

Background

People at work are drowning in notifications, but starving for clarity on what actually matters.

On any given day, an employee is pinged across Slack, email, Jira, Google Docs, calendar invites, and more. Each app has its own notification system, its own priority logic, and its own way of demanding attention. The result: people spend their mornings scanning 5-6 different tools just to figure out "what do I need to do today?"

This creates real problems. Important work slip through the cracks. Before any real work starts, most people spend 30-40 minutes each morning just catching up between apps, mentally sorting what's urgent vs. what can wait. There's no single source of truth for "what needs my attention." Every tool thinks it's the most important.

Problem Statement

How can Glean, use its access to 100+ work apps, enterprise knowledge graph, and AI capabilities to become the single intelligent surface that helps people stay on top of what matters?

Why Past Products Failed?

Every major platform has attempted attention management.

Product	Approach	Why It Failed
Google Inbox	Smart categories, bundles, snooze	Email-only scope; no cross-platform intelligence
Cortana Briefing	Daily email digest from O365 data	Passive delivery; no cross-platform signals; deprecated 2023
Unified Inboxes	Aggregate email + chat + social (Spike, Front)	Aggregated noise without intelligence; created another destination to check
Motion / Reclaim AI	AI calendar + task auto-prioritization	Calendar-centric. Optimizes time blocks but doesn't surface cross-app attention signals.

Why to Solve now?

Employees lose significant time every day to manual triage and context switching. Reducing this directly improves productivity, lowers missed deadlines, and helps teams execute faster and higher ROI.

Aggregating notifications was always the easy part, making sense of them was hard. With the latest AI models, we can now understand intent, urgency, and context from unstructured messages. An LLM can read a Slack thread and get context that this is a blocker from your manager on a project due Friday and assign importance.

Glean already connects to 100+ work tools, understands permissions, and maps people, projects, and documents. The Enterprise Graph maps people, projects, and teams. The Personal Graph knows what you're working on and how you operate. This means Glean can respond intelligently, based on who's asking, what project it's tied to, and your deadlines.

User Personas

In most organizations, work is divided in three broad ways: coordinating work, executing work, and making decisions. Regardless of job title, employees usually fall into one of these roles based on how they spend their time and how they handle attention and priorities.

- Coordinators focus on aligning people, projects, and timelines.
- Executors focus on completing core work with high concentration.
- Decision-makers focus on approvals, escalations, and strategic direction.

In a technology company, these roles are commonly represented by Product Managers, Developers, and Senior Management.

	Product Manager (Coordinator)	Developer (Executor)	Senior Management (Decision-Maker)
Tools	Slack (15+ channels), Gmail, Google Calendar, Jira, Confluence, Google Docs, Figma, Notion	Slack (5-8 channels), GitHub, Jira, IDE, Google Docs, Calendar	Slack (10+ channels), Gmail, Google Calendar, Dashboards, Google Docs, Salesforce
Pain	Spends 45 min every morning triaging across 8 apps; misses dependency blockers buried in Slack threads; forgets to follow up on stakeholder emails; context-switches 100+ times/day	Gets pulled out of deep work by constant pings; can't distinguish real blockers from noise; PR review requests get buried; spends 30 min before writing first line of code just catching up	Tagged in 50+ threads daily but only 5 need a decision; approvals sit unnoticed blocking entire teams; spends an hour daily filtering signal from noise across email and Slack
Goals	Know what's stuck, what's waiting on me, and what's at risk across all projects — without checking every app.	Get a clear "here's what actually needs you" list and silence everything else during focus time.	See only the decisions and escalations stuck on me, with enough context to act immediately.

The PM sits at the intersection of engineering, design, leadership, and stakeholders, when they miss something, entire teams get blocked. Improving their attention management has a multiplier effect across the organization. They're also the most measurable: we can track decision latency that is, time from when an issue surfaces to when the PM acts on it. The other personas are weaker starting points as, executives often already have assistants who map out their day, and developer workflows are shifting rapidly with AI-assisted coding.

The PM is where the pain is sharpest and the impact is widest.

North Star Vision - Glean Pulse (12 Months)

User Flow

Checkpoint	JBTD (Platform)	User Job	Success Metric
1. System Login	Authenticate user, sync with all connected apps, compute initial priority queue within 30 seconds of login	"I want to know at a glance if anything urgent happened without opening any app"	>90% of sessions where bubble loads within 30s of login
2. Open floating Dashboard	Maintain a persistent, low-friction entry point across applications and rank all pending items; surface missed @mentions and upcoming events	"I want a single view that tells me what needs me right now, what I missed, and what's coming up"	>70% of users open dashboard at least once per morning
3. Take Action in Dashboard	Route user actions back to source platforms (Slack reply, Jira update, email send); use LLM to draft messages and create tasks from natural language input	"I want to respond, approve, or delegate without switching to another app"	>50% of action items resolved from within dashboard
4. Exit & Resume	Persist state; continue monitoring in background; update priority queue in real-time	"I've triaged everything, now let me focus without worrying I'll miss something"	>80% of users report feeling "caught up" after closing (survey)
5. Ambient Nudge	Evaluate incoming signals in real-time; only interrupt for: urgent messages from key people, time-sensitive approvals, overdue escalations	"Don't bother me unless it's truly important"	<5 nudges per day per user; >80% of nudges acted on (low nudge = high precision)
6. Revisit on Demand	Re-rank queue based on latest signals; generate "while you were away" summary if user was in a meeting or away >1 hr	"I want to quickly catch up after a meeting or focus block without re-triaging from scratch"	>3 dashboard opens per user per day

THE ONE-LINER

Glean Pulse is a system that tells you what matters across every connected app, with full context of the project, persons, deadlines & enables them to triage, respond, and take action without leaving the dashboard.

Features

Feature	Pain Point Addressed	Opportunity (Glean + AI)	Tradeoffs
Floating Bubble + Dashboard	PM doesn't want another app; attention layer must be ambient	- Glean browser extension already supports overlays and new tab takeover - Needs to be floating on desktop, to feel like live dashboard and "not another app"	- Desktop install = adoption friction - Needs to be responsive and snappy for high adoption
AI Priority Ranking	PM gets a message but can't tell if it's a P0 blocker or casual FYI	- Enterprise Graph → sender seniority, reporting chain via org chart - Personal Graph → PM's active projects, deadlines - LLMs classify urgency combining: who + deadline + project + sentiment	- Bad ranking erodes trust fast - Needs "why is this ranked high?" explanation - Cold start mitigated by existing Personal Graph data
Cross-App Deduplication	Same issue appears as Slack thread + email + Jira comment = 3 items for 1 event	- Glean indexes 100+ apps, raw data for overlap detection already exists - Enterprise Graph maps project-ticket-people relationships	- False merges / missed connections dangerous - Will only improve with user corrections
"What You Missed" Temporal Tracking	PM checks Slack but not email; dashboard shows everything again including already-handled items	- Connectors ingest read receipts, view timestamps, interaction logs in near real-time - Need to maintain persistent state of as "seen" and "not-seen" categories	- Unreliable persistent memory state - May miss something important
Inline Actions + Sync Confirmation	PM acts from dashboard but has no confidence it actually updated on Jira/Slack	- Extend bidirectional sync capabilities: execute on dashboard → platform acknowledgment → show status in real-time	- Not all platforms confirm reliably (Slack fast, Jira lags) - Delay needs to be kept minimum and all guidelines to be followed while filling
Deep Links to Source	PM sees a comment but finding the exact location in source platform is friction	- Connectors already store source URLs for all indexed content - Surface per-item deep links	- Links break if content moves/deleted - Some platforms have inconsistent deep link support
Contextual Summary per Item	PM sees a comment but doesn't know: who is this, what project, why now?	- Enterprise Graph, Personal Graph Model Hub LLM generates grounded one-liner: "Rahul (Backend Lead) — checkout bug follow-up, sprint ends tomorrow"	- LLM might hallucinate context - Every summary must link to verifiable source
Natural Language Task Creation	Creating tasks requires opening Jira/Slack separately; breaks flow	- Help user in filling a connector platform forms, recalling context from projects, deadline information, feature requested in meetings	- Needs confirmation step ("Is this right?")
Smart Notification Routing	Casual FYI and P0 blocker both ping the same way — or not at all	- Priority score + real-time user state: Calendar API (in meeting?), Personal Graph (focus pattern?) - Routing notifications: passive (badge) → light (bubble pulse) → heavy (push/Slack DM via Actions)	- Too aggressive = user disables - Too passive = misses critical items - Needs per-user calibration
Behaviour Personalization	One-size-fits-all urgency model doesn't work; preferences differ by person and time of day	- Every snooze/dismiss/engage feeds back into Personal Graph - Learns: time-of-day patterns, topic preferences, focus blocks - Model Hub fine-tunes priority scoring per user	- Takes 2-4 weeks to become useful - Over-fitting risk (dismiss one VP ping → deprioritize all VP pings)

MVP Scope (3 Engineers, 3 Months)

4 features selected based on: **high impact** on PM's daily pain and feasible with 3 engineers in 3 months. Features like cross-app deduplication and behaviour personalization need longer data cycles, better suited for V2.

AI Priority Ranking launches in Month 1 as a web view inside the existing Glean browser extension, no new install. This gets core value out immediately and generates behavioural data that makes Month 2-3 features smarter.

The real problem: Adoption

- The biggest risk isn't building the product, it's getting PMs to change their morning routine
- Prove value with data. Send emails: "You saved 3.2 hours this week using Glean." Concrete proof drives retention.
- Tie to KPIs. Track decision latency improvement so when PMs can show their manager "I'm responding to blockers 40% faster," adoption becomes self-reinforcing.
- Seed with early adopters. Launch Month 1 to power-user PMs. Use their time-saved metrics as internal proof for wider rollout.

MVP Feature	User Story	What success looks like?
AI Priority Ranking (Ship in Month 1 as web view in browser extension)	As a PM, I want to see all items needing my attention ranked by urgency in one list, so that I don't waste 45 min every morning triaging across 6 apps.	- PM opens one view and knows their top 5 priorities within 60 seconds
Link to Source App (Ship in Month 1)	As a PM, I want to open the exact message, ticket, or document from the dashboard with one click, so that I don't waste time searching inside other tools.	- PM reaches the correct source location (thread, comment, ticket) in one click
Smart Notification Routing (Ship in Month 2)	As a PM, I want to be nudged only when something is truly urgent, so that I can focus on deep work without worrying I'll miss a blocker.	- >80% of nudges are acted on (high precision) - PM feels in control of notifications
Floating Bubble (Ship in Month 3)	As a PM, I want an always-visible bubble on my screen that shows me how many items need attention, so that I don't have to remember to check any app proactively.	- PM interacts with bubble multiple times per day without needing to open Slack/Email/Jira directly
Contextual Summary per Item (Ship in Month 3)	As a PM, I want to see the context behind every item, who, what project, why now, so that I can decide how to act without digging into each platform.	- Context summaries are accurate and grounded

Metrics to Track

Input Metrics (Leading — tracked from Month 1)

- **DAU/WAU ratio** : Are PMs coming back daily? Target: >60% by Month 3.
- **Dashboard opens per user per day** : Is the dashboard becoming the PM's starting point? Target: >3 opens/day.
- **Nudge-to-action rate** : When we nudge, does the PM act? Measures notification precision that is, how often our nudges are actually useful. If we nudge 10 times and the PM acts on 8, our precision is 80%. If they ignore most nudges, we're sending noise, not signal. Target: >80% acted on within 30 min.

Output Metric

- **Productivity increase** : Is the time from a request being created to it being completed reducing? Measured through TAT between request created to completed (average time from task/request creation to resolution) and for a more holistic view number of projects where deadlines are breached (percentage of projects missing committed timelines). Together, these indicate whether better attention management is translating into faster execution and more reliable delivery. Target is ~10-20% reduction in TAT and ≥20% reduction in deadline breaches.


Open Questions

MVP Feature	User Story
How do we handle ranking disagreements?	Start with a thumbs up/down feedback mechanism per item. The model adapts gradually so that single overrides do not heavily influence rankings. Define a minimum number of consistent signals before changing priority patterns.
What's the right nudge threshold?	Begin conservatively with minimal active nudges in Month 1. Allow users to opt into stronger notifications. Introduce a "notification confidence score" and trigger nudges only when urgency confidence is above 90%.
Browser extension vs. native desktop app?	Launch first as a browser extension overlay using existing Glean infrastructure. This enables faster rollout, no install friction, and cross-platform support. Limitation: no presence when the browser is closed. Re-evaluate native app investment after 6 months based on adoption.
What happens when connectors lag or fail?	Display "last synced" timestamps on the dashboard. Show warning banners when connectors are delayed for more than 10 minutes. Explore caching last-known state to avoid empty dashboards during outages.

Explicit Scope Cuts (Not in MVP)

- **Cross-app Deduplication** : Requires extensive entity resolution tuning and user correction data; ship after 3+ months of usage data
- **"What You Missed" Temporal Tracking** : Read-status ingestion unreliable across platforms; needs connector-level improvements first
- **Inline Actions + Sync Confirmation** : Write-back APIs vary by platform; start with read-only dashboard, add actions in V2
- **Natural Language Task Creation** : LLM misinterpretation risk too high without enough user context; defer until Personal Graph is richer
- **Behaviour Personalization** : Needs weeks of behavioural baseline per user; can't deliver value on day one; defer 3-6 months
- **Deep Links to Source**: Direct link to connected platform is implemented. Link to relevant project files across entire database being linked automatically based on context, is being left out.
- **Mobile Surface** : Start with desktop browser extension; mobile in Phase 2
- **Auto-drafting Replies** : Too much trust risk in V1; users need to trust the ranking before trusting AI to act on their behalf

Prototype & User Flows – [Link to Prototype](#)



- The user logs into their laptop and sees a persistent floating bubble on the screen.
- Clicking the bubble opens a unified dashboard that aggregates updates from Slack, Email, Jira, and Docs, ranked by priority.
- The dashboard highlights urgent items, missed signals, and upcoming deadlines with AI-generated summaries and context.
- Users can open source threads, documentations directly from this view without switching apps.
- After resolving items, the user closes the dashboard and resumes work with minimal context loss.
- Throughout the day, high-confidence urgent updates are surfaced via subtle alerts (red), while non-critical items remain passive(golden)

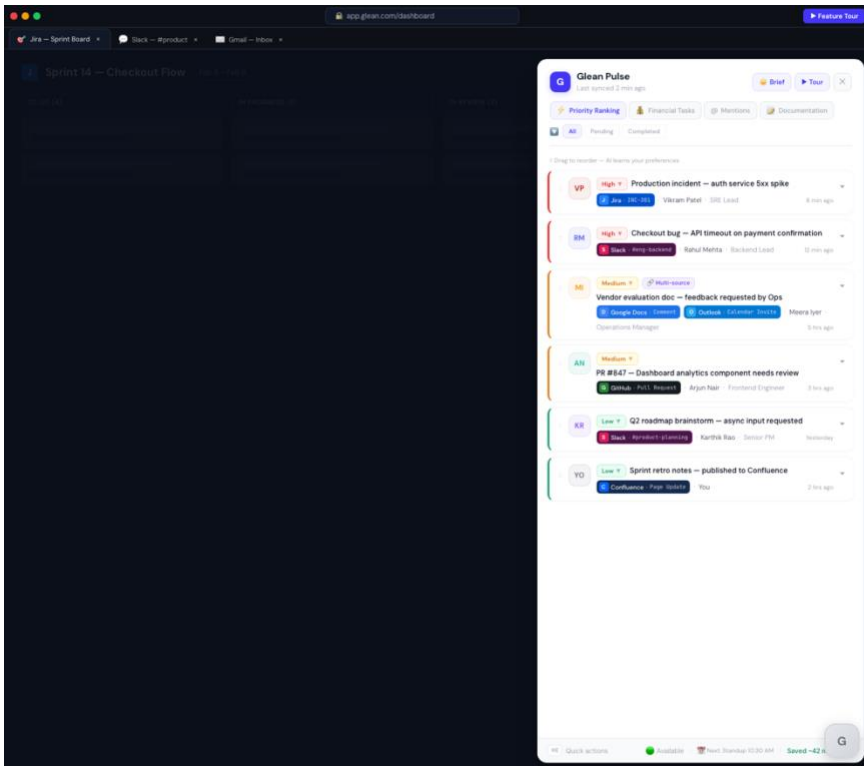


Figure 1 - Priority Ranking : Default dashboard view showing items sorted high → medium → low, with color-coded left borders for instant visual triage. Floating bubble (bottom-right) persists across any app/work.

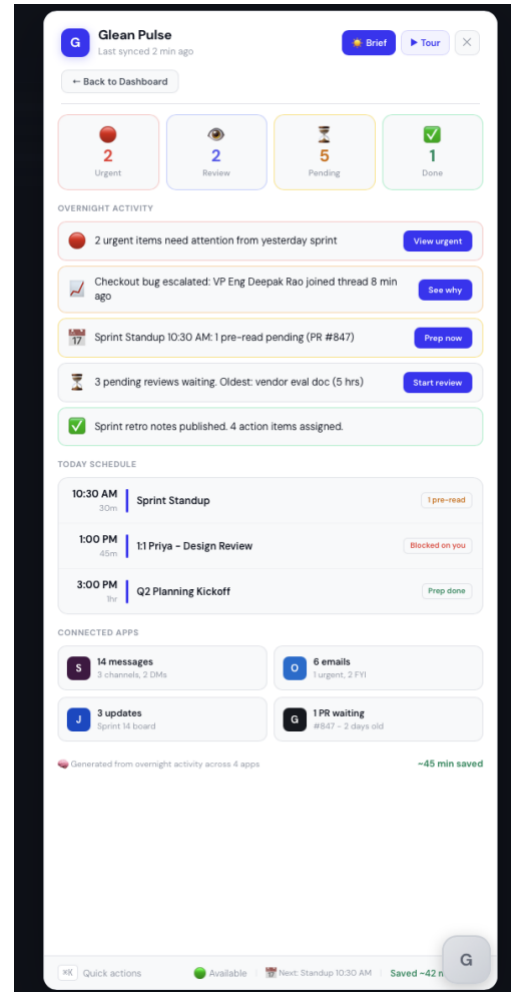


Figure 2: Morning Brief , an AI-generated overnight summary showing urgency stats, activity feed with actionable CTAs, today's schedule with blocker flags, and connected app status across Slack, Outlook, Jira, and GitHub.

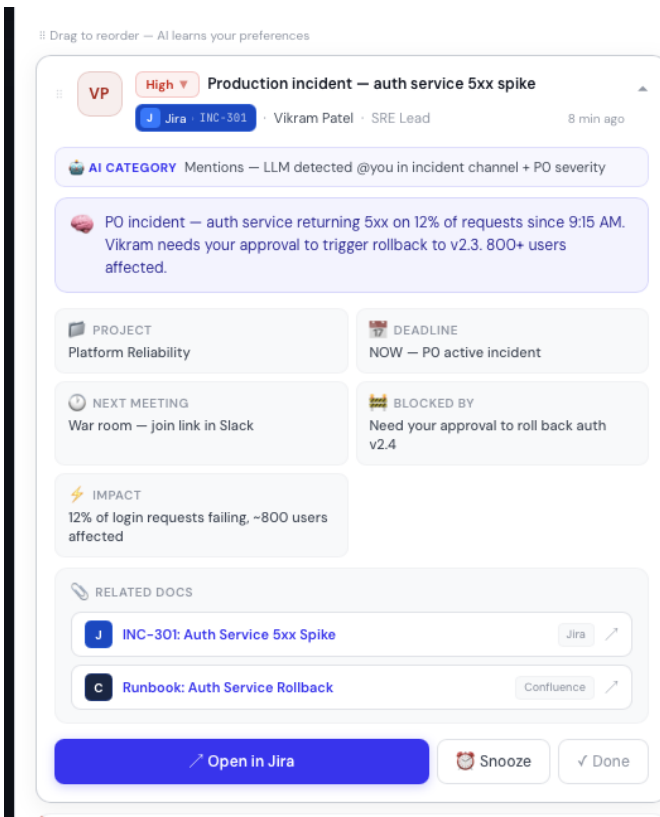


Figure 3: Expanded Card. Full context pane for a P0 incident. AI-generated summary, structured context grid (Project, Deadline, Next Meeting, Blocked By, Impact), related docs with platform icons, and inline action buttons — all without leaving the dashboard.