



CRF CERTIFICATE COURSE · DENVER · AUGUST 10, 2026

AI Applications in Credit & Order-to-Cash

From hype to **execution** — in one hour

Ali Kidwai · Senior Director — Product, Implementation & Engineering · Bectran



BEFORE WE START

Three things vendors have said this year

- A "Our AI reduces bad debt by **40% in 90 days.**"
- B "AI cuts financial statement review from **30 minutes to 3.**"
- C "Autonomous agents will **replace your credit team's decisions** by 2027."

One is real. One is half-true. One is hype. **You'll vote at 9:25.**



HOW THIS HOUR WORKS

Less lecture. More screen.

3 live demos

Real builds, live on screen — no pre-baked videos.

5 audience calls

You vote, you edit the rules, you pick the tone.

6 prompts to keep

Printed handout. Fill in the blanks, run them Tuesday.

Everything you'll see runs on data you could assemble **this week**.

Real vs. hype — in every corner of O2C

Disputes

- ✓ Classification & triage
- ✓ Evidence checking

- ✗ Auto-settling disputes

QUESTION

USE CASE 1 · CREDIT DECISIONING

Where does your credit policy live?

A binder. A PDF nobody opens. Three people's heads — one of whom retires in March.

And how many of last quarter's decisions actually followed it?

- LIVE DEMO 1 · POLICY-AS-CODE

Let's make the policy do the work

30-page policy

Real structure, fake company.
Nothing pre-trained on it.

1 page of collector notes

The tribal knowledge that never
made it into the binder.

5 pending applications

Decided live, each citing the
policy section that drove it.

Then **you** change a rule — and we watch the decisions change.

WHAT JUST HAPPENED · DEMO 1'S ACTUAL OUTPUT

Every decision cites the policy — then **you** changed the policy

APPLICANT	DECISION	BECAUSE
Meridian Supply	Approve	Sec. 4.2 · trade refs ✓
Bluff City Materials	Decline was Approve	Sec. 4.6 — your new NSF rule
Torres & Sons	Review	Sec. 6.3 · guarantor needed
Lakeshore Industrial	Decline was Review	Sec. 4.6 — your new NSF rule

One rule edited by this room — **two decisions flipped**. Full audit trail, zero black box.

TAKEAWAY

Every decision traces to a sentence **your team wrote.**

Auditability isn't a feature bolted on afterward — **it is the output.**

Collector notes → written policy → executed decisions. Tribal knowledge just became an asset.

HANDS UP · APPROVE OR DECLINE?

You make the call

Account 1

\$18M revenue · 22 yrs in business

D&B 42 (**weak**)

Pays 8 days late, **every** invoice

Same pattern for 6 straight years

Account 2

\$4M revenue · 3 yrs in business

D&B 71 (solid)

Perfect payment history

Orders doubled in 60 days

Account 3

\$45M revenue · 15 yrs

D&B 58 · slowing 2 days/mo

Long-time top-20 customer

New ownership last quarter

You make the call

Clean pay since restructuring

Your rep: "something feels off on calls"

Progress

THE SCORE-ONLY AI'S VERDICT

It declined **every account** this room just approved

Account 1**DECLINED**

"D&B 42 — high risk"

Missed: 6 years of **predictable** lateness**Account 3****DECLINED**

"Score deteriorating"

Missed: 15-year relationship context

Account 4**DECLINED**


"Bankruptcy on record"

Missed: clean pay since 2019

The score wasn't wrong. It was **alone**.

THE BACK-TEST · 962 REAL CREDIT DECISIONS

What context does to AI agreement

**100%**

of human-approved "high-risk" accounts were declined by score-only AI

57%

of those same accounts paid clean — your gut was right

THE LESSON

Your gut was right — it just **doesn't scale**.

AI's real job in credit is **context assembly**, not judgment replacement.

It reads everything. You decide.

- LIVE DEMO 2 · THE PAPER GAUNTLET

This is what "AI-ready data" actually looks like

Crooked scan

Photographed on a desk, at an angle

Handwriting

Margin notes from their controller

Missing items

No interest expense line anywhere

Coffee ring

Genuinely. On the balance sheet.

Watch the clock — full extraction and ratio analysis, live. **Target: under 2 minutes.**



AUDIENCE VOTE

One of these bank statements is **doctored**

Statement A

First National · Nov 2025
Ending balance \$284,116

Statement B

First National · Nov 2025
Ending balance \$302,516

Hands up for A... hands up for B. **Now let's ask the machine.**



WHAT JUST HAPPENED

Extracted from the messy scan — 94 sec Document integrity — Statement B

RATIO	FY25	HEALTH
Current ratio	1.42	Green
Quick ratio	0.71	Amber
Debt / equity	2.9	Red
DSCR	1.1	Amber

⚠ TAMPERING SUSPECTED — 3 signals

1. Ending balance $\neq$ txn sum ($-\$18,400$)
 2. Font inconsistency — line 14 figure
 3. Modified 11 days after issue date
- EIN differs from Sec. of State filing

20–30 min → 2–5 min

per financial statement, in production — Carrier Enterprise (~90% reduction)

"We've effectively eliminated human error in our calculations." — Ralph DeMassa

- LIVE DEMO 3 · THE INBOX

Your inbox is an early-warning system nobody is reading

Tone shifts precede payment problems. At portfolio scale no human can read it all. **AI can — and we'll watch it catch one, then do the same for a dispute.**

Tone turned 6 weeks before the first late payment



Dispute triage — same pattern, different queue

Verdict: INVALID — POD signed for full quantity; shortage claim contradicts their own receiving doc. Response letter drafted, evidence attached.

IF YOU REMEMBER ONE SLIDE

**Credit files. Financials. Emails. Disputes.
Remittances.**

**It's all one pattern:
unstructured mess in →
structured, cited, actionable out.**

AI VS. SOMEONE TELLING YOU IT'S AI

Is it AI — or a rules engine in a trench coat?

1 · Bring your messy data

Rules break on the crooked scan. Real AI degrades gracefully — you watched it.

2 · Ask to see it fail

Real AI has confidence levels and a "couldn't read this" list. Fakes are always 100% sure.

3 · Ask "why this one?"

Real AI explains per case — like Sec. 4.6 citations. Rules give the same canned answer every time.

4 · Change something live

We changed a rule mid-demo and reran in seconds. Hard-coded logic needs a release cycle.

Instant + identical output every single time = **a lookup table, not intelligence.**

THE HONEST MATRIX

Two tiers — know which one your problem belongs to

	Tinker — build in-house	Buy — production platform
WHAT	Analysis, drafts, one-off tools — today's demos	Core O2C workflows
USERS	You and your team	Customer-facing — your customers touch it
STAKES	Wrong = redo it	Wrong = bad credit decision at scale
AVAILABILITY	Runs when you run it	Uptime SLAs — 24/7, month-end doesn't wait
TRUST	You eyeball every output	Audit trail · security review · SOC 2 · support
PROOF	A cool demo	Real-world applications in production — reference customers, live volume

THE FIVE QUESTIONS

Which tier does this problem belong to?

- 1 Does a **customer** ever see it?
- 2 Does it need to run at **2am on month-end**?
- 3 Does a wrong answer cost money **before a human checks it**?
- 4 Does compliance need an **audit trail**?
- 5 Does it need to work **10,000 times**, not 10?

Any **yes** → buy tier. All no → tinker away, **today**. The demos you watched are the tinker tier — *that's the point*.

FIVE WAYS THIS GOES WRONG

Every one of these, you just watched us dodge

- **Dirty data first.** The crooked scan worked — but garbage policy in = garbage rules out.
- **Black-box vendors.** If it can't cite a section like Demo 1 did — walk away.
- **No exception handling.** Demos show the clean match. Ask about the other 20%.
- **The pilot-to-production gap.** A 6-minute demo $\neq$ month-end at scale.
- **"AI is confident when it's wrong."** Remember 31%. Verify before you trust — always.

LESSONS FROM THE FIELD

How teams actually get this working

Start small

One dataset, cleaned first. One workflow. Then expand.

Appoint a tinkerer

Someone on the team owns experimentation — officially.

Enterprise plans only

Customer data never touches consumer/free AI tools.

Expect to iterate

The first prompt is never the one that works. That's normal.

Measure before & after

Or the ROI conversation never ends.

Human decides

AI drafts, assembles, flags. Sign-off stays with your team.

VOTE AGAIN

The three claims — now you know

- A "40% bad debt reduction in 90 days" — **HALF-TRUE**. Real vendor, cherry-picked baseline.
- B "30 minutes to 3" — **REAL**. You watched it happen at 9:04.
- C "Agents replace your credit team by 2027" — **HYPE**. Remember 31% — context assembly, not judgment.

That detection skill? **That was the certificate.**



Take-home: **six prompts**

Q&A

