

The Escalation Authority Runbook

Run escalations without the executive rescue.

A support leader kept pushing the customer to follow the best practice his team had been given for configuring a feature. The customer's senior director of operations said that configuration does not work in their environment. The support leader, backing up his support engineer, kept repeating the best practice like a mantra.

Then the Subject Matter Expert (SME) and the senior director talked for 10 minutes. They found a solution that kept the best practice and adapted it to the customer's environment.

A week and a half of back and forth. Ten minutes to solve it, once the right two people talked.

The answer was one call away for a week and a half. Nobody owned the call.

What a bad escalation costs

Five costs, and none of them show up on the ticket:

1. Executive time.
2. Leadership time.
3. Trust.
4. The future relationship, across the whole product line.

5. Future development time.

The first cost shows up in hours. The leader in the example below spent 15 to 25 hours a week on customer issues. Multiply those hours by what that leader's time is worth to your organization. Then add what those hours were supposed to go to, the work that got deferred while they sat on calls.

Why authority collapses under pressure

Pressure does not eliminate authority. It redistributes it. When an escalation hits, decisions do not disappear. They move toward whatever is safest to defend: the most senior person in the room, or the official answer. In the story above, it was the official answer.

When it is the senior person, it looks like this. Before I took on the escalation leadership role, my leader spent 15 to 25 hours a week on customer issues. On those calls, customer and technical, the team waited for his okay before moving forward. Once he stepped off the calls, he got that time back. The team brought him in only when a decision needed his approval: budget, commitments, and roadmap changes.

Stepping in feels like leadership. What it does is rescue. Rescue solves the moment, but it can train dependency: the team learns the safe move is to wait. Every time an executive steps in and the team lets them, the system teaches itself that the right owner is whoever is senior enough to absorb the risk. The next escalation waits longer, because now there is a precedent.

A braver team does not fix this. A system does, one where deciding is the safe choice at every level, because the decision rights were granted before the pressure started. That is what this document installs.

How to use this runbook

What this document is: the procedure a team follows during a live customer escalation, so that authority, ownership, and customer trust are decided in advance instead of negotiated while the fire is burning. Five steps, plus one afternoon of preparation. None of them climbs to the executive by default.

What this document is not: a maturity model or a philosophy of leadership.

Who it is not for: teams who want a policy to point at after the fact. This is written to be used during the event.

The one line: one repeatable way to run any escalation, so the decision lands on the right owner instead of on the senior leadership, and every event leaves the system stronger.

STEP 0

Before it starts: decide the hard calls in calm weather

Every step below depends on decisions made before the first phone rings. This step takes an afternoon, once, and gets revisited after every review in step 5.

- Write the escalation categories and severity levels (step 1).
- Write the decision rights per severity level (step 2).
- Preapprove the recovery moves. Wayne McCulloch, in *The Seven Pillars of Customer Success*, describes teams that had preapproved terms ready (paused payments, approved discounts) so the renewals team could act the day a customer hit hardship, instead of taking weeks. Customers remembered who helped and who held to the contract. The same applies to escalations: decide now what the owner may offer without asking:
- Service credits, up to a written limit.

- Engineering or onsite resources, from support or a Technical Marketing Engineer (TME), coordinated with TME leadership.
- Roadmap statements, limited to what is already committed.
- Name the watch signals that tend to come before an escalation: repeat cases on the same issue across different systems, a pattern in case load, a change in customer behavior, a sponsor or champion change, a usage drop, a low satisfaction score, an unanswered request that keeps coming back. Each one gets an owner and a response before it becomes an escalation.

NEVER leave a recovery move undecided until a customer is waiting for it. An undecided move goes to whoever is most senior, by default.

STEP 1 **Name it**

The first sixty seconds of any escalation decide the next six hours. Get the category and the severity right, fast, and do not relitigate them mid event.

What to answer, in order:

- What kind of escalation is this. Pick from a fixed, short list your team already agrees on. If your list has more than six categories, it is too long to use under pressure.
- What is the severity, on a fixed scale tied to customer and business impact. Loud and severe are different variables. A loud customer on a minor issue and a quiet customer on a major one get different severities. Confuse the two and your cadence and your executive threshold both break.
- Who is naming it. One person, decided in advance by role. Whoever picks up the phone first does not get to set severity by default.

Severity levels and customer update cadence:

Level	Common label	Typical example	Support engagement	C
Sev 1 / P1	Critical / Emergency	Customer-facing platform unavailable; global network outage; data-loss/security event	24×7 active incident response; dedicated incident lead; senior engineering engagement	E
Sev 2 / P2	High / Major	Core workflow intermittently fails; one region/site is offline; severe performance degradation	Continuous work during defined coverage, business-hours or 24×7 depending on contract and impact	E
Sev 3 / P3	Medium / Normal	Reporting feature unavailable; a department has an application issue; defect with workaround	Standard queue and business-hours investigation	D
Sev 4 / P4	Low / Informational / Request	Configuration question; documentation issue; cosmetic interface defect	Business-hours, best-effort or scheduled handling	V a n

The rule: severity is decided once, at the start, by the named role. If new information changes the picture, the same role revises it on purpose, out loud. Severity does not drift upward because someone in the room is anxious.

What earns an executive automatically: define this in writing before you need it. A short list, tied to severity level and customer tier. The seniority of the person complaining is not on the list. If the list does not exist yet, the default owner is whoever owns the account. Examples that earn an executive automatically: legal requirements, accounting and finance, and the company's reputation.

NEVER re-argue severity in the middle of the event, or let it be set by volume.

STEP 2 The right owner decides, now

"No one owns the customer, but someone always owns the moment." Scott Huggins, Walt Disney World, quoted in The Seven Pillars of Customer Success

This section stops the leak. Everything else in this document supports one rule: the decision sits at a set level, with permission already granted, before the escalation happens.

Define, in advance, for every severity level:

- Who owns the decision. By role, so the runbook survives turnover.
- What that role is authorized to decide alone. Refunds up to what amount. Workarounds of what kind. Commitments of what nature. Write the number. Vague authority is no authority.
- What that role is required to escalate, to whom, and how fast.
- What an executive is asked to do when they are pulled in. This is the line that matters most. The executive backs the team and tells the customer who on the team is in charge. Taking over the work and becoming the incident responder undoes both.

An executive on the call can support: unblock a resource, make a call the owner is not authorized to make, absorb a relationship conversation the owner should not be having yet. An executive on the call cannot take over a decision the owner already had the right to make, run the call, or become the single point of contact from then on. The moment the executive starts making the owner's calls, the system has quietly retrained itself, and it will happen again next time.

The sentence. Every team needs a version of this that someone can say out loud, to a person more senior than them, safely: "I have this. I will flag you if I need you." The executive needs the matching sentence: "You have this. Tell me what you need." If neither sentence exists in your culture yet, that is the gap, and no cadence document fixes it by itself. I made these two sentences a rule in my own escalation team. The directors adopted them and carried them to the rest of the organization.

When nobody has decided in 10 minutes: name a default. Either the decision escalates automatically to the next role up, on a clock, or the most senior person present makes a visible, temporary call and hands it back the moment the right owner is available. Silence is the only wrong answer. A wrong decision, named as temporary, is recoverable. A decision nobody made is harder to recover from.

If the customer's champion or sponsor changes mid-escalation: that is a new risk, so name it (step 1) and confirm who owns the relationship with the new person before the next customer update.

NEVER let an executive make a decision the owner was already authorized to make.

STEP 3 **One voice, on a clock**

No cadence, no trust. The customer hears from one person, on a promised schedule. Six voices and silence both break it.

The customer sets the external clock. The team sets the internal one. Keep a cadence per severity, and when the customer asks for a different one, follow theirs.

- First contact with the customer within 20 minutes of the escalation being named.
- The first substantive update: what is known, what is being done, when the next update arrives. Never promise a fix time you are not certain of.

- The interval after that: Sev 1 every hour, Sev 2 every 2 hours, Sev 3 daily or weekly (the table in Step 1), unless the customer sets a different one. Hold it until resolution, and write down the conditions that allow it to relax.
- One named person for the customer to hear from. If ownership transfers internally, the customer's point of contact does not change without being told, in the same message that introduces the handoff.
- What gets said when there is nothing new: that there is nothing new, and when the next update comes. A customer who hears nothing assumes the worst. A customer who hears "still working, this is what we are trying, next update at 3pm" does not.

Internal cadence runs on the same discipline, at a different interval. During an active escalation, executive and stakeholder updates go out every hour for Sev 1 and every 2 hours for Sev 2, or match the customer's cadence if they set one. Use a fixed short format: what happened, what is being done, what the customer has been told, when the next update comes. This stops leadership from interrupting the work to ask for status. When the update arrives on schedule, they stop pulling it out of the team.

NEVER miss a promised update, even when there is nothing new to say.

STEP 4 **Close it twice**

Fix the issue, and repair the trust. A resolved ticket with a quiet customer is still open.

Technical resolution closes the issue. Trust repair closes the escalation. These are two different jobs, done by two different motions.

Technical close: the fix is verified, the customer is told it is verified, and any lingering risk or data impact is stated plainly. Never guess. Never estimate a

resolution time you do not have.

Trust close, within 5 business days of technical resolution:

- A named person calls or writes the customer, separate from the "it's fixed" message.
- What is said about what failed: plain, specific, no more than the customer needs to trust the account is being handled, no defensive language.
- What is not said: blame of an individual, speculation, anything that has not been confirmed.
- What changes from here. Customers stay for the future they believe in more than for the past you delivered (Wayne McCulloch makes this case in *The Seven Pillars of Customer Success*), so the trust close tells them what will be different, and step 5 delivers it.
- What the customer receives in writing, if anything, and who owns sending it.
- The test of a trust close: does the account team believe the relationship is stronger, the same, or weaker than before the escalation. If the answer is weaker, the trust close has not happened yet, whatever the ticket says.
- For high-severity, high-impact cases, hold a follow-up meeting with the customer.

NEVER let the "it's fixed" message stand in for the trust close.

STEP 5 Bank one lesson

One owned fix, with a name and a date, so the same escalation does not come back. One thing: never a list.

We inspect the system before we blame the person. The review exists to change the system.

The review, briefly:

- Held within 2 business days of resolution, while the details are fresh.
- The people in the room: everyone who touched the decision. Everyone who touched the incident is a bigger room than you need.
- The question that matters more than any other: where did authority wait for someone senior when it did not need to, and why did that feel like the safe choice.
- One change, owned by one named person, due by one named date. Five action items produce zero changes. One produces one.
- Tell the customer when the change lands. That completes the promise made in the trust close.
- Tag the review with where in the customer journey the escalation started: onboarding, launch, adoption, renewal, expansion. After several reviews the tags show you where escalations cluster, and that is where step 0 needs the most work.
- What the review never produces: a list of who did what wrong, in writing, attached to a name.

NEVER leave the review without one owner and one date.

What you are probably using today, and what it misses

- **The incident process.** It measures time to resolve. It does not record where the decision sat, or who waited for whom.
- **A stronger escalation manager.** It works until that person is on vacation, promoted, or gone.
- **The executive who jumps in.** It solves the moment and trains the team to wait for the next one.

This runbook adds the layer none of them cover: who owns the decision, decided before the pressure starts.

Why this works

People wait because acting is risky to them personally. This system makes the right call the safe call, so waiting stops being the smart move. Authority flows to the right owner under pressure, instead of climbing to whoever is senior enough to absorb the blame.

AI can support the work described in every step above. It cannot own the consequence. Whatever tooling you route these steps through, the decision rights, the sentence in step two, and the trust call in step four still belong to a named person.

If you are the executive who wants your time back and a team that owns the next escalation without you, this is how you get both.

Where to go from here

Step 0 is the afternoon teams skip, and every number in this runbook belongs to a specific team. This document is the open version of the operating loop, and you can start running it today. The full version, customized to your organization with a severity taxonomy, a decision rights matrix, and a pressure simulation, is the Authority Under Pressure OS engagement.

To see where authority is leaking in your team today, take the Authority Leak Scorecard.

Take the free Scorecard →