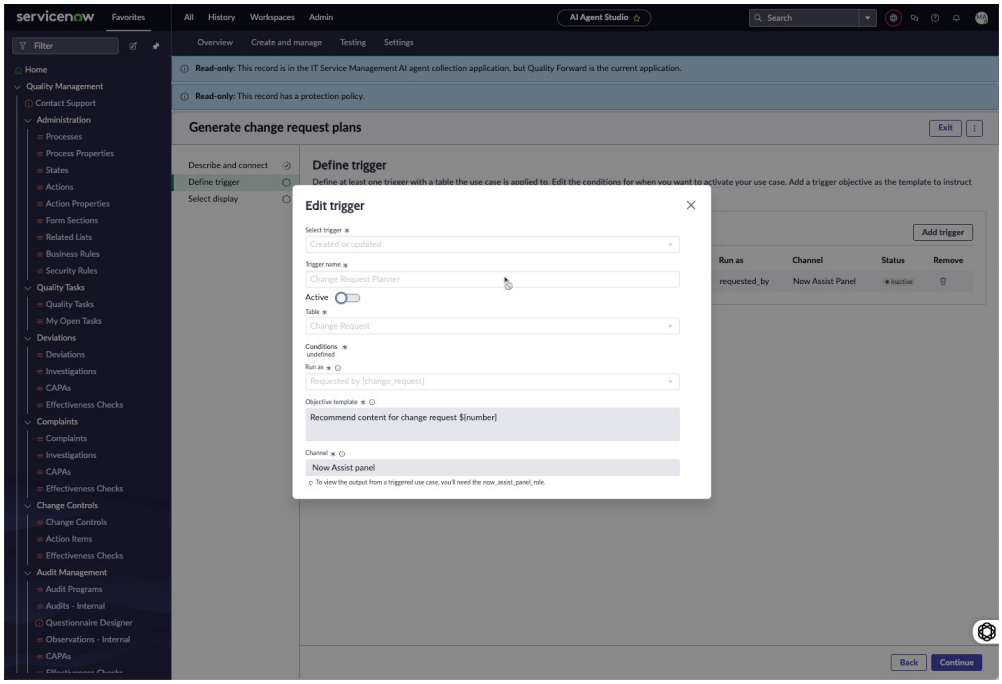

Figure 7: Real-time triggers can be defined and activated to launch AI Agents automatically, keeping work flowing while users focus on higher-value tasks.

Security and access controls are built at an enterprise level, not layered on top. Teams can assign user permissions, track approvals, and hold detailed audit trails without relying on manual sign-offs scattered across folders or spreadsheets. This is all-important for regulated industries where each piece of the record must be traceable. This example illustrates how AI can handle tasks like pattern detection and prediction for CAPA cycle times, but still keep final sign-off in human hands when it matters.

Uptime and performance are other points often overlooked. ServiceNow's single-instance cloud architecture provides real-time data access globally, so quality data stays consistent whether teams are based in the EU, US, or Asia. This guarantees that AI models and agents work from the same live data, reducing version-control headaches and delays during inspections or product launches.

What makes ServiceNow unique for a life sciences QMS is that it does not force companies to choose between pre-built automation and flexibility. The Creator Studio allows quality teams to deploy AI agents that follow best practice, then modify or expand them as new products, sites, or suppliers come online. It is not just about automating small tasks, it is about layering intelligence into every stage, from smart routing of deviations to predictive risk scoring that prioritises what needs human attention first.

This combination of enterprise-grade configurability, built-in security, a robust workflow engine, and native AI means teams can automate what makes sense while staying flexible as products and sites grow. With proven reliability and guaranteed uptime across regions, ServiceNow's approach keeps life sciences QMS data connected, accessible, and inspection-ready everywhere it is needed.

4. How Quality Forward Brings It All Together

Quality Forward was designed to make advanced AI tools feel practical, not intimidating. When building entirely on the ServiceNow platform, the system uses proven enterprise-grade security, scale, and configurability as its base, but shapes it specifically for the demands of life sciences quality teams.

One of the biggest advantages is how easy it is to adapt. Quality Forward takes the no-code designer in ServiceNow and makes it useful for everyday quality workflows. If teams need a new AI agent to handle a different type of complaint or adjust how deviations are routed, they can clone what works and fine-tune it without waiting on outside developers. This keeps changes under control and fully visible, which is exactly what a regulated environment needs. These AI agents can classify quality events automatically, prioritise CAPAs based on real data, and highlight patterns that might not show up in routine reports. Each insight feeds back into the daily work, helping people make faster, better decisions without losing traceability.

What sets Quality Forward apart is how it keeps compliance at the centre. The same platform handles security, audit trails, and user access controls, so every update remains within a validated framework. When an AI recommendation is made, the system records why it happened, who approved it, and what action was taken.

Because it all runs on ServiceNow's cloud, there's no extra overhead to manage. Teams work in one place, data stays connected, and adding new workflows doesn't mean rebuilding what already works. For life sciences companies balancing rapid change with strict oversight, this balance of flexibility, control, and real-time AI support makes Quality Forward stand out.

5. Outcomes Life Sciences Teams Are Seeing

Companies adding AI and machine learning (ML) to their QMS report measurable gains that traditional systems simply can't deliver at the same speed. When repetitive tasks are automated and quality data is analyzed in real time, QA teams gain time and insights that help keep operations inspection-ready.

- Cycle times for deviations and CAPA drop by 30% to 50%. ML models flag similar past incidents, suggest likely root causes, and help route tasks to the right people sooner.
- Predictive algorithms use trend data to identify where repeat issues are likely to occur, so preventive actions can be taken before they impact production or patient safety.
- Anomaly detection runs continuously on quality data streams, highlighting outliers in test results, supplier lots, or manufacturing runs. This reduces the risk of hidden compliance gaps.
- Automated audit trails link every change, sign-off, and training record to the relevant SOPs, which means inspection preparation becomes routine.
- Smart classification and routing of complaints and deviations eliminate manual re-entry, reduce human error, and cut down on bottlenecks in review cycles.
- Better traceability connects each quality event to its associated documentation, approvals, and training. This gives auditors a clear line of sight from issue to resolution, improving trust in the system.

This is where Process Mining comes in. Helping teams visualize their workflows, spot trends, and sustain these improvements over time.

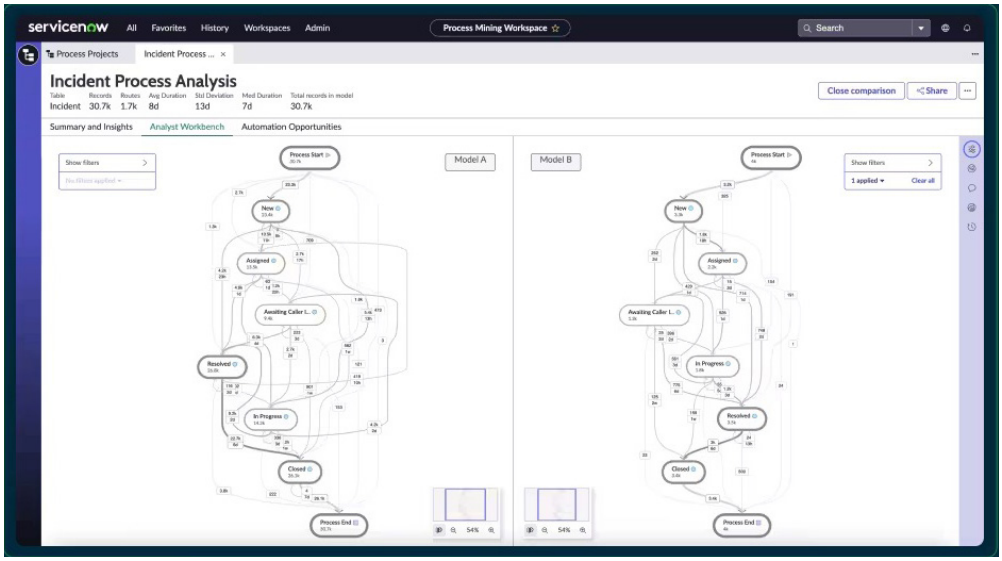

Figure 8: Process Mining AI capabilities help quality teams visualize workflows, uncover trends, and detect anomalies or bottlenecks to sustain improvement over time.

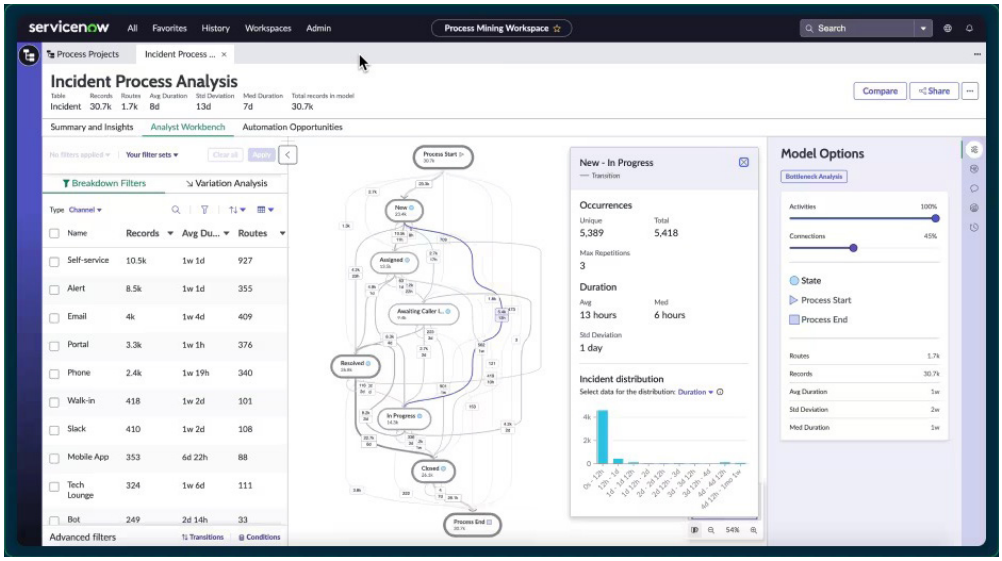

Figure 9: Finalized processes can be analyzed using AI-powered Process Mining tools to identify performance trends, anomalies, and improvement opportunities.

Using Continuous Process Mining to Sustain Measurable Outcomes

Process mining in a modern, AI-enabled QMS is a continuous practice that helps quality teams monitor real operations, detect issues, analyze root causes, and drive meaningful improvements. It is not a single AI action, it is a connected cycle that supports both people and automation over time.

Objectives supported by process mining include:

reducing case reassignments, increasing self-service rates, reducing Mean Time to Resolve (MTTR), lowering the number of reopened cases, improving response times, deflecting unnecessary cases, maintaining a smaller backlog, increasing productivity, and avoiding SLA breaches or penalties.

Detection activities cover:

performance analytics and reporting to find outliers with KPI Signals, identify the most urgent cases with Spotlight, forecast proactive trends with time-series analysis, compare KPIs over time with snapshotting, and track SLA or survey performance. Unstructured text analytics help derive problem symptoms. Predictive Intelligence clusters similar customer cases and uncovers knowledge management gaps. Benchmarking shows performance compared to industry or internal targets.

Analysis includes:

process discovery that maps every process step across systems, process highways showing the most common routes, conformance checks to understand deviations, and rule findings to resolve compliance breaches. Enhancement features include advanced statistics to detect process variance, AI-powered root cause analysis, cluster analytics to understand what cases cause the most issues, automation discovery for repetitive steps, and bottleneck analysis to compare slow-performing transitions.

Improvements come through:
predictive case assignments, Virtual Agents for self-service, RPA to automate routine tasks, skill- and capacity-based work dispatching with Advanced Work Assignment, regular workforce performance reviews, better scheduling to match workload demand, and removing unnecessary approval steps with a flow designer.

This cycle should be continuously monitored, refined, and reviewed so that quality teams can respond quickly to changes in operations, spot trends and anomalies early, and keep processes inspection-ready.

The visual below illustrates how Detect, Analyze, and Improve activities connect in a continuous cycle to sustain measurable QMS outcomes.

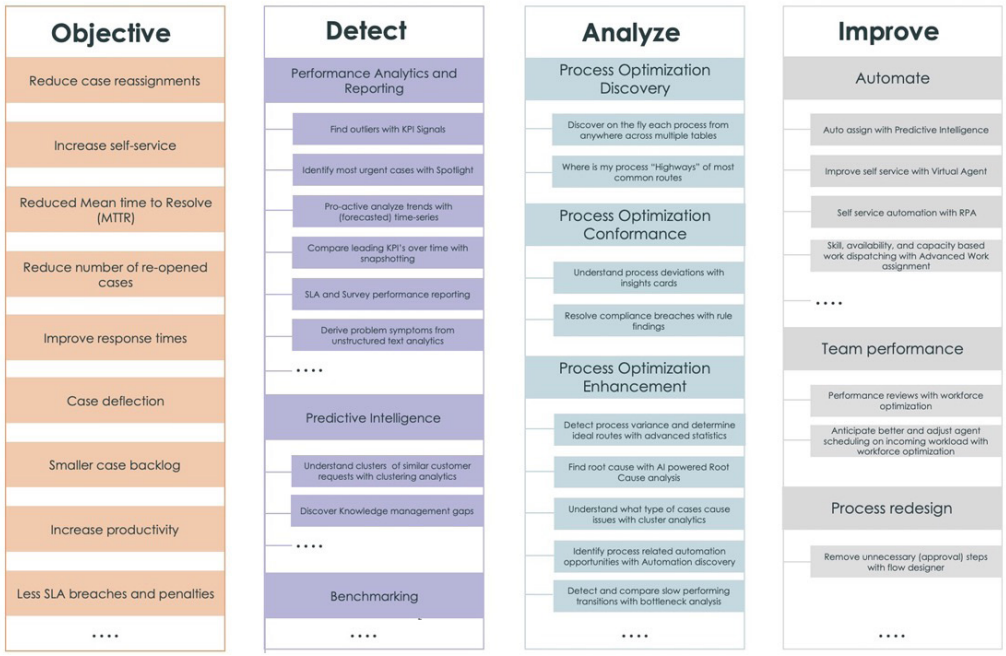

Figure 10: *Continuous Process Mining framework. Seeing the entire objective–detect–analyze–improve loop helps teams understand that process mining is not just a single AI action but an ongoing practice that must be monitored and refined for maximum impact.*

Teams can even set up their process mining objectives and KPI dashboards at project creation to keep improvements on track.

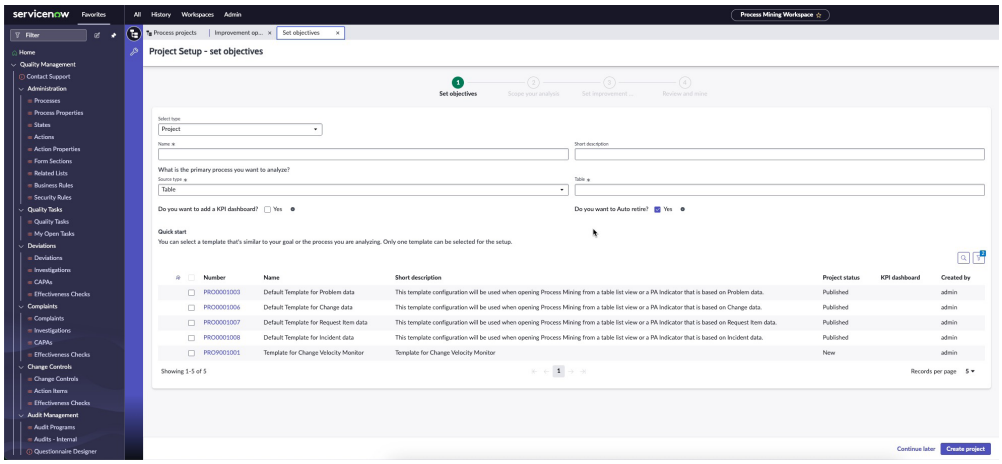

Figure 11: *Project setup for process mining. Example of how teams can configure objectives and dashboards at the start, ensuring continuous measurement and improvement.*

6. Getting Started

Moving to an AI-enabled QMS requires a platform that fits how regulated work gets done, and that can adapt as the business changes. Teams need confidence that AI and machine learning are practical, explainable, and linked directly to daily quality tasks.

What to look for in an AI QMS:

- Built-in AI and ML that do not rely on bolt-ons or extra modules
- No-code configuration so quality teams can adjust workflows as needed
- Clear traceability between AI recommendations and documented actions
- Complete audit trails that show how decisions are made
-

Questions for vendors:

- How is AI validated to meet GxP and industry standards?
- Who configures AI models and workflows? Does it need IT every time?
- What proof is there that AI improves cycle times or issue detection?
- How is ongoing model training handled as your data grows?

How Quality Forward helps:

Quality Forward uses the ServiceNow platform to connect enterprise-grade performance with compliance-ready AI. Pre-configured workflows, predictive tools, and smart routing are built in from day one. Teams use proven best practices without coding or custom builds. This approach shortens the time to deployment, keeps data traceable, and supports quality operations as they grow.

servicenow

Quality Forward, built on ServiceNow, combines powerful digital workflows with a purpose-built eQMS for regulated industries.

YOKOGAWA

Quality Forward is proud to be backed by Yokogawa, a global leader in industrial automation and digital transformation.

Founded in 2017, Quality Forward was established to solve a critical challenge in life sciences: enabling organizations to migrate from outdated, on-premise quality systems to a secure, cloud-native eQMS; without disrupting existing workflows or losing historical data. The solution combines a best-in-class digital experience with AI-driven insights, automated validation, and real-time reporting. Built on the ServiceNow platform, it empowers quality teams to digitize, integrate, and manage critical quality processes; deviations, CAPAs, audits, training, document control, risk, and more, with traceability and ease.