



STANDARD OPERATING PROCEDURE

Severity Tier System

How we classify work – P0 / P1 / P2 / P3 + Stability

Owner	Aman · Product FleetRobo
Date	12 May 2026
Version	1.0 (Draft)
Status	For review — PMO Lead, Product PMs, Eng Lead

INSIDE THIS DOCUMENT

- 1. **Why this exists**
- 2. **Four dimensions** of severity — the scoring rubric
- 3. **Four tiers** — P0, P1, P2, P3
- 4. **Classification flow** — who proposes, confirms, vetoes
- 5. **Promotion & demotion** rules
- 6. **Stability Track** — chronic infra carve-out
- 7. **VIP Override** — controlled escape hatch
- 8. **Worked example** — the 11 PMO items re-tiered
- 9. **Cheat sheet** — 60-second decision tree
- 10. **Metrics** for tier-system health
- 11. **Rollout plan**
- 12. **Open items**

DISTRIBUTION
 Product · PMO · Eng · QA · CS

CLASSIFICATION
 Internal Use

NEXT REVIEW
 2026-06-09

A living document. Use it, break it, tell me where it breaks. v1.1 in 4 weeks.

How We Classify Incoming Work – and What Happens At Each Tier

Companion document to: PMO ↔ Product ↔ Engineering Delivery Process SOP v1.0 **Version:** 1.0 (Draft) **Date:** 2026-05-12

0. Why This Exists

Today, every painful item gets labelled "P0." When everything is urgent, nothing is. Engineering lives in firefighting mode, real incidents get diluted, and there is no shared language between PMO, CS, Sales, and Engineering for *how bad is this actually*.

This document fixes that with:

1. **A scoring rubric** — four dimensions of severity, not one
2. **Four tiers** — P0, P1, P2, P3 — each with explicit SLA, comms, and capacity reservation
3. **A classification flow** — who proposes, who confirms, who can veto
4. **Promotion + demotion rules** — tiers move both directions, and fast
5. **A Stability Track** — chronic infra debt does not belong in the severity ladder
6. **A VIP override gate** — the realistic escape hatch, controlled

1. The Four Dimensions of Severity

Every item is scored on these four axes **before** being assigned a tier. No score → no tier → no ticket movement.

Dimension	What it measures	Scoring
Blast Radius	How many are affected right now	High: >10% of fleet OR >5 clients OR all users on a module Medium: 1–5 clients OR single module degraded for many Low: Single client OR single user
Functional Impact	What the user cannot do	Outage: Cannot use the platform / module at all Degradation: Can use it, but slow / partial / requires workaround Annoyance: Works, but ugly / inconvenient
Business Impact	Money + contract risk	Severe: Active SLA penalty clause triggered OR churn signal OR revenue stopped Material: Client visibly frustrated, escalation risk Minimal: Internal-only OR no client visibility
Recoverability	Workaround availability	None: No way around it, data loss possible Manual: Workaround exists, painful, not scalable Easy: Clean workaround, low cost

Plus a time gate:

- **Acute** = symptom started within the last 2 hours OR is actively worsening
- **Chronic** = has been broken for >24 hours OR has recurred ≥ 3 times this month

Acute issues climb the tier ladder. Chronic issues drop to the **Stability Track** (see §5). This is the most important rule in this document.

2. The Four Tiers

Tier-at-a-Glance

Tier	Name	Response SLA (Phase 1)	Dev Build Window (Phase 2)	Rollout Type (Phase 3 default)	Capacity	Comms Cadence
P0	Critical Incident	30 min ack · 1 h confirm · 2 h in motion	Same-day hotfix	Type 1 — Direct	Hot-swap from current sprint	Every 60 min until mitigated
P1	High Impact	4 hours	Next Tuesday Dev Release	Type 2 — Notified	20% sprint reserve	Daily until released
P2	Standard	2 business days	1–2 sprints (Dev decides)	Type 2 or 3 — PM picks	30% sprint reserve	Weekly status
P3	Low / Backlog	5 business days	Next quarterly planning	Type 2 or 3 — PM picks	Surplus capacity	Monthly summary

Three SLAs, not one. Response SLA is how fast PM picks up the ticket. Dev Build Window is how fast Dev ships internally. Rollout pattern is how the client experiences it. The three are separate by design — see PMO ↔ Product ↔ Engineering SOP v1, §2 "The Three Phases."

P0 – Critical Incident

Definition: Production is on fire. Drop everything.

Qualifies only if ALL of:

- Blast Radius = **High** (>10% fleet OR >5 clients OR full module unusable)
- Functional Impact = **Outage** (not degradation)
- Recoverability = **None** OR **Manual-only** with no scalable workaround
- Time gate = **Acute** (started <2h ago OR actively worsening)

AND at least one of:

- Active SLA penalty clause is being breached right now
- Data is being lost or corrupted
- A paying client cannot run their business

Response (realistic for BSL today — single Eng Lead, no 24/7 rotation):

- Acknowledged in `#incidents-bsl` within **30 minutes**
- On-call Engineering Lead pulled in immediately
- Severity confirmed (P0 vs P1 demoted) within **1 hour**
- Mitigation work explicitly in motion within **2 hours**
- Current sprint work paused for affected squad
- War-room call at the 1-hour mark if not mitigated
- If Eng Lead unreachable beyond 30 min, incident escalates to Head of Product who pages a backup

Comms (to client):

- Initial acknowledgement within **1 hour** of detection
- Status update **every 60 minutes** until mitigation
- Mitigation announcement same channel
- Full RCA within 72 hours

Capacity: Hot-swap. Any developer can be pulled. Current sprint owners notify their PM that their ticket slips.

Examples (real-world):

- Login broken for all clients
- Server unreachable / 500 errors across the platform
- Data ingestion stopped (no telemetry coming in)
- Mass vehicle offline event after a deploy (e.g. 1,000+ vehicles)
- Trip creation completely failing for all users
- Active client SLA penalty being triggered right now

NOT P0 (common mistakes):

- One client can't log in → P1 (single-client outage)
- Vehicles randomly offline daily (chronic) → Stability Track
- Page is slow → P2 (degradation, not outage)
- Auto-trip-closure not working but manual closure works → P1 (workaround exists)

P1 – High Impact

Definition: Something important is broken, but the platform is up. Fix it inside one sprint cycle.

Qualifies if:

- Blast Radius = **Medium or High**
- Functional Impact = **Outage OR Degradation**

- Business Impact = **Severe or Material**
- A workaround exists (manual or otherwise) so we are not bleeding right now
- Time gate = **Acute** OR recently-acute (within the last 7 days)

Response:

- PM owner assigned within **4 hours** of intake
- Triaged in the next Tuesday Intake Sync
- Locked into the next 7-day cycle (released the following Tuesday)
- Hot-flagged in Odo

Comms (to client):

- Acknowledgement within **4 hours**
- Daily status until released
- Release notification on the day of fix

Capacity: 20% of every dev squad's weekly capacity is reserved for P1s. This is non-negotiable. If P1 volume exceeds 20%, the overflow goes to P2 OR triggers a capacity escalation.

Examples:

- Dashboard data missing for one client
 - Single-client login failure
 - Report column showing wrong value for a paying customer
 - Trip creation broken for one client's vehicle type
 - API endpoint failing for one integration partner
-

P2 – Standard

Definition: Real issue, real fix, real timeline. No drama.

Qualifies if:

- Blast Radius = **Low or Medium**
- Functional Impact = **Degradation** OR **Annoyance** on a paid feature
- Business Impact = **Material or Minimal**
- Workaround exists

Response:

- PM owner assigned within **2 business days**
- Triaged in Tuesday Intake Sync
- Scheduled into the next available 15-day cycle slot
- Goes through full Definition-of-Ready and Definition-of-Done gates

Comms (to client):

- Acknowledgement within **2 business days**
- Weekly status
- Release notification on the day of fix

Capacity: 30% of every dev squad's weekly capacity is reserved for P2 work.

Examples:

- Video Telematics history view is slow
 - Dashboard widget shows stale data after refresh
 - Filter doesn't persist across sessions
 - Small UX bug on a paid feature
-

P3 – Low / Backlog

Definition: Worth fixing eventually. Not worth disrupting a sprint for.

Qualifies if:

- Blast Radius = **Low**
- Functional Impact = **Annoyance** OR cosmetic
- Business Impact = **Minimal**
- Clean workaround exists, or the feature has low usage

Response:

- Logged in Odoo
- Reviewed in monthly product review for prioritisation
- Scheduled when sprint capacity allows

Comms (to client):

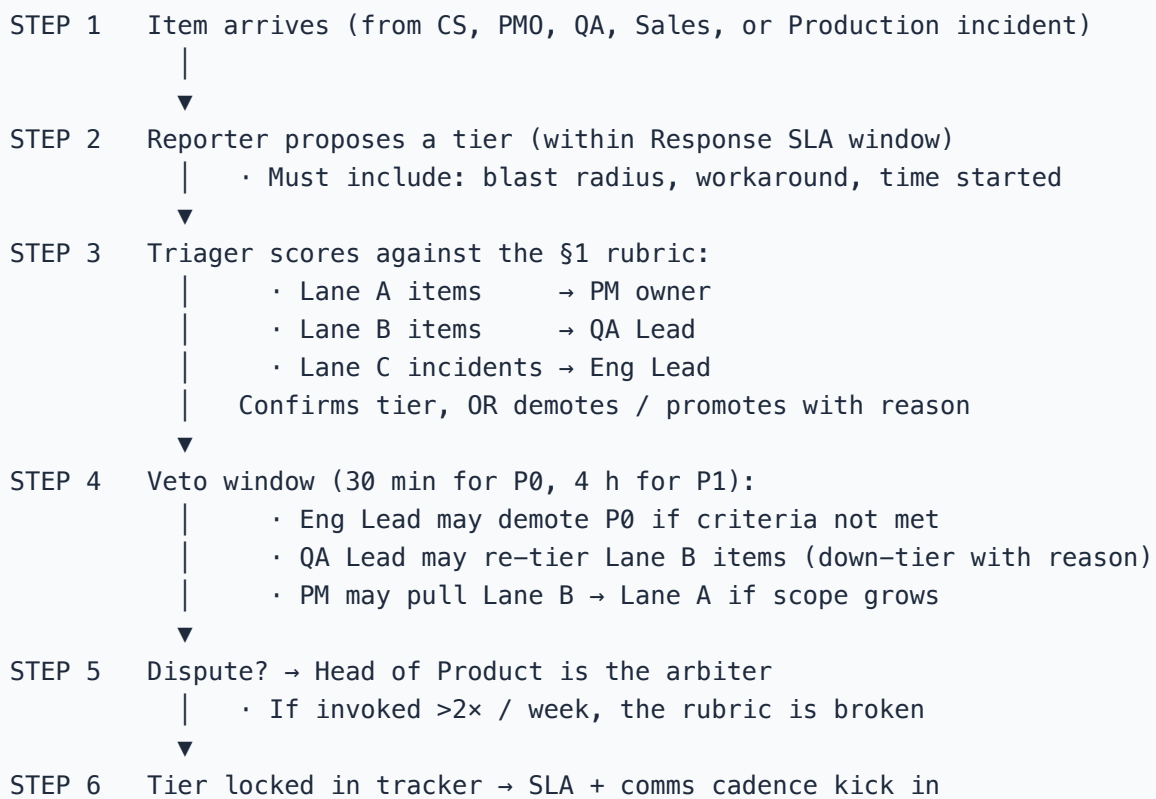
- Acknowledgement within **5 business days**
- "We have logged this and will consider it in our roadmap"
- Monthly summary if requested

Capacity: Surplus only. Picked up when P0/P1/P2 buffer is clear.

Examples:

- Typo in a label
 - Button alignment off by 2px
 - Old field still showing in an unused admin screen
 - Feature request from a single low-tier client
-

3. Classification Flow



Ownership:

- **Proposer:** Whoever raises the ticket (CS, PMO, QA, Sales) — proposes a tier with justification
- **Triager:** PM (Lane A) or QA Lead (Lane B) — confirms or re-tiers within the SLA window
- **Vetoer (P0):** Engineering Lead — for P0 only, can demote within 30 minutes if criteria not met
- **Vetoer (Lane B):** QA Lead — can down-tier Lane B items with documented reason
- **Arbiter:** Head of Product — final say on disputes

Hard rule: Proposer cannot override Triager. Triager cannot override Vetoer. Either party can escalate to Arbiter.

Friction signal: If the Arbiter is being pulled in >2 times per week, the rubric is broken — schedule rubric review.

4. Promotion + Demotion Rules

Tiers move both directions. Aggressively.

Promotion (something got worse)

A ticket is **promoted up a tier** when:

- Blast radius grew (more clients now affected)
- Workaround broke or stopped working
- Business impact escalated (SLA clause triggered, client threatening churn)
- Active worsening trend

Owner: PM owner notifies PMO + Eng Lead immediately. Tier change logged with reason.

Demotion (it turned out to be smaller)

A ticket is **demoted down a tier** when:

- Investigation shows blast radius was overestimated
- Workaround discovered
- Issue turns out to be field/SIM/hardware (not platform) → routes to field-tech, not engineering
- Time has passed (Acute → Chronic, see below)
- Item has been "P0" for >24 hours without mitigation → automatic demote to P1 (because by definition, it's no longer a same-day fix)

Owner: PM owner OR Eng Lead. Demotion is logged with reason. PMO is notified within the response SLA of the new tier (so the client comms cadence updates).

The 24-hour rule: An item cannot stay P0 indefinitely. If it has been P0 for 24 hours, it either gets mitigated or it gets re-tiered. No exceptions. This prevents "P0 inflation" where tickets sit at the top forever.

Recurring P0 → Stability Track

If a P0 category fires **≥3 times in 30 days**, it stops being an incident. It becomes a **Stability Track item** (§5). Engineering owns the root cause, not the symptom-fix.

5. The Stability Track

Severity tiers are for **acute** issues. Chronic infrastructure debt belongs somewhere else.

Why this matters: PMO's current P0 list contains items like "vehicles randomly offline daily" and "operational pages hanging during live use." These are NOT incidents — they are *predictable consequences of the current architecture*. Labelling each flare-up as P0 burns engineering and never fixes the root cause.

The Stability Track is:

- A parallel work stream owned by Engineering Leadership (not by ticket queues)
- Funded by **20% of every dev squad's weekly capacity** (yes, on top of the 20% P1 reserve, so 40% total reserved for non-feature work)
- Visible in a separate Odo project, not in the severity queue
- Reviewed bi-weekly with Engineering Leadership as part of the architecture roadmap

Items that go here:

- Chronic vehicle-offline pattern (the 2,000–2,500 daily offline issue)
- Platform slowness / page-load times >5s
- Session timeout misconfigurations
- Cross-product login inconsistencies
- Recurring data-sync gaps
- Any P0 that has fired ≥ 3 times in 30 days

This track is linked to the broader infra plan — Kafka + Kafka Stream + Redis + S3 data lake + PostgreSQL migration. The Stability Track items either get fixed in patchwork now OR get absorbed by the migration. Engineering decides which.

Communication to PMO: "This is in the Stability Track. Engineering owns the root cause. You will not get a 7-day ticket-level ETA, but you WILL get a bi-weekly progress update." This is the conversation that ends "everything is P0."

6. The VIP Override

Reality: Sometimes a big client calls a VP, the VP calls Engineering directly, and the tier system gets bypassed. If we don't have a controlled path for this, people will do it informally — which destroys the system.

The VIP Override is:

- An explicit path to bypass the tier rubric
- Gated by **Head of Product approval** (not a VP, not a PMO request, not a CS escalation — only Head of Product)
- Logged with: requesting party, business justification, expected resolution, displacement cost
- Has a maximum frequency: **2 active VIP overrides at any time, platform-wide**
- Auto-expires after 7 days (can be renewed once with re-approval)

Rules:

- A VIP Override does NOT make a ticket P0 — it gets fast-tracked through whatever tier it actually belongs to with extra capacity attached
- The displacement cost (what got de-prioritised) is documented and communicated to the displaced work's owner
- VIP Overrides are reviewed in the monthly product review — if any party is using this path >1x/month, that pattern is escalated to leadership

Why this gate matters: Without it, the system collapses on the first big-client escalation. With it, the escape hatch is auditable and rare enough to stay healthy.

7. Reclassifying the Current PMO List

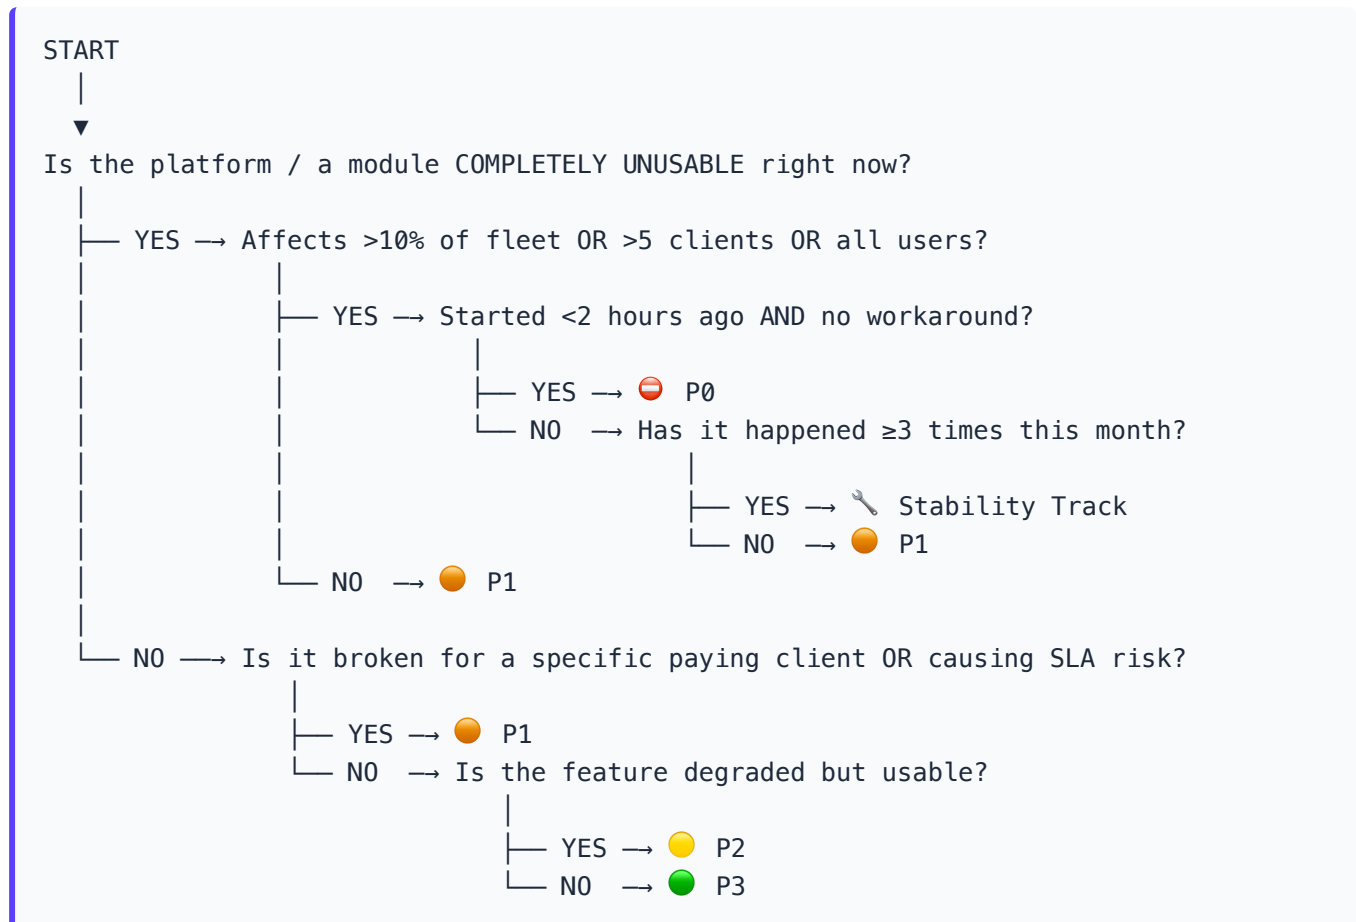
A worked example using the 11 items PMO labelled as P0.

#	Item	Real Tier	Reasoning
1	Trip Creation Failure — Server Error	P0	All users on the module, outage, no workaround, acute
2	Auto Trip Closure Failure	P2	Manual closure works (workaround exists), no outage, no SLA risk
3	User Login Failure — Auth Error	P0 if cross-client, P1 if single client	Tier depends on blast radius — confirm at intake
4	Dashboard Data Not Visible	P1	Data exists, display issue, single-module degradation
5	All Vehicles Showing Offline	Stability Track	Recurring pattern, chronic. Root cause is infra, not a ticket
6	Server Error / App Page Not Loading	P0 if acute, Stability if chronic	If the platform is down right now → P0. If pages are slow daily → Stability
7	Trip Creation / Submission Error	Merge with #1	Same root issue, deduplicate
8	Data Sync Failure (device ↔ server)	P1 if mass, field-tech if single device	Triage step: is it the platform or one device?
9	Operational Page Hanging During Live Ops	Stability Track	Performance issue, chronic, structural fix needed
10	API — Issues with Running APIs	Not classifiable	Too vague. Needs which API, which client, what symptom — kick back to PMO for detail
11	VT Dashboard — Live Streaming / History	P2	Feature-specific, degradation, workaround = use history view
12	Command Center — Alerts Not Generating	P1	High-impact for ops users, but platform is up. Sprint-cycle fix

Result: Of 11 "P0s" → 2 are genuine P0, 4 are P1, 2 are P2, 2 are Stability Track, 1 needs more detail, 1 is conditional. This is the kind of breakdown engineering needs to actually function.

8. PMO Cheat Sheet – Decision Tree

A one-screen flow for PMO to classify an item in under 60 seconds.



Three questions to answer before raising:

1. **How many are affected right now?** (count of clients / fleet %)
2. **Is there a workaround?** (yes/no — and what is it)
3. **When did this start?** (timestamp, even approximate)

If you cannot answer all three → it is not a ticket yet, it is a complaint. Get the info first.

9. Metrics for This System

Tracked monthly. Tier system health metrics, separate from delivery metrics.

Metric	Target	Why
% of P0s mitigated within 24 hours	≥ 90%	P0 must mean same-day
% of P0s demoted within 4 hours	≤ 30%	Some demotion is healthy. High demotion = PMO over-escalating
P0 false-positive rate (demoted to P2/P3)	≤ 10%	If higher, rubric is failing
Stability Track items closed per quarter	Trending up	Real infra progress
Recurring P0 categories	0	Each repeat = systemic failure
VIP Overrides per month	≤ 2 active	Escape hatch stays rare
Head of Product disputes per week	≤ 2	Rubric must self-resolve

10. Rollout

Week	Action	Owner
W1	Circulate this doc. Walk PMO through §1 rubric and §8 cheat sheet	AJ
W1	Re-tier all 11 PMO items per §7	PMs + PMO
W2	First Tuesday Intake Sync uses new tier system	PMs
W2	Stability Track project created in Odoo	Eng Lead
W3	First Stability Track bi-weekly review with Engineering Leadership	Product + Eng Lead
W4	First monthly review using §9 metrics	AJ
W6	Rubric review — adjust thresholds based on first month	AJ

11. Open Items

- Capacity reserve realism** — 20% P1 + 20% Stability + 30% P2 = 70% non-feature reserve. Realistic, or do we start at 15/15/25 and grow? Need Eng Lead input.

2. **"Needs detail" status** — unclassifiable items (e.g. "API issues") need a holding tier before being scored. Add to Odoo?
 3. **Field-tech routing** — single-device issues should not hit engineering. Does a field-tech queue exist?
 4. **SLA penalty signal ownership** — who flags when a client's SLA clause is being breached: Sales, AM, or PMO?
 5. **VIP Override channel** — Teams DM to Head of Product, email, or Odoo workflow? Pick one.
-

This document replaces the simplistic "PO / not PO" model. It is a living rubric. v1.1 after first month of real use.