



SAMPLE CLIENT DELIVERABLE · OPERATIONAL READINESS

Building a Change Intake System

Jessica Sumner

Sumner Solutions



THE PROBLEM

Change reaches the customers before change reaches the front line

Ops needs a change management process that is:

- repeatable, scalable, sustainable
- leverages automation
- lightweight enough to be workable, heavy enough to do the job
- drives customer value

A CHANGE impacts anything a customer touches - agent conversations (both human & AI), KB content, email, SMS



ASSUMPTIONS

Working assumptions

Set so the design could move forward; each is easy to change per client.

PEOPLE

Most agents work for an outsourced partner: train-the-trainer, about two weeks' notice.

A lean readiness team: two people

No dedicated CX engineering resource: delivery leans on Engineering, Product, Compliance.

Dependencies on cross-functional stakeholders to buy-into process and artifacts

TOOLING

Jira is the system of record — Product and Engineering already work there.

A knowledge base exists that agents open during a live call.

An AI skill-building capability exists at the corporate level.

PROCESS & SCOPE

Tier timetables of 30 / 10 / 3 days are a stake in the ground, not a finding.

Regulated product: consumer-protection rules (e.g., Reg E / Reg Z) govern disputes.

No readiness metrics exist today. Contact-reason data does.



THE FIX

Change: from capture, to launch, to validate

File it

Change owner files.
Slack form & email + AI listener:

- What's changing?
- Who does it affect?
- Customer visible?
- Go-live date?
- Change type?
- Change owner?
- Impact to agent?
- Links to docs?

Size it

Auto-create Jira ticket from the form answers.
Jira because already work there

- Tier
- Priority
- Due date
- Artifacts (KB article, macro, QA scorecard, etc.)
- Training requirements

Build it

Only what the tier earns — AI drafts, human approve.

- Agent update, plain words
- KB article and macros
- Escalation path named
- Training where errors cost
- QA scorecard updated

Launch it

The readiness gate

- Operational Readiness gives green-light to go live
- Sign-off by stakeholders
- Trainers scheduled, floor briefed

Measure it

Set cadence after go-live & baseline

- Contacts on the topic
- AHT
- Escalation Rate
- QA fails on that topic
- Questions asked twice

Learn from it

On set cadence, across every change

- Themes by team and change type
- What the misses have in common
- Rules and templates updated
- Next one costs less than the last




LEARN FROM IT

Iterate until steady state

Agent has answer BEFORE customer contact.

Learn from it closes the loop. What we find in one change rewrites the next.



HOW THE TIER IS ASSIGNED

RICE: The tier is calculated, not guessed

Four scores, auto-calculated from the intake form

Reach

Who does it affect?

- 1 One partner, one queue
- 3 Several partners, one channel
- 5 All partners, all channels

Impact

Customer visible? Change type?

- 1 Cosmetic, no consequence
- 3 Wrong answer, recoverable
- 5 Money, credit, dispute rights

Confidence

Spec linked? Date firm? Owner named?

- 1 A rumour in a thread
- 3 Scoped, date still soft
- 5 Written, dated, owned

Ease

Impact to agent?

- 1 Nothing about it is intuitive
- 3 Some new steps to learn
- 5 Nothing new to learn

Score = (Reach × Impact × Confidence) ÷ Ease

TIER 1 · BIG

25 and above

TIER 2 · MIDDLE

6 to 24

TIER 3 · SMALL

under 6

Ease inverts: what nobody can handle divided by 1, and scores high.

What beats the score — straight to Tier 1

- A regulator or compliance rule is behind it
- New client partner, new program, or new channel
- Irreversible, or it touches customer money

Anyone can push a tier up. Only we push one down — in the ticket, with a reason.



HOW WE SIZE IT

Matching work to impact

TIER 1

30 days notice

New partner client, new channel, anything with a regulator behind it.
Requires: Full training, QA scorecard before launch, someone signs off.

TIER 2

10 days notice

Changes what an agent says or does.
Requires: An update, a macro, a fifteen-minute huddle.

TIER 3

3 days notice

Good to know, nothing to do differently.
Requires: One line in Friday's note. No meeting.


RICE decides the size, not volume: decides what to work on in what order.

Tier 0 - Exceptions. Expedited changes happen and need to be accounted for. Nothing gets blocked, the trade-off is named and recorded to track impact. After 90 days we can show what expedited changes cost in terms of escalations & repeat contacts.



HOW PEOPLE SEE IT

Single roadmap: What's coming, who owns it, and whether it's on track.



Bars are calculated, not drawn: Big needs 30 days' notice, Middle 10 business days, Small 3. Notch = materials lock, 21 days before go-live.

What's the value? One view of every change impacting front-line agents, human and AI, so capacity, sequencing, and business impact are tracked simultaneously.

The tools file it, chase it, and write the first draft.

Two paths, running in parallel. Automation moves the tickets; Ops Readiness makes the calls.

What we hand off to automation:

Filing and sizing — from the moment a change appears.

- 1 **Captured**
Slack form, email, or the listener catches a channel post
- 2 **Ticket opens**
Form answers map straight into Jira — nobody retypes anything
- 3 **Scored and tiered**
Reach, Impact, Confidence, Ease — from those same answers
- 4 **Dated backwards**
Go-live minus 30 / 10 / 3 days; lock at minus 21
- 5 **Work spawned**
Artifact checklist and training scope by tier; owner assigned
- 6 **Chased and measured**
Nudges at 5, 2, 0 days; metrics at day 7 and 14

What humans keep:

Judgment, wording, and the relationships behind the process.

- 1 **Triage the flag**
Really a change? Confidence under 3 requires review
- 2 **Set the tier**
Anyone pushes one up. Only we push one down, with a reason
- 3 **Edit the drafts**
With CX Enablement: review all artifacts for veracity and completeness
- 4 **Clear regulated language**
Anything a regulator or compliance rule governs
- 5 **Call it ready**
The go-live decision, with a name on it
- 6 **Work the room**
Cross-functional buy-in, and the team that skips the form

WORKED EXAMPLE: THREE CHANGES IN FLIGHT

Same door. Very different work.

RICE scorecard: the numbers behind each tier, and the sanity check.

NEW PRODUCT LAUNCH · 45 DAYS

Tier 1

R	I	C	E
5	5	4	1

$$(5 \times 5 \times 4) \div 1 = \mathbf{100}$$

Materials final three weeks out. The outsourced trainers need time to deliver the plan we've developed. No training of front-line agents = no go-live decision.

LIVE CHAT · ALL PARTNERS

Tier 1

R	I	C	E
5	4	3	1

$$(5 \times 4 \times 3) \div 1 = \mathbf{60}$$

Needs a chat QA scorecard, new artifacts, training and baseline metrics. Start with one partner and a small group. Full rollout after the product launch settles.

DISPUTE POLICY

Tier 2

R	I	C	E
2	3	5	2

$$(2 \times 3 \times 5) \div 2 = \mathbf{15}$$

Internal policy, not a legal deadline. QA scorecard updated, and agent notification.

Sanity check: a new program and a new channel force Tier 1 whatever the score says. The dispute policy is internal, so nothing overrides it and the score stands on its own.

Stakeholder pain-point identification
Create framework that mitigates the pain
Launch v1 process
Measure & Learn

What does success look like? Whether an agent taking a call on day 91 has the answer in front of them and got it before the customer called, not after.

DAY 91

What does success look like?

A password reset change went live, with no warning to CX: before and after

Day 1: current state

Password reset feature updated. What actually happened.

- 1 Eng ticket opens**
Nothing asks whether a customer or an agent will notice
- 2 Change ships**
On the release schedule, not a readiness schedule
- 3 Written up in Slack**
Clear, detailed, and written for an engineering audience
- 4 No one owns the handoff**
Nobody in Customer Ops is tagged or accountable for it
- 5 Customers hit it first**
Resets behave differently that morning; no one is warned
- 6 Agents learn from the caller**
Handle time up, escalations up, scored as agent error
- 7 Downstream, in the numbers**
Increase in refunds, decrease in online star ratings

Day 91: to-be state

The same update, implemented through the framework.

- 1 Flagged at the source**
The release template asks: does this change what an agent says?
- 2 Readiness ticket opens itself**
At "ready for release" — the listener catches the Slack post too
- 3 Scored and tiered**
Tier 2: ten business days, prep date set, owner named
- 4 Drafts inside a day**
Agent update, KB article, macro, check questions — we edit
- 5 Pushed, not posted**
Pinned in the channel and sent to the BPO team leads
- 6 Ready before live**
Same date for the release and the floor; verified at day 14
- 7 Downstream, in the numbers**
Decrease in refunds, increase in online star ratings



APPENDIX

Sample password reset BPO agent comms

Password reset — what changed · 2 minute read

The short version

People now reset their password with a link we email, instead of a code we text. The link is good for 15 minutes.

Script

"I'll get that started for you. You'll get an email with a reset link in the next minute — it's good for 15 minutes. Want to stay on with me while you look for it?"

Questions & Answers

"I didn't get the email." Spam first, then check the address on the account is the one they're looking at.

"It says the link expired." Send another and stay with them while they open it.

"Can't you just give me a code?" We can't. Be clear: "we moved to email links because they're harder for someone else to grab." No need to apologize.

What you can't do

You can't set it for them or extend the link. You can resend as many times as they need.

Short version first, because that's all anyone reads mid-shift.

Words to say out loud, not a description of the feature.

The issues - that's what the call actually is.

What you can't do; false promises are what create the escalation.

Illustrative example: details are representative, not drawn from any single client.

