

The Coffee Break Test

25 June, 2026 By **Mark Cane**

Should Salesforce Duplicate Management Take Longer Than a Coffee Break?

Every Salesforce administrator has a rhythm to their day.

There are notifications to check, dashboards to review, users to support, automations to monitor and integrations to keep track of. Somewhere amongst all that activity sits another familiar task: Salesforce duplicate management. It rarely demands immediate attention, but it has a habit of never quite going away.

A duplicate report needs attention. A queue has grown since the last review. A set of records needs comparing. A few merges need completing. It is easy to imagine an admin opening Salesforce, picking up a coffee and thinking, “I’ll just clear a few duplicates.”

That led us to a simple question. Should routine Salesforce duplicate management take longer than a coffee break?

Introducing the Coffee Break Test

The Coffee Break Test is a simple way to think about whether a duplicate management process is really working. If routine duplicate review consumes hours each week, creates a growing backlog or feels like another recurring administrative task, the process may be asking humans to make too many decisions.

If however, the routine work has already been handled and the admin is only reviewing a small number of genuine exceptions, duplicate review can naturally fit into a short, focused window of time.

The objective is not to rush important decisions. The objective is to remove unnecessary ones.

A duplicate management process passes the Coffee Break Test when most of the work has already been resolved before a person needs to get involved.

Why Queues Feel Productive

Humans like clearing queues because they provide visible progress. Inbox Zero feels satisfying. Unread notifications create a small itch. Open tasks invite

attention. Duplicate reports trigger the same feeling and every record removed from the queue feels like an achievement, even if the decision queue will grow again tomorrow.

When five hundred duplicate records become four hundred, it feels like order is being restored. That progress is genuine, but it can also create a false sense of completion.

Duplicate queues have a habit of returning. New records are created, data is imported, users enter information differently and integrations introduce variation.

Before long, the queue begins to grow again.

The administrator hasn't necessarily solved the problem. They've successfully processed another batch of duplicate records.

That's an important distinction; processing work isn't the same as eliminating work.

The real opportunity isn't to help administrators review duplicate records more efficiently. It's to reduce the number of duplicate decisions that ever require human attention.

That is why the Coffee Break Test matters. It shifts the question from "How quickly can we clear the queue?" to "Why is so much work reaching the queue in the first place?"

From Duplicate Management to Duplicate Resolution

This distinction is central to how we think about the category. Salesforce already does an excellent job of finding duplicate records through Matching Rules and

Duplicate Rules. Identifying duplicate records is no longer the hardest part of the problem.

The work begins afterwards. Someone still has to decide whether the records should be merged. Someone has to determine which values should survive. Someone has to process the merge and make sure the outcome is consistent.

That is duplicate management. We think the next evolution is duplicate resolution: Salesforce finds duplicates. OneMerge resolves them.

It is a subtle change in language, but a significant change in philosophy. The goal is not to help administrators manage more duplicates. The goal is to reduce the amount of duplicate management they have to do.

The Operating Model: Setup, Automate, Review

The organisations that pass the Coffee Break Test do not simply work faster. They work differently.

First, they set up the rules. They define how duplicate records should be resolved, which values should survive and what governance requirements need to be respected. Then they automate resolution. Routine duplicate decisions are applied consistently through merge automation, rather than being remade manually every time a similar scenario appears. Finally, they review exceptions only. Human judgement is reserved for the small number of cases where context, sensitivity or business knowledge genuinely matters.

This operating model is intentionally simple.

- Setup the rules.
- Automate resolution.
- Review exceptions only.

That is how duplicate management becomes duplicate resolution.

A Better KPI: Human Review Rate

Traditional duplicate management metrics often focus on activity.

- How many duplicates were identified?
- How many records were reviewed?
- How many merges were completed?

Those numbers are useful, but they mostly tell us how busy the process was. They do not tell us how efficient it was.

A better KPI is the **Human Review Rate**. In other words, how many duplicate decisions actually required human involvement?

Imagine a month where 1,000 duplicates were identified. If 790 were resolved automatically, 180 were resolved by an AI agent and only 30 required human review, the most interesting number is not 1,000. It is 30.

That means only 3% of duplicate decisions needed a person. That is the difference between a duplicate management process and a duplicate resolution process. One creates work. The other removes it.

Visibility Creates Trust

Automation only works when people trust it. That is why visibility matters as much as automation.

Admins need to understand what happened, which records were resolved, why a decision was made and whether anything requires follow-up. Without that visibility, automation can feel like a black box.

The OneMerge feature “Merge Activity” is designed to provide that confidence. It gives admins a clear view of duplicate resolution activity, including automated merges, AI agent-assisted decisions, human-reviewed exceptions, merge history and rollback information.

That visibility creates trust. Trust allows teams to automate more routine decisions and more automation means fewer duplicate decisions reach a human in the first place.

This is where the Coffee Break Test becomes more than a usability idea. It becomes a measure of confidence. If an admin can quickly review what happened, understand the exceptions and move on with peace of mind, the process is working.

The Future: Automated, Agent-Assisted and Human Reviewed

The future of duplicate resolution is not a single workflow where every duplicate record follows the same path. It is a layered model. Obvious merge decisions should be resolved automatically using configured rules and automation. More complex but still routine scenarios can be reviewed and resolved by trusted agents working within customer-defined instructions and guardrails. Only the remaining exceptions should go to humans.

That gives us three levels of duplicate resolution.

- Automated resolution.
- AI Agent-assisted resolution.
- Human review.

This model does not remove control. It moves control to the right place. Humans define the rules, monitor the outcomes and review the exceptions. Automation and agents handle the routine work.

The UX Pattern: High-Volume Decisions with Low Friction

As the number of human decisions becomes smaller, the quality of the review experience becomes more important.

When a person does need to review a duplicate exception, the experience should be simple, contextual and fast.

The Tinder Pattern

One UX concept we are exploring for a future OneMerge release is inspired by what we sometimes describe as the Tinder Pattern. Not because duplicate management should resemble a dating app, but because Tinder is one of the best-known examples of an interface designed for high-volume, low-friction decision making. It presents clear context. It reduces the decision to a simple action. It moves naturally to the next item. It creates momentum without requiring the user to navigate complex screens. Those same usability principles are relevant to duplicate exception review.

If an admin needs to decide whether records should be merged, the interface should help them understand the situation quickly, make the decision confidently and move on. The Tinder Pattern is forward-looking for OneMerge, but the principle behind it is already central to the product: if human judgement is required, the experience should be effortless.

Passing the Coffee Break Test

Perhaps the future of Salesforce duplicate management is not about becoming faster at reviewing duplicate records. Perhaps it is about reaching the point where reviewing duplicates is no longer a significant part of the day.

Imagine starting your morning knowing that the routine cases have already been resolved, AI agent-assisted decisions have reduced the workload further and only genuine exceptions require your attention. Instead of working through hundreds of records, you're simply reviewing the handful of decisions that genuinely benefit from human judgement.

You can see exactly what has happened via the Merge Activity tab, understand why each decision was made, undo a merge if necessary and move on with confidence. Duplicate management no longer feels like an endless queue waiting to be cleared. It becomes a well-governed process that quietly takes care of itself while keeping you informed.

That is what the Coffee Break Test is really about.

It isn't about trying to review duplicates while drinking a coffee. It's about building a duplicate management process that is efficient enough, trusted enough and intuitive enough that routine review naturally fits within the time it takes to enjoy one.

Does your Salesforce duplicate management process pass the Coffee Break Test?

We'd love to hear how your team measures duplicate management success.