



BlueCallom COMPETE - Always On. Always Evolving.

SAP and **BlueCallom** in the New Age of Competitive Enterprise AI

| A Strategic Use Case on the Future of Enterprise Competition

Strategic Findings for Executive Review

2025-10-03 — 12:00–21:45 UTC

Introduction

This executive summary consolidates strictly factual insights drawn solely from the provided research reports and the Cross Department Actions document for BlueCallom AG. The research shows that agentic AI and embedded, workflow-aware intelligence are moving from experimentation to mainstream adoption, while enterprises simultaneously demand hybrid deployment, transparent pricing, and strong governance. Incumbents (notably SAP) are infusing multi-agent capabilities across their suites, but customers report readiness, usability, and cost-transparency gaps that shape buying behavior. Regulatory pressures (EU AI Act) and reliability concerns are redefining technical and commercial requirements. These shifts matter for BlueCallom's future strategy—especially where we differentiate on hybrid readiness, out-of-the-box solutions, orchestration, and pricing clarity.

Market Trends Change

Overview: The market is transitioning toward agentic, embedded AI that closes the “insight-to-action” gap, while enterprises seek orchestration to curb AI sprawl. Adoption momentum is strong but tempered by execution complexity and governance demands—setting the stage for platforms that connect agents, data, and controls across the enterprise.

All Findings in Market Trends Change

Agentic AI is moving from pilots to platform standard.

Explanation: Analysts and trade press indicate agentic AI capabilities are on track to be common in enterprise software within the next few years. Leaders increasingly plan to scale agents beyond proofs of concept, reflecting confidence that autonomous, context-aware agents can drive material efficiency and responsiveness.

- By 2028, about one-third of enterprise software is projected to include agentic AI capabilities.
- 61% of business leaders report integrating AI agents with plans to scale usage.
- Roughly half of GenAI-adopting companies could deploy agentic solutions by 2027.

Rationale: The shift validates investment in multi-agent orchestration, persistent memory, and governance. It supports prioritizing out-of-the-box, production-grade agentic workflows to meet near-term demand.

From standalone chat to embedded, workflow-aware AI.

Explanation: The market is moving beyond generic chatbots toward AI embedded in enterprise processes and tools. Embedded AI provides context-rich decisions and recommendations where work happens, improving adoption and outcomes.

- Embedded AI offers real-time, context-aware insights inside existing business tools.
- Shifts from conversational interfaces to proactive, process-integrated actions.
- Minimizes context loss and increases operational impact versus isolated bots.

Rationale: Embedding BlueCallom agents within core workflows (and partner ecosystems) strengthens differentiation and time-to-value.

Orchestration is the cure for AI sprawl.

Explanation: Enterprises face fragmented AI toolchains and siloed departmental experiments. The trend is toward interoperability and centralized orchestration for shared context, unified governance, and consistent outcomes.

- AI sprawl undermines collaboration and governance when tools remain siloed.
- Platform approaches are preferred to align AI usage with business goals.
- Enterprise orchestration hubs are emerging to coordinate cross-functional agents.

Rationale: BlueCallom's Enterprise AI Operating Platform and Network alignment with orchestration positions us to consolidate pilots into scalable enterprise programs.

Closing the “insight-to-action” gap with autonomous execution.

Explanation: Traditional analytics often stall at recommendations; an estimated 73% of insights never get executed. Agentic AI can autonomously act within guardrails, accelerating response times and operational agility.

- 73% of insights fail to translate into action in traditional approaches.
- Agentic AI executes tasks within policies, reducing administrative delays.
- Examples span supply chain adjustments and real-time sales personalization.

Rationale: Investment in multi-agent execution, embedded actions, and governance-by-design directly addresses a high-value market need.

Sources

1. [TechRadar – Trust in agentic AI, scale projections](#)
2. [TechRadar – AI agents and silo risks](#)
3. [TechRadar – Future of intelligent work with agentic AI](#)
4. [TechRadar – Tackling AI sprawl](#)
5. [TechRadar – Embedded AI sets the standard](#)
6. [TechRadar – From insight to action](#)

Customer Behavior Changes

Overview: Enterprise buyers are vocal about deployment flexibility (hybrid/on-prem), feature parity with cloud, pricing transparency, integration guidance, and better usability with richer out-of-the-box content. These preferences strongly influence vendor selection and time-to-value expectations.

All Findings in Customer Behavior Changes

Hybrid/on-prem AI with full feature parity is a top customer demand.

Explanation: User groups emphasize that cloud-only AI disadvantages on-prem estates. Customers want AI embedded where they are today, with equal functionality across deployment models.

- ~Two-thirds of surveyed SAP users prioritize on-site AI independent of the vendor's cloud.
- On-prem customers insist they must not be disadvantaged versus cloud editions.
- Cloud-only AI strategies have triggered visible customer frustration.

Rationale: BlueCallom's hybrid/on-prem runner and parity focus (as per Cross Department Actions) directly align to buyer requirements and unlock near-term pilots.

Buyers insist on transparent, predictable pricing.

Explanation: Transaction- or quote-based pricing creates budgeting anxiety. Customers request clear calculators, guardrails, and published models to prevent overruns.

- Transaction-based licensing (e.g., RPA) can trigger unpredictable costs if not tightly governed.
- Users call for appropriate, transparent pricing and clear roadmaps for AI features.
- SAP AI pricing is quote-based; lack of list pricing adds opacity.

Rationale: BlueCallom's public, usage-based, zero-subscription pricing and ROI/TCO tooling (Cross Department Actions) meet a well-documented market ask.

Customers want richer out-of-the-box content and simpler UX.

Explanation: Users report complexity and limited prepackaged activities in incumbent tools. Enterprises value starter kits, templates, and intuitive UIs to shorten time-to-value.

- Reviewers cite "not easy to use" UIs for enterprise data/AI tools.
- Fewer prepackaged activities/templates compared to leaders slows adoption.
- Desire for consumer-grade UX and preconfigured agents/workflows is recurrent.

Rationale: Expanding BlueCallom's OOTB library and UX modernization (Cross Department Actions) accelerates adoption and contrasts with complexity pain points.

Integration guidance and standards are repeatedly requested.

Explanation: Customers seek vendor-provided reference architectures and standardized rules to reduce integration risk and avoid new AI silos.

- User groups ask for standardized integration rules and best practices.
- Reviewers cite connectivity and integration hurdles with external data lakes/systems.
- Data fabric and agent-to-agent standards are seen as enablers for scalable agentic AI.

Rationale: BlueCallom's integration factory, connectors, and reference architectures (Cross Department Actions) address a critical buyer need and reduce project failure risk.

Sources

1. [CIO – SAP user demands: on-prem/hybrid, parity, transparency](#)
2. [TechTarget – Unpredictable transaction-based costs](#)
3. [TrustRadius – Usability and integration complaints](#)
4. [SAP – AI pricing requires quote](#)
5. [TechRadar – Data fabric and A2A standards](#)

Technology Trends

Overview: Scaling enterprise AI requires multi-agent orchestration, robust data fabric, reliability/observability, and governance-by-design to meet regulatory thresholds. At the same time, caution is warranted: complexity and "agent washing" risk project cancellations without disciplined execution.

All Findings in Technology Trends

Multi-agent systems, A2A protocols, and persistent memory are becoming foundational.

Explanation: Emerging standards and architectures (e.g., MCP/A2A) aim to enable agents to collaborate effectively across workflows. Unified data fabric and persistent memory are seen as prerequisites for scalable, context-rich agent networks.

- Data fabric enables trustworthy, shared context and reduces silos.
- Agent-to-agent communication frameworks are being explored for interop.
- Persistent memory and scalable infrastructure are essential for enterprise-grade agents.

Rationale: Prioritizing interop, memory, and data fabric in BlueCallom's platform strengthens differentiation as an orchestration-first provider.

Reliability and observability are executive-level gating factors.

Explanation: Many organizations lack mature AI deployment frameworks; scalability and reliability issues raise outage and churn concerns. Enterprises require tracing, health checks, SLOs, and failure recovery to trust mission-critical AI.

- Only ~39% report strong internal frameworks for AI deployment.
- Scalability/reliability issues hinder production scaling and increase overhead.
- Leaders worry about customer loss due to AI outages/instability.

Rationale: Shipping a reliability/observability package and 24/7 AI-Ops (Cross Department Actions) is strategically vital to win enterprise trust.

Governance-by-design and EU AI Act readiness are mandatory.

Explanation: Trust, compliance, and auditability are front-and-center as regulators tighten controls. Fines for high-risk AI violations can be severe, pressuring enterprises to implement guardrails and reporting from day one.

- EU AI Act penalties can reach up to 7% of global revenue for violations.
- Human-in-the-loop, explainability, and audit logs are cited as requirements.
- Enterprises seek transparent, auditable AI actions to satisfy oversight.

Rationale: BlueCallom's COMPLY-AI integration and governance-by-design roadmap (Cross Department Actions) directly mitigate regulatory risk and accelerate adoption in regulated sectors.

Reality check: complexity and "agent washing" heighten project failure risk.

Explanation: Analysts warn that high costs, integration challenges, and mislabeling of conventional tools as "agents" can derail initiatives. Disciplined pilots and value-focused execution are critical.

- Over 40% of agentic AI projects could be canceled by 2027 due to complexity/cost.
- "Agent washing" obscures true capabilities; only a subset of vendors offer genuine autonomy.
- Standardized delivery playbooks reduce failure risk.

Rationale: Emphasizing prescriptive delivery, reference architectures, and outcome benchmarks (Cross Department Actions) increases success rates and credibility.

Sources

1. [TechRadar – Data fabric, governance, memory](#)
2. [ITPro – Reliability gaps and outage risks](#)
3. [Reuters – EU AI Act duties and fines](#)
4. [Reuters – 40% of projects risk cancellation; agent washing](#)

Competition

Overview: SAP is embedding collaborative agents across its cloud suite while acknowledging customer readiness gaps and

advocating regulatory balance. Broad incumbents (Microsoft, Salesforce, Oracle, IBM, ServiceNow) and well-funded startups (Sierra, TinyFish, DualEntry) intensify competition and shape buyer expectations around bundling, orchestration, and deployment flexibility.

All Findings in Competition

SAP is rolling out collaborative AI agents at scale.

Explanation: SAP's leadership positions agentic AI as the next step for enterprise software, embedding multi-agent capabilities (via Joule and domain agents) into core processes across finance, supply chain, and more.

- Announced collaborating agents for sales optimization and supply chain operations.
- Joule is embedded across SAP's portfolio, targeting support for ~80% of common tasks.
- Approach keeps "humans in the driver's seat" while automating routine steps.

Rationale: BlueCallom must differentiate on deployment flexibility (hybrid), OOTB solutions, and orchestration across non-SAP estates to win where SAP leads with footprint.

Readiness gaps and cloud-first stance create openings.

Explanation: SAP notes ~80% of its customers aren't ready for advanced agents; user groups also voice frustration with cloud-only AI. Sovereign and regulated sector moves show SAP targeting trust, but hybrid needs remain acute.

- ~80% of SAP customers not technologically ready for agentic AI today.
- Cloud-only delivery has been poorly received by on-prem customers.
- Partnerships (e.g., Mistral/Capgemini) and sovereign AI offerings address compliance.

Rationale: BlueCallom's "AI without forced cloud migration" pilots (Cross Department Actions) and COMPLY-AI for regulated sectors directly exploit this gap.

Incumbents raise the bar: bundling, build-your-own agents, orchestration.

Explanation: Oracle bundles hundreds of AI features at no extra cost; Microsoft enables custom autonomous agents; Salesforce launches Agentforce; IBM and ServiceNow push orchestration and acquisitions—reshaping expectations for value, governance, and speed.

- Oracle NetSuite: 200+ AI features bundled at no extra cost.
- Microsoft Copilot Studio: build autonomous agents for enterprise tasks.
- Salesforce Agentforce and IBM watsonx Orchestrate elevate agent orchestration; ServiceNow buys Moveworks to accelerate AI service agents.

Rationale: BlueCallom must maintain transparent pricing while selectively bundling baseline capabilities (Cross Department Actions) and emphasize orchestration-first differentiation.

Startup momentum intensifies niche competition.

Explanation: New entrants (Sierra, TinyFish, DualEntry, Lamatic) quickly attract capital by targeting focused workflows with agentic automation, increasing noise and raising the bar on speed-to-value.

- Sierra raises significant funding for supervisor-led service agents.
- TinyFish automates complex web tasks (price/competition tracking).
- DualEntry compresses ERP data migration from months to 24 hours; Lamatic offers low-code agent builders.

Rationale: BlueCallom should leverage partner co-marketing and expand prebuilt solutions (Cross Department Actions) to maintain velocity and credibility.

Sources

1. [Axios – SAP agents for sales and supply chain; readiness gap](#)
2. [SAP News – Joule AI copilot embedded across portfolio](#)
3. [SAP – Multi-agent approach and 80% tasks vision](#)
4. [CIO – Cloud-only frustration among SAP customers](#)
5. [Reuters – SAP/Capgemini/Mistral in regulated sectors](#)
6. [TechRadar – Sovereign OpenAI for Germany with SAP](#)
7. [Axios – Oracle bundles AI at no extra cost](#)
8. [Reuters – Microsoft enables autonomous agent development](#)
9. [Axios – Salesforce Agentforce](#)
10. [Reworked – IBM watsonx Orchestrate](#)
11. [Axios – ServiceNow acquires Moveworks](#)
12. [Axios – Sierra funding and positioning](#)
13. [Reuters – TinyFish funding and approach](#)
14. [Reuters – DualEntry funding and ERP workflows](#)
15. [F6S – Lamatic low-code agents](#)

Business Model and Pricing

Overview: The market is bifurcating between bundled AI (no incremental fee) and premium add-ons metered by opaque models. Customers push for transparent, usage-guardrailed economics. BlueCallom’s published, zero-subscription, usage-based model directly addresses this transparency gap and aligns with CFO-level budgeting needs.

All Findings in Business Model and Pricing

Bundled AI versus premium add-ons: buyers compare value across models.

Explanation: Some incumbents are bundling expansive AI features at no additional cost, while others monetize through premium packages and usage meters. This divergence shapes buyer expectations about baseline versus premium AI value.

- Oracle NetSuite bundles 200+ generative AI features at no extra cost.
- SAP offers premium “Joule” AI packages; pricing is quote-based.
- SAP AI often metered via AI units/credits, scaling with usage.

Rationale: BlueCallom should maintain clarity on what’s included versus metered, and selectively bundle baseline AI to remain competitive while monetizing high-value “flights.”

Demand for transparent pricing and guardrails is explicit.

Explanation: Customers cite unpredictable costs under transaction-based models and call for calculators, budgets, and caps to avoid overruns and speed approvals.

- Users request appropriate, transparent pricing models and clear roadmaps.
- Transaction-based licensing has led to unexpected spend without strict governance.
- Quote-based AI pricing reduces predictability for procurement.

Rationale: BlueCallom’s published pilots, runtime fees, and proposed usage caps/ROI tools (Cross Department Actions) directly

reduce purchasing friction.

BlueCallom's zero-subscription, published pricing differentiates.

Explanation: BlueCallom lists pilot and runtime prices publicly and ties fees to delivered usage. Expected outcome ranges provide CFOs with value framing up front.

- Pilot example: CHF 38,800 (e.g., TRANSFORM); runtime/flight ~CHF 18,750 with volume discounts.
- Other pilots (e.g., DISRUPT CHF 56,500) and runtimes are published for planning.
- Expected annual gains ranges are provided to support business cases.

Rationale: Staying transparent and adding calculators, budget caps, and tiered discounts (Cross Department Actions) reinforce trust and accelerate deals.

Hybrid deployment preferences influence economic choices.

Explanation: Customers seeking hybrid/on-prem AI with feature parity also want pricing predictability across deployments. Vendors that avoid forced cloud migrations and keep parity can reduce switching costs and speed adoption.

- Two-thirds of SAP users prioritize on-site AI independent of vendor cloud.
- On-prem customers expect equal functionality and clarity about costs.
- Cloud-only AI policies create perceived "taxes" on customers' existing estates.

Rationale: BlueCallom's hybrid runner (Cross Department Actions) plus transparent, deployment-agnostic pricing aligns with buyer economics and lowers barriers to entry.

Sources

1. [Axios – Oracle bundles AI at no extra cost](#)
2. [SAP – Business AI premium packages](#)
3. [SAP – Request a quote for AI pricing](#)
4. [SAP Community – AI units/credits metering](#)
5. [CIO – Pricing transparency expectations](#)
6. [TechTarget – Unpredictable costs with transaction-based licensing](#)
7. [BlueCallom – Published pilot pricing](#)
8. [BlueCallom – Published runtime pricing](#)
9. [BlueCallom – Expected gains ranges](#)
10. [CIO – On-prem/hybrid demand with parity](#)

Cross Department Actions

Introduction

This one-page summary consolidates the highest-priority cross-department actions to align execution across Marketing, R&D, Sales, Operations, and Finance. It focuses on initiatives that span functions and require coordinated delivery to accelerate enterprise adoption. Each department includes up to three actions with concise descriptions and supporting sources. Use this as a working reference for planning, resourcing, and status tracking.

Marketing

Position BlueCallom as “Hybrid-Ready, Out-of-the-Box Agentic AI” with Transparent Pricing - Priority 1

- **Description:** Refresh web, decks, and messaging to spotlight hybrid/on-prem deployment, OOTB enterprise workflows (e.g., DISRUPT, TRANSFORM, COMPLY-AI), and zero-subscription, transparent usage pricing—meeting buyer demand for hybrid parity, prebuilt solutions, and pricing clarity while differentiating vs. cloud-only incumbents.
- **Sources:** [CIO](#) | [fair-news.de](#) | [BlueCallom Pricing](#)

Launch “AI Without Forced Cloud Migration” campaigns to SAP user groups - Priority 2

- **Description:** Run targeted campaigns in SAP communities offering hybrid pilots (COMPLY-AI, TRANSFORM) and a migration-safe path—addressing frustration with cloud-only AI and enabling value now with lower risk.
- **Sources:** [CIO](#) | [Axios](#)

Publish ROI/TCO calculators and pricing guardrails - Priority 3

- **Description:** Provide online ROI/TCO tools, sample consumption scenarios, and optional spend caps—reducing anxiety over unpredictable costs and speeding procurement.
 - **Sources:** [TechTarget](#) | [CIO](#) | [SAP AI pricing](#)
-

R&D

Ship a Hybrid/On-Prem Runner for Agentic Workflows - Priority 1

- **Description:** Deliver a VM/K8s runner with secure connectors and full cloud feature parity for private VPCs/on-prem—meeting top enterprise requirements for hybrid flexibility without lock-in.
- **Sources:** [CIO](#)

Build Integration & Data Fabric with Reference Architectures - Priority 2

- **Description:** Provide standardized integration patterns, connectors (SAP, data lakes), and reference architectures with A2A/MCP-style interop and persistent memory—unlocking cross-functional workflows and reducing failed projects.
- **Sources:** [TechRadar](#) | [TrustRadius](#)

Reliability and Observability Package for Agents - Priority 3

- **Description:** Add tracing, SLO dashboards, retries/circuit breakers, failure recovery, cost monitoring, and audit logs—reducing outage risk and addressing AI-ops skills gaps for mission-critical use.
 - **Sources:** [ITPro](#)
-

Sales

“Hybrid Pilot” Play for SAP On-Prem Accounts - Priority 1

- **Description:** Sell packaged pilots (COMPLY-AI or TRANSFORM) delivered hybrid/on-prem with an integration blueprint and fixed pilot fee—addressing cloud-only pain and speeding entry with transparent pricing.
- **Sources:** [CIO](#) | [BlueCallom Pricing](#)

Proposals with Usage Caps and Outcome Framing - Priority 2

- **Description:** Include optional monthly consumption caps and quantified outcome narratives in MSAs—mitigating cost unpredictability and accelerating approvals.
- **Sources:** [TechTarget](#) | [BlueCallom](#)

Regulated Sector Focus with COMPLY-AI - Priority 3

- **Description:** Build vertical plays for finance, public sector, and healthcare centered on COMPLY-AI and EU AI Act readiness—shortening cycles where compliance is a gate.
 - **Sources:** [Reuters](#) | [EU AI Act \(Reuters\)](#) | [PR.com \(COMPLY-AI\)](#)
-

Operations

Standardized Delivery Playbooks to Reduce Project Failure Risk - Priority 1

- **Description:** Implement prescriptive playbooks (discovery → pilot → production), risk checklists, and acceptance criteria across integration, data, governance, and reliability—boosting throughput and success rates.
- **Sources:** [Reuters/Gartner](#)

Establish AI-Ops with 24/7 Monitoring and Runbooks - Priority 2

- **Description:** Stand up centralized observability, SLOs, rollback plans, canaries, incident response, and cost alerts per “flight”—reducing outage risk and customer attrition fears.
- **Sources:** [ITPro](#)

Governance Council, Policy Templates, and COMPLY-AI Ops - Priority 3

- **Description:** Provide standardized governance templates (risk assessments, human oversight, audit logs) and an operating cadence powered by COMPLY-AI outputs to meet documentation and oversight expectations.
 - **Sources:** [TechRadar](#) | [EU AI Act \(Reuters\)](#)
-

Finance

Reinforce Transparent, Zero-Subscription Pricing with Usage Caps - Priority 1

- **Description:** Maintain public pilot/runtime pricing, add optional spend caps and tiered discounts, and publish budgeting guides—delivering cost predictability buyers seek.
- **Sources:** [BlueCallom Pricing](#) | [TechTarget](#) | [CIO](#)

ROI Calculator and Outcome Benchmarks for Pilots - Priority 2

- **Description:** Fund ROI/TCO calculators and publish benchmark ranges for pilot outcomes to strengthen business cases and speed approvals.
- **Sources:** [BlueCallom \(benchmarks\)](#) | [TechRadar](#)

Invest in Compliance and Reliability Roadmap - Priority 3

- **Description:** Allocate budget for governance-by-design, COMPLY-AI enhancements, and observability features—mitigating fines/outage risks and enabling enterprise scale.
- **Sources:** [EU AI Act \(Reuters\)](#) | [ITPro](#)

Marketing Department Findings

2025-10-03 — 21:45

Introduction

This report distills only those research insights that directly support the Marketing actions in Cross Department Actions. The research synthesized online market analyses, customer/user feedback, technology and trend reports, competitor/self reviews, and ecosystem news, focusing on verifiable evidence that informs positioning, campaigns, pricing transparency, thought leadership, and co-marketing.

Findings

Buyers demand hybrid/on-prem AI with cloud parity, out-of-the-box solutions, and transparent pricing

Enterprise user groups clearly state that AI must run on-prem or hybrid with equal functionality to cloud. They also ask for transparent pricing and more prebuilt, immediately deployable solutions. This directly supports positioning BlueCallom as “Hybrid-Ready, Out-of-the-Box Agentic AI” with published, usage-based pricing. The evidence shows BlueCallom already offers ready-to-use agentic solutions and public pricing, matching what buyers are explicitly requesting.

Sources: [CIO \(hybrid/on-prem, parity, pricing transparency\)](#) | [fair-news.de \(BlueCallom ready-to-use solutions\)](#) | [BlueCallom Pricing \(public pilots/usage\)](#)

SAP user groups are frustrated by cloud-only AI; most aren't ready for advanced agents

Customer feedback shows strong resistance to forced cloud moves for AI features. SAP's own messaging acknowledges that ~80% of its customers aren't ready for agentic AI yet. This validates Marketing's “AI Without Forced Cloud Migration” campaigns to SAP communities, offering hybrid pilots (e.g., COMPLY-AI, TRANSFORM) that unlock value now without risky migrations.

Sources: [CIO \(cloud-only frustration\)](#) | [Axios \(80% not ready\)](#)

Unpredictable, transaction-metered costs create anxiety—publish ROI/TCO calculators and budget caps

Users report cost overruns with transaction-based licensing and a lack of pricing clarity for AI features. SAP's AI pricing is commonly quote-based, while some competitors now bundle AI at no extra cost. Providing online ROI/TCO tools, sample consumption scenarios, and optional caps will address procurement concerns, speed approvals, and differentiate BlueCallom's transparent, usage-based model.

Sources: [TechTarget \(unpredictable costs\)](#) | [CIO \(pricing transparency expected\)](#) | [SAP AI pricing \(quote-based\)](#) | [Axios \(Oracle bundles AI at no extra cost\)](#)

AI sprawl, compliance risk, and reliability fears are top of mind—lean into thought leadership with COMPLY-AI

Enterprises struggle with siloed AI tools and governance gaps; EU AI Act fines can reach up to 7% of global revenue; and reliability/observability issues threaten adoption. Marketing can lead with webinars and guides on “taming AI sprawl,” governance-by-design, and reliability patterns—spotlighting COMPLY-AI for EU AI Act readiness and BlueCallom's orchestration capabilities to build trust and credibility.

Sources: [TechRadar \(AI sprawl, orchestration\)](#) | [Reuters \(EU AI Act, high fines\)](#) | [ITPro \(reliability/observability issues\)](#) | [PR.com \(COMPLY-AI launch\)](#)

Co-marketing with solution partners accelerates credibility and time-to-value

Ecosystem validation plus concrete use cases reduce perceived integration risk. BlueCallom's partnership with Log-hub showcases embedded agentic workflows in supply chain scenarios, aligning with the co-branded use-case campaign strategy to drive adoption and generate proof points for Marketing content.

Sources: [Log-hub partnership \(embedded agentic workflows\)](#)

Marketing Department Action Plan

Introduction

This action plan details the prioritized Marketing initiatives for BlueCallom AG, expanded into practical steps that can be executed immediately. It focuses strictly on the actions defined for the Marketing department in the Cross Department Actions document. Each action includes an overview, a step-by-step list, and a concise rationale to guide execution and alignment.

Actions

The Marketing team will reposition BlueCallom's messaging and assets, run targeted hybrid-readiness campaigns to SAP audiences, operationalize ROI/TCO transparency, lead governance and reliability thought leadership, and execute partner co-marketing. Execution emphasizes clarity, speed, measurability, and reuse across website, campaigns, and sales enablement. All actions are designed to produce directly shippable assets, trackable performance, and support for cross-functional go-to-market motions.

Action 1 — Position BlueCallom as “Hybrid-Ready, Out-of-the-Box Agentic AI” with Transparent Pricing

Refresh core messaging and assets to make “Hybrid-Ready, Out-of-the-Box Agentic AI with transparent, usage-based pricing” the primary narrative across the website, decks, and collateral. Emphasize BlueCallom's solutions (e.g., DISRUPT, TRANSFORM, COMPLY-AI), deployment flexibility (cloud, hybrid, on-prem), and zero-subscription, predictable pricing. Ensure parity messaging clearly states that features are consistent across deployment modes and that customers can start fast with prebuilt enterprise workflows.

- Create a messaging framework that defines the primary value proposition, three core pillars (Hybrid-Ready, Out-of-the-Box Solutions, Transparent Pricing), and supporting proof points in one page of copy blocks.
- Rewrite the homepage hero and sub-hero to lead with “Hybrid-Ready, Out-of-the-Box Agentic AI” and add a secondary line on transparent, usage-based pricing with zero subscription.
- Build a Solutions overview page that clearly lists DISRUPT, TRANSFORM, and COMPLY-AI with 3–5 bullet outcomes each and direct CTAs to start a pilot.
- Publish a Deployment page that contrasts cloud, hybrid, and on-prem with feature parity statements, security posture, and typical implementation timelines.
- Update the Pricing page to spotlight public pricing, usage-based model, and pilot options, and add a simple “How pricing works” explainer with 2–3 illustrated scenarios.
- Revise the corporate and product pitch decks to align with the new narrative; include one slide each for Hybrid-Ready, OOTB workflows, and Transparent Pricing.
- Create a one-page sales leave-behind that summarizes the three pillars, available solutions, deployment options, and the “start in weeks” message.
- Run an A/B test on homepage headlines and CTAs for two weeks to validate conversions on the new positioning versus the current baseline.
- Enable Sales and Customer Success with a two-page positioning brief and a 15-minute enablement recording to ensure consistent delivery.

Rationale: Enterprise buyers prioritize hybrid/on-prem options without compromising features, want fast time-to-value via prebuilt workflows, and demand clear, predictable pricing. A unified narrative that explicitly addresses these priorities increases relevance, trust, and conversion across web, campaign, and sales motions.

Action 2 — Launch “AI Without Forced Cloud Migration” campaigns to SAP user groups

Run targeted campaigns in SAP communities and industry associations that promote hybrid pilots and a migration-safe path to agentic AI. Lead with COMPLY-AI and TRANSFORM pilots that can be deployed without forced cloud moves and with low integration risk. Use precise targeting, credible content, and strong CTAs that invite prospects to start pilots now while preserving their current SAP landscape.

- Draft a campaign brief that defines audience segments (SAP on-prem stakeholders, IT leaders, compliance owners), the core promise (“AI now, no forced migration”), and the primary CTAs (Book a pilot, Request deployment blueprint).
- Develop a landing page that explains the hybrid pilot approach, lists eligible pilots (COMPLY-AI, TRANSFORM), and includes a short intake form for scoping calls.
- Create two email sequences of five touches each: one for technical leaders (deployment, security, integration) and one for business leaders (outcomes, risk, time-to-value).
- Produce three ad creatives and three LinkedIn post variants emphasizing “No forced cloud migration,” hybrid readiness, and pilot fast starts.
- Prepare a 30-minute webinar titled “Agentic AI for SAP Without Cloud Migration” including a live demonstration and a downloadable pilot checklist.
- Offer a two-page “Hybrid Pilot Starter Pack” PDF that outlines prerequisites, timeline, data requirements, and success criteria.
- Coordinate community placements and sponsorships with SAP user groups and relevant industry associations to distribute the webinar and starter pack.
- Set measurable goals for form fills, webinar registrations, pilot requests, and opportunity creation, and track weekly performance in a shared dashboard.
- Establish a lead handoff SLA with Sales for sub-24-hour follow-up and a feedback loop to refine targeting and messaging after the first two weeks.

Rationale: Many SAP customers resist cloud-only mandates and want a safe, incremental path to value. Highlighting hybrid pilots that avoid forced migrations reduces perceived risk and accelerates adoption.

Action 3 — Publish ROI/TCO calculators and pricing guardrails

Provide interactive calculators and clear budgeting aids that explain expected value, usage ranges, and optional spend caps. Include simple consumption scenarios with sliders and presets to model pilot costs and expansion plans. Present guardrails plainly so procurement teams understand how to control spend without complex negotiations.

- Define calculator inputs and outputs, including workload assumptions, expected productivity outcomes, usage “flights,” and optional monthly caps.
- Design a web-based calculator experience with presets for common scenarios (e.g., compliance reporting, productivity optimization, innovation pipeline) and downloadable results.
- Create a Pricing Guardrails section on the Pricing page that explains caps, volume discounts, and how spend is controlled during a flight.
- Publish three example budgeting scenarios as short case vignettes with data ranges and a simple “what-if” view.
- Embed the calculator on the Pricing page and in relevant solution pages, and add UTM tracking to measure engagement and funnel progression.
- Produce a one-page PDF procurement guide that summarizes pricing logic, caps, and approval-friendly language for MSAs.
- Train Sales on how to use the calculator live during calls and how to export a customer-ready summary in under five minutes.
- Set monthly review cadences with Finance to validate assumptions and update presets as real usage and benchmarks evolve.

Rationale: Buyers often struggle with unpredictable AI costs and need transparent, controllable pricing. Calculators and guardrails reduce anxiety, speed approvals, and reinforce BlueCallom's zero-subscription, usage-based clarity.

Action 4 — Thought leadership on Governance, Reliability, and “Taming AI Sprawl”

Lead the market conversation on enterprise-grade governance, orchestration, and reliability by producing webinars, guides, and playbooks. Showcase COMPLY-AI for EU AI Act readiness, multi-agent orchestration patterns, and reliability/observability practices as practical templates. Align content with concrete workflows and checklists that prospects can apply immediately.

- Publish a quarterly editorial calendar covering governance-by-design, EU AI Act readiness, agent orchestration patterns, and reliability/observability foundations.
- Create a flagship “AI Governance Starter Guide” that includes a readiness checklist, oversight roles, audit requirements, and sample reports produced by COMPLY-AI.
- Host a three-part webinar series: “EU AI Act in Practice,” “From Sprawl to Orchestration,” and “Designing for Reliability,” each with live Q&A and downloadable templates.
- Develop two detailed solution briefs illustrating cross-functional agent workflows and how to prevent AI sprawl with an enterprise data fabric and guardrails.
- Produce a five-page Reliability Playbook that explains tracing, health checks, SLOs, and failure recovery patterns in clear, non-technical language for business stakeholders.
- Capture each webinar and convert into a blog summary, a short video highlight, and a LinkedIn carousel to extend reach and engagement.
- Enable Sales with a three-slide “Why Governance and Reliability Matter” mini-pack tied to COMPLY-AI and orchestration capabilities with suggested talk tracks.
- Collect attendee questions and objections, and update an FAQ document that marketing, sales, and product teams can use for consistent responses.

Rationale: Enterprises need confidence that AI can be governed, reliable, and orchestrated across silos. Providing practical guidance and assets positions BlueCallom as a trusted, enterprise-ready partner and drives qualified demand for COMPLY-AI and orchestration-led solutions.

Action 5 — Partner co-marketing with solution ecosystems

Run co-branded campaigns with strategic partners to showcase embedded agentic workflows and rapid time-to-value. Focus on clear, repeatable use cases and shared proof points that reduce integration risk. Deliver assets jointly and coordinate distribution to maximize reach and credibility.

- Select two to three partners with complementary solutions and clear enterprise use cases, with at least one supply chain example.
- Co-create a campaign brief that defines the joint value proposition, use-case outcomes, proof points, and combined CTAs.
- Develop a co-branded one-pager, a short demo video, and a solution brief that explains the end-to-end workflow and business impact.
- Host a joint webinar featuring a live demo of the embedded agentic workflow and a customer-style walkthrough of the deployment steps.
- Publish a case-style article or success vignette highlighting time-to-value, integration steps, and performance metrics.
- Coordinate LinkedIn posts, email campaigns, and partner marketplace listings to drive consistent messaging and traffic to a shared landing page.
- Define shared KPIs, including registrations, qualified leads, pilot requests, and influenced pipeline, and review results in a joint monthly meeting.

- Create a partner enablement kit with messaging, visuals, and a quick-start guide so partners can independently promote the campaign.

Rationale: Ecosystem credibility and real-world use cases lower perceived integration risk and speed adoption. Co-marketing demonstrates practical value, expands reach, and supports faster deal cycles with concrete, verifiable outcomes.

Research and Development Department Action Plan

Introduction

This plan translates the Research and Development actions into clear execution steps for BlueCallom AG's R&D team. It covers delivery of a hybrid/on-prem runner, integration and data fabric with reference architectures, reliability and observability for agents, prebuilt solutions and UX modernization, governance-by-design and EU AI Act readiness, and multi-agent orchestration. The scope is limited to the actions enumerated in the Cross Department Actions and focuses on concrete delivery work for the R&D function.

Actions

The R&D team will ship deployable capabilities and content that enable enterprise-grade agentic AI, ready for hybrid and on-prem environments at feature parity with cloud. The work includes platform components, integration patterns, operational tooling, solution templates, and in-product governance. Each action below lists the execution approach, detailed steps, and a concise rationale to guide prioritization and delivery.

Action 1 — Ship a Hybrid/On-Prem Runner for Agentic Workflows (Priority 1)

Deliver a secure, deployable runner that executes BlueCallom agents inside private VPCs and on-prem data centers with full feature parity to our cloud runtime. Package for both VM and Kubernetes targets, include hardened images, and provide zero-egress options. Ensure a smooth install, upgrades, observability hooks, and enterprise identity integration.

- Define a parity matrix versus cloud: agent execution, secrets management, connectors, policy enforcement, audit logging, usage metering, model routing, and update channels; mark each as GA, beta, or backlog with target dates.
- Produce artifacts for K8s (Helm charts, Operators) and VM (OVA/AMI images) with CIS/STIG-aligned hardening, rootless containers, and image signing (Sigstore/Cosign).
- Implement identity and access: native SAML/OIDC SSO, SCIM provisioning, role-based access control, and service accounts; include just-in-time user provisioning and per-tenant isolation.
- Offer network profiles: standard (limited egress), private registry pull, and fully air-gapped; include offline license activation and deferred metering sync.
- Bundle secure connectors (SAP, major data lakes, message buses) with policy-controlled data paths, secrets vault integration, and customer-managed keys (HSM/KMS) support.
- Add preflight validation and post-install health checks; output a readiness report with cluster resources, endpoints, certificates, and SLO baselines.
- Integrate observability: OpenTelemetry traces/metrics/logs, syslog/SIEM export, and run-level audit trails; provide Grafana/Prometheus dashboards and alert rules.
- Provide lifecycle ops: in-place upgrades, canary channels, rollbacks, and configuration drift detection; document backup/restore of stateful components.
- Validate on common targets: EKS, AKS, GKE, OpenShift, Rancher, and vSphere; certify minimum versions and publish a support matrix.
- Define acceptance criteria: install under 60 minutes, zero critical CVEs at release, 99.9% runtime availability in soak tests, and 100% parity for GA-listed features.

Rationale: A production-ready hybrid/on-prem runner removes deployment friction for enterprises with strict data and networking controls while maintaining the same capabilities they expect from our cloud runtime.

Action 2 — Build Integration & Data Fabric with Reference Architectures (Priority 2)

Create a standardized integration and data fabric that unifies agent access to enterprise systems and data with clear patterns and templates. Provide connectors, canonical entity models, memory/persistence options, and A2A/MCP-style interoperability. Publish reference architectures with runnable examples to reduce time-to-integration.

- Define canonical entities (customer, supplier, product, order, ticket) and schemas; supply mapping guides for SAP modules and common data lake formats.
- Stand up a persistent memory layer: vector search, key-value store, and relational store; expose CRUD and search APIs with retention and encryption policies.
- Ship high-priority connectors: SAP (RFC/OData/Event Mesh), major data lakes (S3/ADLS/GCS), messaging (Kafka/Event Hub), and identity (Azure AD/Okta); include throttling and backoff controls.
- Introduce A2A/MCP-style interop contracts for agent-to-agent and tool invocation; version the contracts and publish a compatibility matrix.
- Provide ETL/ELT templates for initial sync and incremental CDC; include data quality checks, lineage metadata, and replayable pipelines.
- Publish three reference architectures (on-prem isolated, hybrid with private link, cloud-first) with diagrams, Terraform/IaC samples, and cost envelopes.
- Release a developer toolkit: SDKs, code samples, Postman collections, and end-to-end examples (e.g., supply chain anomaly → SAP update → notification).
- Add observability for integration flows: per-connector health, throughput, errors, and retries; provide dashboards and common alerts.
- Document security posture: least-privilege service principals, credential rotation, network ACLs, and data classification handling in pipelines.
- Set definition of done: connector passes contract tests, latency under target for top flows, and sample references deploy in under two hours.

Rationale: Standardized patterns, connectors, and references shorten integration cycles and give customers a predictable path to connect agents across silos.

Action 3 — Reliability and Observability Package for Agents (Priority 3)

Embed operational reliability and visibility into every agent run so teams can trace, monitor, and control runtime behavior. Provide tracing, health checks, retries, circuit breakers, SLO dashboards, and cost monitoring, including exportable audit logs. Deliver recovery workflows to resume or replay runs safely.

- Instrument agents with OpenTelemetry traces and metrics: step timings, token usage, model calls, cache hits, and external API latencies; propagate correlation IDs across services.
- Implement health checks, retry strategies with jitter, circuit breakers for flaky dependencies, and bulkheading to isolate failing components.
- Add deterministic replay and checkpointing to resume long-running workflows; persist intermediate state securely with idempotent write semantics.
- Provide SLO dashboards (availability, latency, success rate, cost per outcome) and alert rules for breach, saturation, and error spikes.
- Enable cost governance: per-agent budget caps, alert thresholds, and per-run cost attribution; export usage to billing systems.
- Create dead-letter queues and automated remediation playbooks (e.g., auto-retry with backoff, route to human review, or downgrade model tier).

- Produce comprehensive audit logs: prompts, tool calls, data sources touched, decisions taken, and human interventions; support tamper-evident storage.
- Introduce safe deploys: canaries, blue-green switching, progressive rollouts, and automatic rollback on health regression.
- Run chaos tests (dependency faults, latency injection, quota exhaustion) in staging and document recovery time objectives.
- Publish runbooks with standard error codes, troubleshooting steps, and escalation paths; integrate with ticketing systems.

Rationale: Built-in reliability and observability reduce outages, speed incident response, and give enterprises confidence to operate agents in mission-critical workflows.

Action 4 — Expand Prebuilt Solutions, Templates, and UX Modernization (Priority 4)

Grow a curated library of ready-to-use agents, activities, and starter kits while modernizing the UX to be simpler and more guided. Deliver function-specific templates with data models, playbooks, and sample integrations. Streamline onboarding and in-app guidance to minimize setup time and errors.

- Select high-impact starter kits by function (e.g., compliance, productivity transformation, supply chain) with clear inputs, outputs, and KPIs.
- Package each starter kit with preconfigured agents, activities, data schemas, example datasets, and step-by-step configuration wizards.
- Introduce a template taxonomy and quality bar: coding conventions, security review, localization readiness, tests, and performance baselines.
- Build an in-product template catalog with search, tags, versioning, release notes, and deprecation guidance.
- Modernize onboarding flows: guided setup, contextual tooltips, inline validation, human-friendly error messages, and progress save/resume.
- Add accessibility and responsiveness checks; ensure keyboard navigation, ARIA labels, and contrast compliance.
- Instrument UX analytics: time-to-first-value, configuration completion rate, and template adoption; use findings for iterative improvements.
- Provide “getting started” tutorials and short videos embedded at key steps; offer one-click import/export for enterprise distribution.
- Localize core templates and UI strings for priority languages and deliver a translation update pipeline.
- Definition of done: a user with standard permissions can deploy a starter kit to a non-prod environment and realize a demo outcome in under 60 minutes.

Rationale: More out-of-the-box content plus a guided UX reduces setup friction and accelerates value realization for enterprise teams.

Action 5 — Governance-by-Design and EU AI Act Readiness (Priority 5)

Embed governance controls, human oversight, explainability, and auditable processes into the platform and integrate them with COMPLY-AI. Provide out-of-the-box compliance reports and incident workflows. Ensure policy enforcement is consistent across cloud and hybrid/on-prem deployments.

- Define governance policies as code: guardrails, model and tool allowlists, data residency rules, and approval workflows; enforce at runtime via a central policy engine.
- Add human-in-the-loop checkpoints with configurable criteria; capture reviewer identity, comments, and decisions in immutable logs.
- Provide explainability artifacts: prompt lineage, decision rationales, model cards, and data source attributions; enable export for audits.

- Integrate COMPLY-AI to generate standard compliance reports automatically (policy conformance, risk register, DPIA inputs, incident summaries).
- Implement incident management: triage forms, severity classification, containment steps, and corrective action plans with due-date tracking.
- Track dataset and model provenance: versions, training/evaluation datasets, change logs, and approval history.
- Ship PII detection and redaction tools, consent capture hooks, and retention/erasure policies with per-jurisdiction settings.
- Offer role segregation (builder, operator, approver, auditor) and enforce separation of duties; integrate with enterprise identity providers.
- Validate compliance with automated checks in CI/CD and pre-prod gates; block deploys that violate critical controls.
- Provide auditor-ready export bundles (PDF/JSON) containing policies, runs, approvals, incidents, and change history.

Rationale: Governance baked into design and operations makes compliance practical at scale and standardizes controls across all deployments.

Action 6 — Multi-Agent Orchestration for “Insight-to-Action” (Priority 6)

Advance multi-agent collaboration so insights lead to safe, governed execution steps in enterprise systems. Enhance memory, planning, and role coordination while enforcing guardrails and approvals. Deliver embedded workflow actions inside tools customers already use.

- Implement a shared task graph with a planner/dispatcher, agent roles (analyst, validator, executor), and arbitration for conflicts; persist state for hand-offs and replays.
- Strengthen memory: short-term task context, long-term knowledge, and retrieval strategies; tune for precision and safety in action decisions.
- Add policy-aware action execution: allowlists, rate limits, dry-run previews, and human approval gates for sensitive steps.
- Integrate with enterprise systems for actioning: SAP (orders, invoices, master data), ServiceNow/Jira (tickets), email/chat (notifications), and document systems.
- Provide an orchestration builder with visual flow editing, reusable subflows, test harnesses, and simulation mode for safe validation.
- Support embedded actions in existing UIs through extensions and side panels; log user overrides and capture feedback for continuous improvement.
- Deliver rollback and compensation actions for partial failures; ensure idempotent updates and reconcile states across systems.
- Expose run and outcome metrics: insight quality, approval rates, execution success, and cycle time; add alerts for stuck or looping flows.
- Package reference orchestrations per function (e.g., compliance case triage, productivity improvement rollouts) with sample data and tests.
- Definition of done: a reference orchestration executes end-to-end with at least one human approval, updates a target system, and logs a complete audit trail.

Rationale: Multi-agent orchestration that moves from analysis to governed action closes the loop and embeds AI directly into enterprise workflows.

Sales Department Findings

2025-10-03 — 12:00–21:45 UTC

Introduction

This summary highlights only those research findings that directly support the Sales actions listed in Cross Department Actions. The underlying research synthesizes industry news, analyst coverage, customer forums/reviews, and official product/pricing pages across our competitive space (BlueCallom and SAP, plus broader incumbents), conducted on the same day across multiple reports.

Findings

Hybrid/on-prem AI is a top SAP customer demand—prime for a “Hybrid Pilot” entry.

Multiple SAP user groups voice frustration with cloud-only AI, explicitly requesting on-prem/hybrid options with feature parity. At the same time, SAP’s CEO states ~80% of customers aren’t ready for advanced agents, signaling an opening for low-risk, hybrid pilots. This directly supports the Sales “Hybrid Pilot” play: lead with a packaged COMPLY-AI or TRANSFORM pilot on-prem/VPC, plus an integration blueprint and fixed pilot fee.

Sources: [CIO \(cloud-only frustration\)](#) | [CIO \(hybrid/on-prem priority & parity\)](#) | [Axios \(~80% not ready\)](#) | [BlueCallom \(transparent pilot pricing\)](#)

Cost unpredictability is a real blocker—counter with usage caps and outcome framing.

Customers report anxiety over transaction-based AI/RPA costs and opaque AI add-on pricing, especially in SAP environments. Buyer communities call for transparent models and clear ROI guidance. This validates offering proposals with optional monthly consumption caps and quantified outcome narratives, differentiating against quote-based or metered competitors and accelerating approvals.

Sources: [TechTarget \(unpredictable licensing costs\)](#) | [CIO \(pricing transparency expected\)](#) | [SAP AI pricing \(quote-based\)](#) | [BlueCallom \(expected gains, zero subscription\)](#)

Regulated sectors face urgent AI compliance pressure—lead with COMPLY-AI vertical plays.

Regulatory risk is intensifying (e.g., EU AI Act fines up to 7% of global revenue), and large vendors partner specifically for sensitive industries. This creates a strong opening to sell COMPLY-AI-centered solutions for finance, public sector, and healthcare, bundling EU AI Act readiness assessments to shorten sales cycles where governance is the gating factor.

Sources: [Reuters \(EU AI Act fines\)](#) | [Reuters \(AI in sensitive sectors\)](#) | [PR.com \(BlueCallom COMPLY-AI\)](#)

Enterprises lack AI readiness and integration patterns—sell Workshops and Integration Clinics.

Evidence shows many enterprises lack internal frameworks, skills, and clear integration guidance, with users explicitly asking vendors for standardized rules and reference architectures. Offering structured readiness assessments and hands-on integration clinics de-risks projects and speeds time-to-value, aligning directly to Sales’ pre-sales enablement packages.

Sources: [ITPro \(skills gaps, lack of frameworks\)](#) | [CIO \(need for standardized integration guidance\)](#) | [Axios \(~80% not ready\)](#)

Incumbents push agents but have exploitable gaps—equip Battlecards: Hybrid + OOTB + Transparent.

Microsoft, Salesforce, IBM, and SAP are rolling out autonomous agents and orchestration—but often tied to cloud ecosystems and quote-based pricing, while Oracle is bundling some AI at no extra cost. Clear, side-by-side battlecards emphasizing hybrid deployment, out-of-the-box workflows, and transparent pricing give reps crisp differentiation to win against suite incumbents.

Sources: [Reuters \(Microsoft autonomous agents\)](#) | [Axios \(Salesforce Agentforce\)](#) | [Axios \(Oracle bundles AI\)](#) | [Reworked \(IBM orchestration\)](#) | [Axios \(SAP agent plans\)](#) | [SAP AI pricing \(quote-based\)](#)

OOTB content wins speed deals—package pilots with ready-to-use workflows.

Customers complain about limited prepackaged activities and templates in incumbent stacks, which slows time-to-value. BlueCallom’s

ready-to-use, immediately deployable solutions are a tangible advantage. For Sales, bundling OOTB workflows in proposals strengthens the “fast start” narrative for pilots and supports the Battlecards story.

Sources: [TechTarget \(limited prepackaged activities\)](#) | [fair-news.de \(BlueCallom ready-to-use solutions\)](#)

Sales Department Action Plan

Introduction

This action plan defines the concrete steps the Sales department at BlueCallom AG will execute to operationalize the cross-department Sales actions. It translates each action into clear tasks, deliverables, and ownership cues that are immediately usable in the field. The scope covers execution of the defined Sales actions only.

Actions

The Sales team will drive a focused motion centered on hybrid pilots for SAP on-prem accounts, proposals that include usage caps and outcome framing, vertical plays for regulated sectors, readiness workshops and integration clinics, and competitive battlecards. Each action is expanded into practical steps with repeatable templates, qualification criteria, and CRM instrumentation. Execution emphasizes speed to pilot, transparent pricing guardrails, compliance-first positioning, and clear competitive differentiation.

Action 1 — “Hybrid Pilot” Play for SAP On-Prem Accounts

Package and sell a fixed-scope, fixed-fee hybrid/on-prem pilot using COMPLY-AI or TRANSFORM, including an integration blueprint and a clear path to production. Focus on SAP on-prem heavy accounts where cloud-only approaches face adoption resistance, and lead with a prescriptive discovery-to-pilot workflow. Standardize the sales kit, SoW, and qualification checklist to reduce cycle time and risk.

- Build a named-account list of 50–100 SAP on-prem targets in CRM, tagging industry, SAP estate (ECC/S4 on-prem), data residency constraints, and key personas.
- Define two pilot offers (COMPLY-AI and TRANSFORM) with 6–8 week timelines, explicit deliverables (integration blueprint, runbook, success criteria), and a fixed pilot fee aligned to published pricing.
- Create a sales kit that includes a discovery questionnaire, solution overview, security/IT checklist, hybrid deployment FAQ, and a draft SoW template.
- Run qualification calls to confirm on-prem requirement, integration surface area, data locations, sponsor availability, and readiness for a time-boxed pilot.
- Schedule a technical discovery with Solution Engineering to validate connectors, access methods, and non-functional needs, capturing risks and mitigation steps.
- Propose the pilot with a milestone plan (kickoff, integration validation, use-case execution, acceptance), named responsibilities, and exit criteria tied to production steps.
- Attach a procurement pack with public pricing, data processing terms, deployment architecture overview, and an MSA addendum referencing hybrid operations.
- Set pilot success metrics (e.g., time saved, error reduction, compliance evidence produced) and agree on a production rollout plan triggered by acceptance.
- Track the pipeline stages in CRM (Identified → Qualified → Design Approved → SoW Signed → Deployed → Completed → Expanded) with stage exit criteria and dates.
- Establish a standard post-pilot expansion offer that converts the pilot into production flights with defined service levels and change management support.

Rationale: A packaged hybrid pilot lowers risk for SAP on-prem accounts, addresses integration ambiguity with a prescriptive blueprint, and uses transparent fixed pilot fees to accelerate buying decisions.

Action 2 — Proposals with Usage Caps and Outcome Framing

Embed optional monthly consumption caps and outcome-based narratives into all proposals and MSAs to remove cost unpredictability concerns. Present simple pricing choices and clear alerting/auto-pause behaviors that protect budgets while enabling scale. Tie economic value to quantified productivity or compliance outcomes that align with enterprise approvals.

- Add a proposal pricing page that offers three options: capped monthly usage with alerts, uncapped with volume discount, and capped with predefined overage approval workflow.
- Include an outcomes section that quantifies expected impact in business terms (e.g., hours saved, manual steps eliminated, audit artifacts produced) for the selected workflows.
- Insert MSA language defining the monthly cap, alert thresholds at 70%/90%/100%, auto-pause behavior, and a simple change-order mechanism to adjust caps mid-term.
- Provide two to three sample consumption scenarios that map use-case volume to estimated monthly spend and show how caps enforce budget guardrails.
- Reference the standard acceptance metrics from the pilot (e.g., productivity uplift or compliance deliverables) and connect them to the proposed production flight tiers.
- Coordinate with Finance on discount thresholds, cap floors/ceilings, and standardized quote templates to ensure consistency and speed during approvals.
- Attach a one-page “cost predictability” explainer that details alerting, reporting cadence, and the cap adjustment process for budget holders.
- Train account teams to narrate outcomes first, then cost, emphasizing predictable guardrails and clear levers for scaling value.
- Record “cap requested” and “cap level” fields in CRM opportunities to forecast revenue and monitor adoption readiness.
- After signature, confirm cap configuration and alert recipients in the onboarding checklist to ensure guardrails are live from day one.

Rationale: Usage caps and outcome framing reduce procurement anxiety about unpredictable consumption while keeping focus on measurable business value, which shortens approval cycles and differentiates from transaction-metered pricing.

Action 3 — Regulated Sector Focus with COMPLY-AI

Run focused vertical plays for financial services, public sector, and healthcare using COMPLY-AI as the anchor solution. Lead with an EU AI Act readiness assessment, a compliance-first discovery, and evidence packs that demonstrate auditability and governance. Align sales artifacts to each sector’s terminology, risks, and decision processes.

- Segment the target list by regulated sectors and map the key regulatory drivers, internal compliance stakeholders, and decision checkpoints for each account.
- Build sector-specific narratives and reference architectures that show COMPLY-AI workflows, human-in-the-loop steps, and audit reporting outputs.
- Use a compliance discovery checklist covering risk classification, data residency, model oversight, incident reporting, and evidence retention requirements.
- Offer an EU AI Act readiness assessment as the first engagement milestone, with a short SoW and a clear artifact bundle as deliverables.
- Assemble evidence packs that include sample compliance reports, policy templates, audit log exports, and oversight workflows to de-risk evaluations.
- Align legal and security responses in a standardized RFI/RFP library tailored to finance, public sector, and healthcare requirements.
- Define a COMPLY-AI pilot package for each sector with domain-specific controls, acceptance criteria, and a production transition plan.
- Map expansion paths from COMPLY-AI into adjacent workflows (e.g., procurement oversight, vendor risk, or operational audits) with named cross-sell plays.

- Track sector attribution in CRM and report pipeline conversion rates and deal cycle time by vertical to refine the playbooks.
- Secure sector references or anonymized outcome summaries to use in future proposals and executive briefings.

Rationale: Compliance is often a gating factor in regulated sectors, and leading with COMPLY-AI plus a readiness assessment meets buyer expectations, speeds trust building, and creates a clear path to broader adoption.

Action 4 — AI Readiness Workshops and Integration Clinics

Productize pre-sales enablement into structured workshops that assess readiness, define reference architectures, and run hands-on integration sessions. Use standardized agendas, checklists, and outputs to compress time from discovery to pilot approval. Convert workshop outcomes directly into pilot SoWs with named owners, dates, and risk mitigations.

- Offer a two-part format: a half-day Readiness Assessment for stakeholders and a full-day Integration Clinic with technical teams.
- Publish agendas, prerequisites, and expected outputs in a one-pager so customers understand time commitments and deliverables upfront.
- Use guided assets: environment inventory sheets, connector matrices, data flow maps, security and compliance checklists, and success criteria templates.
- Staff each session with a Sales Lead and a Solution Engineer, and confirm access to necessary sandbox or staging resources before the clinic.
- Score each account on readiness dimensions (data access, integration posture, governance, sponsorship) and document remediation actions with owners.
- Draft a pilot SoW on the spot using standardized language, pre-filled milestones, and a list of required customer inputs and timelines.
- Capture technical decisions and risks in a living integration blueprint that becomes part of the pilot deliverables.
- Send a follow-up pack within two business days, including the blueprint, SoW draft, timeline, and a checklist to reach pilot kickoff readiness.
- Record workshop outcomes, readiness scores, and next steps in CRM, and track workshop-to-pilot conversion as a key KPI.
- Maintain a quarterly backlog of integration learnings to refine connector priorities and update the clinic materials.

Rationale: Structured workshops reduce uncertainty, align stakeholders, and transform pre-sales discovery into actionable pilot commitments with clear technical and governance readiness.

Action 5 — Competitive Battlecards: Hybrid + OOTB + Transparent versus Incumbents

Equip the field with side-by-side battlecards against SAP, Microsoft, Salesforce, Oracle, IBM, and ServiceNow that highlight hybrid deployment, out-of-the-box solutions, transparent pricing, and orchestration strengths. Provide objection handling, discovery questions, demo flows, and proof points tailored to each competitor. Keep cards current with win/loss insights and circulate updates through regular enablement.

- Create a standard battlecard format: positioning summary, claim/counter, discovery questions, proof points, landmines to avoid, and recommended next steps.
- Include a comparison grid for deployment flexibility (hybrid/on-prem), feature parity across deployments, OOTB solution coverage, and pricing transparency.
- Develop objection-handling scripts for common concerns such as on-prem parity, integration complexity, governance readiness, and reliability operations.
- Attach short demo flows that showcase hybrid runners, COMPLY-AI audit outputs, and multi-agent orchestration embedded in existing workflows.
- Add customer evidence snippets from pilots and productions, focusing on measurable outcomes and decision factors that tipped the deal.

- Publish the battlecards as CRM quick links and track usage with enablement analytics to correlate with win rates and cycle times.
- Run monthly role-play sessions to practice discovery and objection handling for each competitor scenario.
- Set a monthly refresh cadence with inputs from field feedback, product updates, and competitive intelligence to keep content current.
- Provide an executive summary version for senior sponsors that clarifies strategic differences in one page.
- Archive retired claims and counters to maintain a defensible history and prevent outdated messaging from circulating.

Rationale: Clear, current battlecards help sellers position hybrid flexibility, OOTB time-to-value, and transparent pricing against bundled incumbent offerings, improving win rates and deal velocity.

Operations Department Findings

2025-10-03 — 12:00–21:45 UTC

Introduction

This report distills only Operations-relevant insights from the research corpus and ties them directly to the Operations actions in Cross Department Actions. We analyzed industry news, analyst notes, product reviews, and governance updates to surface what impacts delivery success, reliability, compliance operations, enablement, and integration—so your team can execute with confidence.

Findings

Standardize delivery playbooks to cut high project-failure risk

Analysts warn that more than 40% of agentic AI initiatives may be canceled by 2027 due to complexity and integration challenges. Prescriptive delivery playbooks (discovery → pilot → production), with risk checklists and clear acceptance criteria across integration, data, governance, and reliability, reduce ambiguity and prevent stalls.

Sources: [Reuters/Gartner: 40% of agentic AI projects scrapped](#)

Build AI-Ops: 24/7 observability, SLOs, and cost guardrails

Enterprises report reliability and scalability gaps, with executives fearing customer loss from AI outages. Instituting centralized monitoring, SLO dashboards, canaries/rollbacks, incident runbooks, and per-flight cost alerts directly mitigates instability and improves customer trust in production agents.

Sources: [ITPro: reliability/scalability issues and outage risk](#)

Operationalize governance: Council, policy templates, and COMPLY-AI ops

The EU AI Act introduces strict duties and fines up to 7% of global revenue. Research highlights the need for built-in trust measures—auditability, human-in-the-loop, and persistent memory—with clear oversight cadences. Establish a governance council, standard policy templates, and run COMPLY-AI outputs as an operational rhythm to stay audit-ready.

Sources: [Reuters: EU AI Act duties and fines](#) | [TechRadar: trust, governance, auditability requirements](#) | [PR.com: BlueCallom COMPLY-AI launch](#)

Close the skills gap with structured enablement and certification

Many organizations lack AI deployment frameworks and skills, causing misalignment and delays. Structured training paths (admin, solution designer, ops), certifications, and department roll-out kits—leveraging GPTBlue—accelerate readiness. Early user feedback also points to easy solution distribution across departments, aiding scale.

Sources: [ITPro: skills gap and misalignment](#) | [BlueCallom users: easy distribution across departments](#)

Stand up an integration factory to remove onboarding friction

User reviews cite integration/connectivity pain (external data lakes, error-prone installs) and weak debug logs. Customers also request standardized integration rules and reference architectures. Prioritize connectors to SAP estates and data lakes, plus versioned integration guides, to cut time-to-value and reduce support load.

Sources: [TrustRadius: integration/connectivity challenges](#) | [CIO: desire for standardized integration rules/guidance](#)

Operations Department Action Plan

Introduction

This plan outlines the concrete actions the Operations department will execute to standardize delivery, implement AI-Ops, formalize governance, enable customers, and scale integrations. It expands each action into practical steps, templates, and procedures ready for immediate adoption. The scope is limited to execution guidance for the actions listed in the cross-department directives.

Actions

The Operations function will operationalize standardized delivery playbooks, establish 24/7 AI-Ops with strong observability and incident response, and embed governance practices and COMPLY-AI outputs into daily operations. It will stand up structured training and certifications via GPTBlue and run an integration factory to produce and maintain connectors and guides. The emphasis is on repeatability, risk reduction, and scalable execution across discovery, pilot, and production phases.

Action 1 — Standardized Delivery Playbooks to Reduce Project Failure Risk

Build and roll out prescriptive playbooks that guide teams from discovery through pilot to production with clear artifacts, roles, milestones, and acceptance criteria. The playbooks must explicitly address integration, data readiness, governance, and reliability to prevent common failure modes. Deploy them in a single, version-controlled repository with training for all delivery leaders.

- Form a cross-functional working group (Operations lead, Solution Architecture, R&D, Security/Compliance, Customer Success) and appoint a single project owner to drive playbook creation and adoption.
- Define an end-to-end lifecycle with stages: Discovery, Solution Design, Pilot, Hardening, Production Cutover, and Hypercare, each with entry/exit criteria and required artifacts.
- Create standard templates: Project Charter, Stakeholder Map, RAID Log, Risk Checklists (integration, data, governance, reliability, cost), Test Plan, Cutover Plan, and Hypercare Plan.
- Establish a Definition of Ready (DoR) and Definition of Done (DoD) per stage, including data availability, connector readiness, security reviews, SLOs, and acceptance criteria.
- Publish a standard Integration Blueprint outline that references approved reference architectures and connector patterns for SAP and common enterprise systems.
- Introduce stage gates with go/no-go scorecards; require sign-offs from Operations, Security/Compliance, and Customer Sponsor before progressing.
- Mandate traceability from requirements to test cases to acceptance criteria; store evidence and approvals in a shared repository.
- Provide a timeline guideline per stage (e.g., Discovery 2–3 weeks, Pilot 4–6 weeks, Hardening 2–4 weeks) and align staffing plans accordingly.
- Run mandatory enablement sessions for Project Managers and Solution Architects; provide quick-reference checklists and a 1-page “playbook at a glance.”
- Pilot the playbook with two customer engagements, capture lessons learned, and release v1.1 within 30 days with improvements and clarified templates.

Rationale: Prescriptive, artifact-driven delivery with stage gates, risk checklists, and clear acceptance criteria reduces project ambiguity and cancellation risk while increasing throughput and repeatability.

Action 2 — Establish AI-Ops with 24/7 Monitoring and Runbooks

Stand up centralized observability and incident management practices to ensure reliable, cost-controlled operations for agents, connectors, and workflows. Define SLOs, implement canary and rollback strategies, and maintain a 24/7 on-call rotation supported by detailed runbooks. Integrate cost monitoring at the “flight” level to prevent overruns and ensure predictable consumption.

- Deploy centralized observability covering logs, metrics, and traces for agents, orchestrations, and connectors; ensure environment separation for staging and production.
- Define SLOs per solution (e.g., availability, latency, cost per run, accuracy thresholds) and track them on dashboards with alerting thresholds.
- Implement health checks, heartbeats, retries, and circuit breakers for critical steps; automate restarts for transient failures.

- Adopt canary releases and blue/green deployments for agent and connector updates; document rollback procedures and test them before production use.
- Create a 24/7 on-call schedule with escalation paths; maintain runbooks for common incidents (integration failures, cost anomalies, model degradation, data access errors).
- Set up cost monitoring per flight with budget thresholds, daily spend alerts, and anomaly detection, and establish actions for throttling or pausing non-critical workloads.
- Define incident lifecycle targets (MTTA, MTTR), require blameless postmortems for Severity 1–2 incidents, and track corrective and preventive actions to closure.
- Conduct quarterly disaster recovery tests, including backup/restore verification and failover drills to validate recovery time and recovery point objectives.
- Integrate security monitoring and audit log reviews into AI-Ops routines, with a weekly evidence checklist and monthly trend reporting.
- Continuously improve runbooks using insights from incidents, near misses, and capacity trends; version and publish updates with change logs.

Rationale: Mature AI-Ops with strong observability, tested rollbacks, and disciplined incident response minimizes outages, controls consumption, and supports reliable scaling into production.

Action 3 — Governance Council, Policy Templates, and COMPLY-AI Ops

Operationalize governance by establishing a customer-facing governance council, issuing policy templates, and embedding COMPLY-AI outputs into routine operations. Standardize approvals, risk ratings, and evidence capture for changes to agents, data, and integrations. Maintain a continuous cadence of reviews, audits, and incident handling aligned to regulatory expectations.

- Charter a governance council with clearly defined roles (Executive Sponsor, Data Owner, Risk/Compliance, Operations Lead, Solution Architect) and a monthly meeting cadence.
- Publish policy templates for risk assessment, human-in-the-loop checkpoints, model and agent change control, incident management, data retention, access control, and third-party risk.
- Configure COMPLY-AI to generate periodic compliance reports, risk registers, and audit-ready evidence packages; schedule distribution and sign-offs.
- Implement a standardized approval workflow for new agents, datasets, and connector changes with documented risk ratings, mitigations, and required sign-offs before deployment.
- Define incident reporting and escalation procedures, including regulator notification criteria where applicable, and store all records in an evidence repository.
- Set audit log retention and access policies; conduct quarterly internal audits with remediation plans and progress tracking.
- Integrate governance checkpoints into delivery stage gates (e.g., compliance review required at Pilot exit and before Production cutover).
- Train customer champions and project teams on governance procedures and evidence expectations, and track completion and acknowledgement.
- Maintain a living risk register and periodically re-assess risks based on changes in use cases, data sensitivity, and operating metrics.

Rationale: Standardized governance with clear policies and automated compliance outputs builds trust, ensures auditability, and de-risks scale-up under evolving regulatory regimes.

Action 4 — Customer Enablement: Training and Certification via GPTBlue

Launch structured role-based enablement to close skills gaps and accelerate production-grade adoption. Provide curricula, labs, and certifications for admins, solution designers, operations, and business stakeholders. Support department-by-department rollouts with ready-to-use communication kits and measurable adoption goals.

- Define role-based learning paths: Administrator, Solution Designer, Ops Engineer, and Business User, with prerequisites and expected outcomes.
- Develop curriculum modules covering platform fundamentals, agent design, integration patterns, governance practices, AI-Ops procedures, and troubleshooting.
- Create hands-on labs in a sandbox environment and align exercises to real implementation artifacts (e.g., building a connector, authoring a runbook, configuring SLOs).
- Establish certification exams with passing criteria and digital badges; require relevant certifications before granting production access.
- Stand up “champion” programs in each customer department, including office hours, community channels, and a recognition framework.
- Provide rollout kits per department containing communications templates, SOPs, quick-start guides, and an adoption dashboard definition.
- Deliver training through blended formats (self-paced, instructor-led, and cohort-based sessions) and schedule quarterly refreshers.
- Track enrollments, completions, certification rates, and practical application in live projects; report progress in governance council meetings.
- Collect learner feedback and update content quarterly to address emerging needs, common pitfalls, and platform updates.

Rationale: Structured, role-based enablement and certification create consistent skills, reduce deployment friction, and empower customer teams to operate solutions confidently in production.

Action 5 — Integration Factory and Connector Backlog

Establish an integration factory to prioritize, build, and maintain connectors and integration guides for SAP estates, external data lakes, and common enterprise systems. Use a productized approach with versioning, SLAs, and documentation to reduce onboarding friction and support scale. Track adoption and quality metrics to guide backlog prioritization and lifecycle management.

- Create and continuously maintain a connector backlog prioritized by customer demand, business impact, and technical feasibility for SAP and other widely used enterprise systems.
- Organize delivery into stable squads with clear KPIs (e.g., time-to-release, defect density, performance benchmarks, documentation completeness).
- Define a Definition of Done for connectors: functionality, security review, performance thresholds, automated tests, versioned documentation, and sample configurations.
- Implement CI/CD pipelines with automated testing, semantic versioning, and a published compatibility matrix for platform and third-party versions.
- Publish integration guides, reference blueprints, sample payloads, and Postman collections to accelerate customer implementation.
- Set SLAs for connector maintenance and patching aligned with upstream vendor release cycles; announce deprecation schedules and migration paths.
- Provide a contribution and certification path for partners to submit connectors, including review criteria, security checks, and publishing standards.
- Instrument telemetry to track connector adoption, performance, and error rates; feed insights into prioritization and quality improvements.
- Establish a triage process for integration issues with severity definitions, response targets, and dedicated escalation channels.

Rationale: A disciplined integration factory with prioritized backlogs, strong documentation, and lifecycle SLAs removes connectivity bottlenecks and enables repeatable, scalable deployments.

Finance Department Findings

2025-10-03 — 12:00–21:45

Introduction

This report distills only those research findings that directly support the Finance actions listed in Cross Department Actions. The synthesis draws on multiple research reports conducted on 2025-10-03, covering customer feedback, technology and trend analysis, competitor pricing moves, governance/compliance risks, and BlueCallom's pricing model and product strategy.

Findings

Transparent, zero-subscription pricing with optional usage caps is a direct answer to buyer anxiety about unpredictable AI costs.

Market evidence shows enterprises are wary of transaction-based or opaque pricing that can trigger budget overruns. Keeping public pilot/runtime pricing, adding optional spend caps, and publishing budgeting guides align with explicit customer demands for clarity and predictability. This approach reduces procurement friction and accelerates approvals while differentiating against quote-only models.

Sources: [BlueCallom Pricing \(public pilots/runtime\)](#) | [TechTarget \(unpredictable, transaction-based costs\)](#) | [CIO \(pricing transparency expected\)](#)

ROI/TCO calculators with outcome benchmarks materially speed financial approvals.

Enterprises seek quantified business cases before scaling AI. Publishing ROI calculators and benchmark ranges (e.g., expected productivity uplift) directly supports Finance's evaluation process and aligns with the market push to close the "insight-to-action" gap—i.e., translating AI recommendations into executed value. Clear, defensible benefit ranges de-risk pilot funding decisions and shorten approval cycles.

Sources: [BlueCallom \(expected gains examples\)](#) | [TechRadar \(closing insight-to-action gap\)](#)

Targeted budget for compliance and reliability is a high-leverage risk mitigation.

The EU AI Act introduces potential fines up to 7% of global revenue for violations, and enterprise leaders fear customer loss from AI outages. Allocating budget to governance-by-design (COMPLY-AI), auditability, and reliability/observability features reduces financial exposure (fines, churn, remediation) and increases buyer confidence in mission-critical deployments.

Sources: [Reuters \(EU AI Act fines\)](#) | [ITPro \(outage/customer loss concerns\)](#)

Partner incentives to expand the out-of-the-box (OOTB) library lower internal cost and meet buyer demand.

Customers want richer prepackaged activities/templates to accelerate time-to-value. Incentivizing certified GPTBlue partners to publish agents/solutions into a shared catalog scales content breadth without heavy internal spend, directly addressing a documented shortfall in OOTB components and improving economic efficiency of solution rollouts.

Sources: [Biz.PRLog \(GPTBlue Development Studio ecosystem\)](#) | [TechTarget \(demand for more prepackaged activities/templates\)](#)

Price guardrails versus vendor bundling protect competitiveness without diluting monetization.

Major vendors (e.g., Oracle NetSuite) now bundle many AI capabilities at no extra cost. Finance can maintain competitiveness by bundling a baseline of AI features while keeping advanced, high-value usage ("flights") paid—leveraging BlueCallom's zero-subscription model to stay simple and clear while preserving margin on differentiated capabilities.

Sources: [Axios \(Oracle bundles AI at no extra cost\)](#) | [BlueCallom \(Zero-subscription model\)](#)

Finance Department Action Plan

Introduction

This action plan translates the Finance actions from the Cross Department Actions document into concrete, sequenced steps for BlueCallom AG. It focuses only on the defined Finance actions and the operational work Finance must lead to implement them. The content below specifies what to do, who to coordinate with, and how to control execution.

Actions

The Finance actions center on pricing transparency and control, decision-support tools for business cases, targeted investment in compliance and reliability, partner incentives to grow the out-of-the-box library, and competitive pricing guardrails. This plan defines the processes, policies, approvals, and monitoring needed to operationalize these areas. Execution is structured so Finance can drive change while enabling Sales, Marketing, R&D, and Operations to adopt the updates.

Action 1 — Reinforce Transparent, Zero-Subscription Pricing with Usage Caps

Finance will maintain public pilot and runtime pricing while adding optional spend caps and tiered volume discounts. The goal is to publish clear budgeting guidance and usage scenarios and to embed caps into contracting, billing, and alerts so customers can control consumption. Finance will also refresh internal tools and training to ensure the field quotes consistently and honor caps.

- Confirm and freeze the canonical list of public prices (pilot, runtime/flight, services) and publish the latest table to the pricing page and internal price book.
- Design optional spend caps with three options: per-flight cap, monthly cap, and departmental cap, each defined by currency, period, and enforcement behavior.
- Define tiered volume discounts with thresholds (e.g., flights per quarter) and a standard approval matrix for exceptions.
- Publish a Budgeting Guide that includes two to three sample usage scenarios per solution (e.g., COMPLY-AI, TRANSFORM) with inputs, expected consumption, and example monthly totals.
- Update the MSA and Order Form to include optional cap clauses, cap reset windows, customer alert contacts, and procedures for lifting or modifying caps.
- Configure billing and metering to enforce caps: create thresholds at 70%, 90%, and 100% of cap with automated alerts to customer and account team.
- Implement a hard-stop/soft-stop policy per customer preference: soft-stop continues service with overage at list rate; hard-stop pauses new runs until the cap is extended.
- Refresh quoting tools (CPQ or templates) to include cap selection, discount tier, and budget summary; validate that totals appear on one page for clarity.
- Run an internal enablement session for Sales, Customer Success, and Legal on how to position caps, how enforcement works, and how to request exceptions.
- Set KPIs and cadence: monthly review of cap adoption rate, percent of deals using public list price, average discount, overage incidents, and escalation volume.
- Establish a quarterly price governance meeting to consider adjustment of tiers, caps, or list levels based on win-loss, margins, and usage data.
- Publish a one-page “Customer Cap FAQ” and include it in the new-customer onboarding kit and renewal packages.

Rationale: Clear public pricing plus optional spend caps gives customers predictable control over costs while preserving our simple, zero-subscription model. Embedding caps into contracts, billing, and alerts makes the policy enforceable and reliable for Finance and customers alike.

Action 2 — ROI Calculator and Outcome Benchmarks for Pilots

Finance will sponsor the creation of ROI/TCO calculators and publish benchmark ranges for pilot outcomes to support fast, consistent business cases. The calculators will be simple enough for pre-sales use but rigorous enough for CFO review. Finance will own the

methodology, data inputs, and periodic updates, with delivery through the website and internal proposal templates.

- Appoint a Finance owner and cross-functional working group (Sales Ops, Product Marketing, Data) to define calculator scope, inputs, and outputs.
- Select two to three standard pilot scenarios per solution (e.g., compliance workflow audits, productivity transformations) and define measurable benefit drivers.
- Create a transparent benefits model including productivity time savings, error reduction, cycle-time improvements, and avoided manual rework.
- Specify data inputs required from prospects (headcount, volumes, baseline cycle times, hourly fully loaded costs) and default values for when data is missing.
- Design and document the calculation methodology, including sensitivity ranges and conservative/likely/aggressive outcome bands.
- Build a web-based calculator and a mirrored offline workbook for regulated or air-gapped buyers; ensure both produce identical results.
- Publish standardized output artifacts: one-page ROI summary, TCO breakdown, and a payback chart that can be inserted directly into proposals.
- Add a governed “Benchmark Ranges” appendix showing typical pilot outcomes by industry and use case; label assumptions and data vintage.
- Pilot the calculator with five recent deals; collect feedback from customers and account teams; refine defaults and presentation for clarity.
- Train Sales and Customer Success on the inputs to collect, how to run side-by-side scenarios, and how to position results responsibly.
- Establish quarterly model refresh: incorporate realized outcomes from live customers, update defaults, and archive previous versions for auditability.
- Instrument usage analytics and define KPIs: calculator adoption in opportunities, approval cycle time, pilot close rate, and variance between forecast and realized outcomes.

Rationale: A standardized, Finance-owned ROI model equips buyers and sellers with consistent, defensible numbers and shortens approvals. Benchmarked ranges set realistic expectations and reduce back-and-forth on assumptions.

Action 3 — Invest in Compliance and Reliability Roadmap

Finance will earmark budget to accelerate governance-by-design, COMPLY-AI enhancements, and reliability/observability capabilities. Funding will be tied to clear milestones, risk reduction metrics, and customer-visible deliverables. Finance will co-lead a stage-gate process with R&D and Operations to ensure benefits are realized and auditable.

- Create dedicated budget lines for: governance controls and auditability, COMPLY-AI features, reliability/observability tooling, and certifications or third-party assessments.
- Align with R&D on a milestone plan that maps budget tranches to specific deliverables (e.g., audit logs, SLO dashboards, incident reporting, report generators).
- Define financial guardrails: capex vs. opex allocation, contingency reserve, and criteria for reallocation if milestones slip or exceed scope.
- Establish stage-gates: release funds upon completion of design reviews, security/privacy assessments, and successful internal acceptance tests.
- Integrate Operations: ensure 24/7 monitoring, runbooks, and rollback plans are budgeted and resourced alongside product features.
- Include legal and compliance reviews in the plan to validate documentation, governance templates, and customer-facing compliance reports.
- Fund reporting automation so customers can self-serve compliance evidence (e.g., audit trails, human-in-the-loop records) without

manual Finance effort.

- Track risk and value KPIs monthly: open compliance gaps, time-to-detect and time-to-restore, incident count, and customer adoption of governance features.
- Schedule bi-monthly steering meetings (Finance, R&D, Operations, Legal) to review progress, unblock decisions, and approve next-gate funding.
- Prepare a customer communication pack summarizing governance and reliability enhancements for use in renewals, upsells, and risk reviews.
- Include a year-end post-implementation review to quantify reduced risk exposure and confirm that planned benefits were delivered.

Rationale: Proactive funding and stage-gated delivery of governance and reliability capabilities reduce operational and regulatory risks while increasing enterprise confidence. Tying releases to measurable milestones ensures Finance visibility and control.

Action 4 — Partner/Ecosystem Incentives to Expand OOTB Library

Finance will design and underwrite a partner incentive program that rewards certified GPTBlue partners for publishing high-quality agents and solutions into a shared catalog. Incentives will be linked to quality, adoption, and sustained usage to ensure durable value. Finance will define the funding pool, payment schedules, controls, and reporting.

- Set an annual incentive pool and quarterly release limits to manage exposure while enabling momentum.
- Define partner eligibility (certification status) and a tiered incentive structure covering submission grants, quality bonuses, and adoption/usage bonuses.
- Create a quality scorecard: documentation completeness, security review passed, usability standards met, and support commitment defined.
- Require a short business case with target industries, expected usage, and integration prerequisites for each submitted agent/solution.
- Establish a review board (Finance, Product, Ecosystem) that approves submissions, assigns provisional awards, and sets success metrics.
- Link payments to milestones: initial acceptance, first customer deployment, and 90-day sustained usage thresholds.
- Define revenue share rules for paid usage and ensure the model is reflected accurately in billing, reporting, and partner statements.
- Budget MDF (market development funds) for co-marketing of approved solutions with simple claim rules and proof-of-performance requirements.
- Publish legal terms: IP ownership, support SLAs, security/compliance obligations, and right to de-list if quality or compliance declines.
- Instrument analytics to track installs, active usage, customer satisfaction, and churn by solution; provide partners a monthly dashboard.
- Run quarterly business reviews with top partners to reconcile payments, validate metrics, and tune incentives for coverage gaps.
- Add a stop-loss clause to pause incentives in categories that do not meet adoption thresholds within two quarters.

Rationale: Paying for quality and adoption ensures the catalog grows where customers need it most while controlling spend. Clear milestones, analytics, and stop-loss rules keep the program economically efficient.

Action 5 — Competitive Pricing Guardrails vs. Bundled AI

Finance will maintain paid advanced “flights” while bundling a limited baseline of AI features at no extra cost to stay competitive with vendors that bundle AI. Guardrails will protect margins, ensure consistent discounting, and prevent confusion in the field. Finance will monitor competitor moves and adjust bundles and thresholds through a formal governance process.

- Define the baseline features to bundle at no added cost (e.g., lightweight insights, limited queries, or basic templates) and list

the exclusions reserved for paid flights.

- Create a simple packaging map that shows what is bundled, what requires paid flights, and examples of when customers cross into paid usage.
- Update SKUs, price lists, and MSA exhibits to reflect bundled inclusions and the metering logic for advanced usage.
- Establish discount guardrails: maximum cumulative discount by segment, approval thresholds, and non-discountable items to protect value.
- Set up a Pricing Governance Board (Finance-led) that reviews win-loss, competitive offers, and margin impacts monthly and authorizes changes.
- Instrument unit economics by package to ensure bundled features do not drive disproportionate support or infrastructure cost.
- Provide enablement to Sales and CS with clear talk tracks, FAQs, and objection handling to avoid accidental over-bundling.
- Run controlled A/B offers in select segments to measure conversion lift from bundling versus margin impact; feed outcomes into governance.
- Publish a “Competitive Offer Review” intake to handle non-standard competitor bundles; respond within 48 hours with approved counters or alternatives.
- Monitor KPIs monthly: attach rate of paid flights, average selling price, win rate versus bundled competitors, and gross margin on flight usage.
- Refresh the bundled baseline and thresholds quarterly to reflect product evolution and market dynamics without confusing customers.

Rationale: Selective bundling keeps us competitive while preserving monetization for high-value usage. Governance and analytics ensure we adapt quickly without eroding margins.