

REBEL SOFTWARE · RebelPSA

SLA User Guide

Understanding Global SLAs, Client SLAs, Operational Level Agreements (OLAs), and Underpinning Contracts (UCs) in RebelPSA

Intended audience	Technicians, supervisors, and administrators who work with tickets and SLAs in RebelPSA
Version	1.0
Date	August 2026

Contents

- 1.** About This Guide
- 2.** Key Terms: SLA, OLA & UC
- 3.** The Task SLA Model: Concurrent Timers on One Ticket
- 4.** The Priority Matrix: Response, Plan & Resolution Targets
- 5.** The SLA Resolution Cascade: Global SLA vs. Client SLA
- 6.** Business Hours, Holidays & Pause Conditions
- 7.** Per-Priority Toggles: What Each One Does
- 8.** Escalation Chains: What They're For
- 9.** Reading SLA Status on a Ticket
- 10.** The SLA Diagnostic Tool
- 11.** SLA Performance Reporting
- 12.** Administering SLAs: General Settings, OLAs & Imports
- 13.** Common Workflows
- 14.** Quick Reference & Glossary

1. About This Guide

This guide explains how Service Level Agreements (SLAs) work in RebelPSA: the difference between a Global SLA and a Client SLA, how SLAs, OLAs, and UCs can all run on the same ticket at the same time, what an Escalation Chain does, and what each per-priority setting controls.

RebelPSA supports several independent timers running on one ticket at once, rather than the single-timer-per-ticket approach found in many help desk tools. This gives you more precision when tracking commitments, but also means there are a few more moving parts worth understanding — this guide walks through all of them.

2. Key Terms: SLA, OLA & UC

Term	Stands For	What It Actually Is
SLA	Service Level Agreement	The commitment made to the customer -- how quickly you'll respond to, and resolve, a ticket of a given priority.
OLA	Operational Level Agreement	An internal commitment between teams inside your own organization — for example, Tier 1 agreeing to hand a ticket to Tier 2 within a set time. It has its own escalation chain and assignee pool, separate from the customer-facing SLA.
UC	Underpinning Contract	A commitment from a third-party vendor that your own SLA depends on — for example, a hardware vendor's response window on a failed part, tracked per asset. Your ability to meet a customer SLA can depend on a vendor's contract with you.

In RebelPSA, all three are full policy types in their own right, not just labels. An OLA or a UC gets its own complete policy — targets, business hours, pause conditions, and an escalation chain — exactly the way a customer-facing SLA does.

3. The Task SLA Model: Concurrent Timers on One Ticket

RebelPSA separates the reusable policy from the live, running instance of that policy on a specific ticket:

- **SLA Definition** — the reusable template: its priority matrix, business hours, pause/start/stop conditions, escalation chain, brand scope, and type (SLA, OLA, or UC).
- **Task SLA** — the actual live instance attached to one ticket, with its own clock, state, and breach tracking.

Because these are separate, a single ticket can carry several Task SLAs running side by side — for example, a customer Response SLA, a customer Resolution SLA, an internal OLA for a handoff, and a

vendor UC for a part replacement, all ticking independently on the same ticket. None of these compete with each other; they simply run in parallel.

Timer Types

- **Response** — time to first meaningful reply to the customer
- **Plan** — time to produce a resolution plan
- **Resolution** — time to fully resolve the ticket
- **OLA** — an internal handoff or commitment timer
- **UC** — a vendor's underpinning commitment
- **Custom** — any additional timer type your organization has defined

Task SLA States

State	Meaning
Active	The clock is currently running.
Paused	The clock is temporarily stopped (see pause conditions, Section 6).
Met	The target was hit — this timer closed successfully.
Breached	The target was missed.
Cancelled	The timer no longer applies -- for example, the ticket was reclassified or merged.

4. The Priority Matrix: Response, Plan & Resolution Targets

Every SLA Definition includes a priority matrix — a target time for each priority level (Critical, High, Medium, Low, or however your organization names them), for each timer type it uses. A Critical ticket might carry a 15-minute Response target and a 4-hour Resolution target, while a Low-priority ticket might carry a next-business-day Response target and a week-long Resolution target — all defined on the same SLA Definition.

Tip. A Global SLA must cover every priority level your organization uses. An incomplete priority matrix on the Global SLA is rejected, since it's the final safety net and can't have gaps.

5. The SLA Resolution Cascade: Global SLA vs. Client SLA

Every ticket needs an SLA the moment it's created — RebelPSA never leaves a ticket without one. To decide which SLA applies, RebelPSA works down a fixed list, or “cascade,” and the first level that produces a complete SLA wins:

- Ticket-specific override -- set manually or by a workflow rule
- Sub-agreement SLA -- if the specific agreement under the client explicitly overrides one
- Master Agreement SLA -- inherited automatically if the sub-agreement doesn't override it
- Tenant Template SLA reference -- if the Master Agreement inherits from a template
- Company default SLA — the client's own default
- Brand default SLA
- Help Desk Global SLA — the final fallback, always available

In everyday conversation, “Client SLA” usually means whichever of levels 2 through 5 supplied the ticket's SLA — in other words, something tied specifically to that customer's account or agreement, rather than the organization-wide default. “Global SLA” specifically means level 7: the safety net that guarantees every ticket has an SLA, even for a brand-new client with no agreement-level SLA configured yet.

Tip. *A Global SLA is a required configuration at tenant activation — at least one Global SLA covering every priority level must be published first. This is deliberate: it means the cascade can never truly run out of options.*

Important. *OLAs and UCs work differently from the main cascade — they run in addition to whichever level supplies the primary SLA, rather than competing for the same slot. A ticket can be running its Client SLA for the customer-facing Response/Resolution timers while an OLA and a UC layer on top for the internal and vendor-side commitments.*

RebelPSA never treats this as a black box: the ticket detail view shows exactly which cascade level supplied the active SLA, so nothing about the result is a mystery — see Section 10 for the tool built specifically to check this before a ticket even exists.

6. Business Hours, Holidays & Pause Conditions

Every SLA runs against a business hours calendar — its own per-day working hours, in a specific time zone — rather than a flat 24/7 clock, unless your organization deliberately configures a round-the-clock SLA.

- Holiday calendars are hierarchical: a tenant-level holiday calendar cascades down to child schedules, so you set a holiday once rather than everywhere.

- SLA clocks pause automatically on a configurable list of ticket statuses — most commonly anything in a “waiting on the customer” state — so a customer’s own delay in responding never counts against your team.

Tip. *Because pause conditions are configurable per status rather than fixed, your organization controls exactly which statuses stop the clock — this is worth checking if a ticket’s elapsed SLA time looks different from what you’d expect.*

7. Per-Priority Toggles: What Each One Does

Beyond the raw response/resolution targets, each priority level on an SLA Definition carries its own set of behavior settings — meaning your Critical tickets can behave differently from your Low-priority tickets on the very same SLA policy.

Setting	What It Does
Extend to End of Day	Rounds a calculated target forward to the end of that business day instead of leaving it at an awkward mid-shift time — this gives a little breathing room without loosening the day’s actual commitment.
Delay Start to Next Working Day	If a ticket arrives outside business hours, the SLA clock doesn’t start until the next working day begins, instead of ticking away while nobody’s on shift.
Breach Justification Prompt	Requires a technician to enter a reason before proceeding on a ticket that’s about to breach or has already breached — this builds accountability and a documented explanation into the record automatically.
Ad-Hoc Override	Allows an authorized user to manually adjust a specific ticket’s calculated target time for a one-off reason, without changing the underlying SLA Definition for everyone.
Respond-By Reset	Allows the “respond by” clock to reset under defined conditions -- for example, the customer replies, or the ticket is reassigned -- instead of continuing to count down against the original first-response window.
Multiple Response Targets	Lets a single SLA Definition define more than one response checkpoint on the same ticket — for example, an automated acknowledgment target and a separate first-human-touch target — instead of just one flat response timer.

Tip. *These settings are configured per priority level, per SLA Definition — so it’s worth checking a specific priority’s settings if a ticket’s behavior (when its clock started, whether it prompted for a reason, whether it reset) doesn’t match what you expected from a different priority level.*

8. Escalation Chains: What They're For

An Escalation Chain isn't a number you read off a dashboard — it's a configured set of rules that lives on every SLA Definition (and separately on every OLA Definition) and determines what happens automatically as a ticket gets closer to, or passes, its target.

An escalation chain is a multi-level sequence of automated actions, and there's no limit to how many levels you can define. A typical chain might look like:

- At 75% of target elapsed — notify the assigned technician.
- At 90% of target elapsed — notify the technician's supervisor.
- At breach — reassign to a senior technician, raise the ticket's priority, and notify the service desk manager.

Each level can independently change the assignee, raise the priority, and route a notification to whoever needs to know — which turns “this ticket is about to breach” from something someone has to notice into something the system handles on its own.

Tip. *OLAs carry their own separate escalation chain and assignee pool from the customer-facing SLA — an internal handoff running late escalates within your own team, without necessarily involving the customer-facing chain at all.*

9. Reading SLA Status on a Ticket

A ticket's detail view shows every active Task SLA running on it at once — each with its own live countdown, state, and breach-risk indicator. If a ticket has a customer Response timer, a customer Resolution timer, and an internal OLA all running, you'll see all three, clearly separated, rather than one blended status.

- Check the state badge (Active / Paused / Met / Breached / Cancelled) on each individual timer — a ticket can have one timer Met while another is still Active.
- A paused timer isn't stalled by accident — it means a configured pause condition (Section 6) is currently in effect, most often a waiting-on-customer status.

10. The SLA Diagnostic Tool

The SLA Diagnostic tool lets you preview exactly which SLA would apply to a ticket before one even exists — useful when you want to double-check a cascade result, or explain to a customer or colleague why a particular SLA is in effect.

- Choose a Brand Board.

- Choose a Priority.
- Optionally choose a Client (company) and a Site.
- Run the preview to see the resolved SLA and exactly which cascade level supplied it.

Tip. Leaving Client and Site blank previews what the Global SLA would resolve to for that brand and priority — a fast way to sanity-check your organization's safety net independent of any specific customer.

11. SLA Performance Reporting

The SLA Performance dashboard reports on actual historical Task SLA outcomes — met versus breached, broken down by timer type, with a trend view over time and a full breach list you can export.

- You select a company and an agreement before the dashboard loads — this keeps every report properly scoped rather than defaulting to an organization-wide view.
- A column chart breaks down Met vs. Breached by timer type (Response, Resolution, OLA, UC, etc.).
- A trend view tracks met-percentage and breach rate over time.
- The breach grid lists every individual breach with the ticket, the SLA Definition involved, and timing detail, and can be exported to CSV for deeper analysis.

12. Administering SLAs: General Settings, OLAs & Imports

A handful of SLA-related settings live outside an individual SLA Definition, at the general settings level — typically an administrator's territory, but useful for any technician to know exists:

- **SLA Hold** — controls whether tickets currently on an SLA pause are visually highlighted in the ticket list (with a configurable highlight color), and whether on-hold tickets are included when you filter for breached SLAs.
- **Out of Hours Settings** — lets your organization set a distinct priority for tickets that arrive outside business hours, optionally restricted to tickets logged by users or by email, with the option to use a separate email template for out-of-hours tickets and to notify an on-call agent specifically for out-of-hours Priority 1 tickets.
- **OLA Configuration** — OLA Templates and OLA Rules are configured separately from the main SLA Definition editor, reflecting that an OLA targets internal team commitments rather than the customer-facing policy.
- **Imports** — SLAs and their associated Priorities can be bulk-imported from a spreadsheet. Priorities depend on SLAs already existing, so SLAs must be imported first.

13. Common Workflows

Figuring Out Why a Ticket Has the SLA It Does

- Open the ticket and check which cascade level is displayed as the source of its active SLA.
- If it's unexpected, use the SLA Diagnostic tool with the same brand, priority, client, and site to reproduce the result.
- Escalate to an administrator if the result still doesn't match what you'd expect for that client's agreement.

Handling a Ticket That's About to Breach

- Check which specific timer (Response, Resolution, OLA, UC) is at risk — a ticket can have more than one running.
- If a breach justification prompt appears, enter a clear, specific reason — it becomes part of the permanent record.
- Expect the escalation chain to notify or reassign automatically once the breach threshold is crossed; you don't need to manually alert a supervisor yourself.

Requesting a One-Off Adjustment to a Ticket's Target

- Confirm the Ad-Hoc Override setting is available for that priority level.
- Apply the override on the specific ticket rather than asking for the underlying SLA Definition to be changed.
- Document the reason for the override in the ticket for anyone reviewing it later.

14. Quick Reference & Glossary

Term	Meaning
SLA Definition	The reusable policy template: priority matrix, business hours, pause conditions, escalation chain, brand scope, and type (SLA/OLA/UC).
Task SLA	The live, running instance of an SLA Definition attached to one specific ticket.
Global SLA	The organization-wide default SLA — the final, always-available fallback at the bottom of the resolution cascade.

Term	Meaning
Client SLA	Common shorthand for whichever cascade level ties an SLA specifically to that customer's account or agreement, rather than the organization-wide Global SLA.
Cascade	The fixed, ordered sequence RebelPSA checks to resolve which SLA applies to a ticket.
Priority Matrix	The set of response/plan/resolution targets defined per priority level on an SLA Definition.
Per-Priority Toggle	A behavior setting (extend to end of day, delay start, breach justification, etc.) set independently for each priority level.
Escalation Chain	The configured, multi-level sequence of automated actions (reassignment, priority increase, notification) that fires as a ticket nears or passes its SLA target.
Pause Condition	A ticket status configured to stop an SLA clock while the ticket is in that status.
Brand Scope	The set of brands an SLA Definition applies to; defaults to All Brands.