



AI at the Stakeholders:

Business function level review

THE HONEST ASSESSMENT

Eight functions. Where the evidence is real, and where the claims run ahead of it.

FUNCTION	PROVEN TODAY	MATURING, PROMISING	EARLY OR OVERHYPED	RISK TO PRICE IN
Procurement	Spend analysis, contract review	Supplier risk monitoring	Autonomous negotiation	Who answers for a false positive
Finance	Invoices, expenses, and T&E	Audit prep, FP&A narrative	Autonomous financial decisions	Data residency
Sales	Call summaries, CRM fill, outreach	Pipeline risk, deal coaching	Fully autonomous outbound	CRM and contact data access
Marketing	Content drafting, email copy	Segmentation, SEO	Brand-safe content with no editor	Editorial governance
Customer Success	Health scoring, renewal signals	Review and summary prep	Autonomous account management	Telemetry plus client comms
HR	Screening, onboarding Q&A	L&D content	AI-led performance management	Inherited bias, DPA required
Legal	Regulatory monitoring	Contract review, redlining	Legal advice without oversight	Hallucination liability
Operations and IT	Helpdesk, code generation	Process documentation	Fully autonomous IT operations	Your code training their model

FOUR CHECKS FROM THE SESSION

Corporate Data, Corporate Account

A spend cube never goes into a personal login.

Make Legal AI Show Its Source

Document, page, and clause behind every answer.

Own the Code and the Bots

Settle IP and exit in writing before anyone builds.

Check Build Before You Buy

Clean data plus a real taxonomy may beat a new tool.

SIX THINGS PROCUREMENT OWNS

Know which function owns the tool

The team using it rarely judges the commercial risk best.

Apply the function-specific risk lens

One contract review does not cover all eight.

Require a data map before signing

What goes in, where it runs, whether it trains their model.

Connect adoption to payment terms

For anything unproven, buy consumption, not seats.

Build a cross-functional AI supplier register

Every tool, owner, renewal date, and exit term, one place.

Make an honest assessment before approving

Proven, maturing, or early sets how hard you push.

Takeaway:

Do not build an AI process you could not run manually. If the tool goes down, or the output looks wrong, someone has to know the work well enough to catch it.