smarteeva

CASE STUDY · RISK MANAGEMENT

They were building risk assessments in Excel. Now a full report generates in 8 minutes.

8 min
TO GENERATE A FULL RISK ASSESSMENT REPORT

700+
HAZARDS ASSESSED AGAINST LIVE COMPLAINT DATA

9
ORCHESTRATIONS RUNNING ONE REPORT

2 mo
FROM FIRST WORKSHOP TO PRODUCTION RELEASE

About the customer

A leading MedTech innovator developing advanced insulin delivery systems. As the product portfolio and installed base grew, the risk management file grew with it, until keeping every occurrence estimate current against real-world complaint data became a project rather than a routine.

The problem

Risk assessments were built by hand. Complaint data was filtered and aggregated in Excel, calculations were performed there, and the final report was assembled in Word.

Each report took several weeks of skilled quality engineering time. Because the work was that heavy, the team produced roughly thirty reports a year across a portfolio spanning 320 product categories and around 700 hazards.

That is not a resourcing failure. It is arithmetic.

Checking one occurrence estimate means counting every complaint mapped to a specific hazard, then normalising that count against units in the field. Repeat across 700 hazards and the annual cycle explains itself.

Industry	**MedTech / Insulin Delivery**
Pain point	**Manual risk assessment in Excel and Word**
Hazards	**~700**
Product categories	**320+**
Key outcome	**Weeks to 8 minutes**
Orchestrations	**9**
Platform	**Salesforce-native**
Deployment	**2 months to production**

APPROACH

Put the risk file where the complaints already are

The manual work was never the analysis. It was the assembly. So the deployment targeted assembly and left judgment where it belongs.

Hazards became records

Hazards, hazardous situations and intended use moved into the platform as objects, alongside complaints, investigations and adverse events. Once a hazard is a record, a complaint can point at it.

Coding became the bridge

Complaints were already coded to IMDRF Annex terms for regulatory reporting. That coding now assigns hazard and severity, so the classification work does double duty as risk evidence.

Exposure data came from the source

Install base and distribution volumes are pulled from the systems that hold them, so occurrence is expressed as a rate rather than a raw count.

People kept the decisions

Risk acceptability, benefit-risk determination and mitigation strategy stayed with qualified reviewers. Nothing is issued without approval.

SOLUTION

Nine orchestrations, one report

Smarteeva Orchestra runs the analysis as nine orchestrations, more than 250 configured steps in total, against the customer's own risk file. Smart Documents applies their report template, and AI drafts the narrative sections from the assembled evidence.

01 **Analysis Process · 12 steps**
Sets the reporting period, product scope and inclusion criteria, so every downstream calculation runs against the same defined population.

02 **Calculate Complaints and Events · 36 steps**
Counts complaints and adverse events mapped to each hazard across the reporting period.

03 **Calculate Install Base · 41 steps**
Draws units distributed and in service to supply the denominator, converting counts into occurrence rates.

04 **Describe Root Cause · 44 steps**
Reads investigation findings and drafts the root cause narrative, with the underlying records attached as evidence.

05 **Describe System Defect · 36 steps**
Assembles the device and system effect sections in the customer's own report structure.

06 **Fill Existing Control Measures · 36 steps**
Pulls the control measures already documented against each hazard, so mitigations sit alongside the evidence.

07 **Fill Supporting Documents · 37 steps**
Attaches source datasets and referenced records, so a reviewer can trace any figure back to where it came from.

08 **Generate Chart · 10 steps**
Produces the trend and distribution charts the report format requires.

09 **Generate Risk Hazard Conclusion**
Drafts the hazard-level conclusion for review. A qualified reviewer approves, edits or rejects before anything is issued.

RESULTS

What changed

- **Report generation fell from several weeks to 8 minutes.** The same assessment, built from the same records, assembled by the platform instead of by hand.
- **Reporting frequency is no longer set by effort.** At roughly thirty reports a year, cadence was a function of how long each one took. That constraint is gone.
- **Nine orchestrations and more than 250 configured steps** replaced the manual filtering, aggregation and calculation that previously ran in Excel.
- **Two months from first workshop to production release,** built against the customer's existing hazard model rather than a fixed template.

The finding nobody expected

Automating the assessment surfaced records that manual review had not connected, including investigations that were never linked to their hazard.

This inverts the usual argument for automation. The claim is normally that software does the same work faster. Here the software did the same work faster and found gaps in work that had already been signed off.

That is what happens when a manual process is large enough. Not carelessness, volume.

Seven hundred hazards, 320 product categories, and a person making the same judgment call several thousand times in a row.

WHAT THIS DOES NOT CHANGE

Judgment about risk acceptability, benefit-risk determination and mitigation strategy remains with qualified people, and no regulator would accept anything else. What automation removed is the weeks spent assembling evidence instead of interpreting it.

See it against your own risk file

Thirty minutes, your hazard model, your complaint data. We will show you where your current occurrence estimates sit against what the field is reporting.

Book a risk management demo · calendly.com/jaime-castillo-smarteeva