Shared data and better decisions

Cross-environment master data

Emma: Welcome to Central Perspectives, a Business Central podcast from ITrators. I'm Emma, here with Alex. We're AI hosts, and our scripts and voices are AI-generated. Today we're looking at a few changes that deserve a closer look before you decide what to do with them: sharing master data across environments, adding another person to an expense review, and making a sandbox a more useful place to test performance. Those three are previews. We'll also revisit an inventory report whose date setting is easy to misread, and finish with a small improvement to payment journals. Different tasks, but a useful question running through them: what does the feature actually let you conclude? Now, let's get started.

What's New or Changed in Business Central 2026 release wave 2 - Update 29.1

Emma: Let's start with master data. Business Central 29.1 lists cross-environment master data management as a public preview. The important expansion is between environments in the same Microsoft Entra tenant, and it's for Business Central online. That means we're not talking about combining unrelated organizations' tenants. Alex, take a company group with two subsidiaries that both sell the same items. They have separate environments, and someone keeps entering the same product details twice. What's the useful change here?

Set up companies to synchronize master data

What's New or Changed in Business Central 2026 release wave 2 - Update 29.1

Alex: They can now evaluate the standard synchronization route across that environment boundary, rather than assuming separate environments rule it out. Think of one company as the place that owns the shared item details. The subsidiaries pull the information they need from it. This isn't a merger of the companies, and it doesn't mean their sales orders or stock movements become shared. We're dealing with master data: the information used to describe and set up the entities the business works with. In your example, that could reduce the repeated entry of item information. But I'd make the first decision before touching the connection setup: which details really ought to be identical? A common product description is one thing. A field that expresses a local business decision may be another. Sharing everything simply because the connection can reach it would miss the point.

Set up companies to synchronize master data

Emma: I'd be tempted to start with the fields people are already copying into emails. There's a fairly good clue that they want the same answer. But the email also gives the receiving person a chance to say, 'That isn't the item we call that here.' If we remove that conversation, we need to be careful about what replaces it. Suppose both subsidiaries already have item records. Does establishing the connection settle which records represent the same thing, or is that a separate piece of work?

Alex: It's separate, and your email example gets to the important part. A connection answers whether the systems can communicate. Coupling answers which record belongs with which. If the companies already have data, the matching criteria deserve a deliberate review before initial synchronization. I'd choose a small set of representative items: one with an obvious match, one with different numbering, and one that exists only in the source. Then establish the intended result for each before running anything. That's a proposed check, not a claim that a matching rule will magically recognize every local convention. The setup also lets you choose tables and fields. So the goal isn't to automate the email unchanged. It's to make the agreement behind the email explicit: this information comes from here, these records correspond, and these local details stay local. Once that agreement is clear, repeated updates become a much more sensible thing to automate.

Set up companies to synchronize master data

Emma: That changes my starting point a little. The repeated email is evidence of effort, but not necessarily evidence of an agreed source. Two people could each believe they're sending the correction. The documentation describes a one-way pull from the source to the subsidiaries, rather than a collaborative editing arrangement. So if the receiving team finds a mistake, I wouldn't want their local correction to disappear into a tug of war. Where should that conversation go?

Set up companies to synchronize master data

Alex: Back to whoever owns that field at the source. Not because the source team is always right, but because the process needs one place where a correction becomes the shared answer. I'd give the receiving team a clear way to raise it and agree who checks the result after the change. Otherwise you've saved data entry and created a different kind of chasing. There are two reasonable choices here. A genuinely shared field can have central ownership, with local teams contributing corrections. A field whose meaning differs locally may be better left outside synchronization. Neither choice is a failure of the feature. The awkward choice is pretending a field is shared while every subsidiary expects to make independent edits. The practical benefit is strongest when the business already wants one maintained answer and is just tired of distributing it manually.

Emma: And this isn't only a task for the person maintaining item cards. The cross-environment setup includes an Entra application and an administrator's involvement. There's also a documentation wrinkle worth making explicit: the page explains the cross-environment route, but still contains an older requirement saying all companies must be in the same environment. We're relying on the explicit new section and the preview release listing for the narrow scope we've described. We shouldn't turn that into a promise about every region or every tenant's rollout. What would you want agreed before the administrator spends time connecting it?

Set up companies to synchronize master data

Alex: A named source company, a deliberately small first scope, the expected record matches, and someone responsible for operating the connection afterward. That last person matters. The setup uses credentials with an expiry, so renewal belongs to the operating plan, not to whoever happens to remember the original project. I'd also agree what a useful result looks like. For example: a reviewed change to a shared item description reaches the intended subsidiary record, while a local field remains as intended. Then check how the team notices and handles a failed synchronization. Those are better acceptance questions than 'did the setup wizard finish?' For a group in one tenant with clear ownership, this preview is a concrete standard option to investigate. For companies in different tenants, or a requirement for two-way collaborative editing, the announcement doesn't establish that fit. We can keep the opportunity without stretching it.

Set up companies to synchronize master data

Emma: Before we leave that, I want to separate two milestones that a project plan might lump together. There's the initial job of bringing existing records into the arrangement, and then there's keeping them up to date afterward. If the subsidiary is brand new, copying records into an empty table sounds fairly straightforward. If it's been operating for years, the same 'start synchronization' instruction sounds much less reassuring. Would you use different acceptance checks for those two situations? I'm thinking of the person who opens the subsidiary the next morning and needs to recognize the records they've been using, not just see a successful status.

Alex: I would, because the risks are different. With an empty destination, I'd check whether the intended records and fields were created and whether the team can use them as planned. With existing data, I'd pay more attention to record identity, matching and which values the initial operation changes. A full synchronization and matching existing records are different choices in the setup; selecting the convenient-looking option isn't a substitute for deciding which outcome you want. I'd also separate initial completion from the ongoing update check. A successful first load doesn't tell us that the next source change will be noticed and handled correctly. Likewise, seeing job queue entries after enabling synchronization doesn't mean uncoupled records have already been synchronized. So the handover should include a small, understandable example of both: these existing records were matched as intended, and this later source change reached the right place. The receiving team can inspect that result without understanding every connection setting. That's a much stronger handover than a screenshot of a green job, and it keeps the people who know the data involved in deciding whether the setup is right.

Set up companies to synchronize master data

Interim expense approval

Emma: There's another preview in 29.1 that changes a question we left open earlier. In our expense discussion, we looked at an employee's travel claim that also needed a project owner's decision. At the time, the documented lack of interim or conditional approvals was a reason not to assume the agent would fit. Microsoft now documents one optional interim approver. That isn't the same as saying every conditional approval route is supported, but it is a meaningful new step. Alex, does this solve the example we were concerned about?

What's New or Changed in Business Central 2026 release wave 2 - Update 29.1

Review and Approve Expense Reports

Alex: It gives us a specific route to assess instead of ruling the agent out on that one limitation. With Expense Agent enabled, an interim approver can be assigned to a submitted report that's pending approval. That person reviews it before the final approver. The interim person has to be different from both the person claiming the expense and the final approver. So in the example, the project owner could be the additional reviewer and the manager could remain the final approver, provided their setup and responsibilities fit. But there are two separate needs hidden in 'project approval.' One is having a second person review the report. The other is automatically choosing that person because a particular project or amount appears on it. The documentation establishes the first route. We shouldn't quietly add the second promise.

Review and Approve Expense Reports

Emma: Then the useful question isn't just 'can we have another signature?' It's whether the right person receives the right decision. In our example, the project owner knows whether the trip belongs to that project; the manager knows whether the employee should be reimbursed. I'd rather keep those two questions visible than give both people an identical instruction to approve everything. And there's a practical restriction in the current documentation: the complete interim flow isn't supported in the web and mobile apps. Both approval stages need to happen in Business Central. That's a different working day for someone who expected to handle this on their phone.

Review and Approve Expense Reports

Alex: And that restriction could decide whether the improvement is useful now. If both reviewers already work in Business Central, the extra step may fit quite comfortably. If the project owner rarely opens it, the process needs more thought than adding a name to a field. The approver needs the appropriate expense-user setup and a linked Business Central identity; matching the email is part of that setup. I'd walk through a complete example with the actual people who would perform those roles, including how they find the report and what information they use to decide. A workflow can be technically valid and still leave a claim waiting because each reviewer assumes the other person has it. Your distinction between project purpose and reimbursement helps here: it gives each stage a reason to exist, and a clear point at which its responsibility ends.

Review and Approve Expense Reports

Emma: I'd include a rejected report in that walkthrough, not just the tidy one that passes. Say the project owner spots a cost that belongs somewhere else. The employee corrects it and sends the report back. It would be easy to celebrate the happy path while leaving everybody uncertain about whether the earlier decision still counts. And for the employee, the uncertainty is very concrete: do I fix the report, or start emailing the two approvers to explain what changed?

Alex: That's the better test. The documented resubmission route sends an eligible interim review through again, rather than treating the earlier interim decision as permanent. I'd check that the people reviewing the corrected report can understand the correction and know what they're deciding now. The distinction in statuses matters as well: interim approval isn't final approval and doesn't make the report ready to post. Even final report approval remains separate from other posting requirements. Where VAT reclaim review applies, an accountant still has that decision. So the extra reviewer isn't a universal stamp that certifies the expense, the accounting, and the tax treatment all at once. It adds one meaningful review stage. Keeping its purpose narrow makes it easier to explain to the employee and less likely that everyone assumes somebody else has checked the important detail.

Review and Approve Expense Reports

Emma: I'm glad you narrowed that. 'More approval' can sound automatically safer, when two people checking the same thing may just produce a longer queue. I'd use this where the second person knows something the final reviewer doesn't. For a straightforward claim that doesn't need that distinction, the extra step would be hard to justify. And if a company needs several project owners, amount-dependent routing, or a fully mobile process, we still have an assessment to do rather than an announcement that settles it.

Review and Approve Expense Reports

Alex: Yes, and an unchanged process can be the right result of that assessment. A business whose present approvals work doesn't need to add another person just because the option exists. The new opportunity is more specific: if the missing second review was the obstacle, revisit that obstacle with a real report and the two intended reviewers. Check the ordinary submission, a rejection, and a corrected resubmission. Agree who assigns the interim reviewer and how an absent reviewer is handled in the actual configuration; don't assume the new field answers those operating questions. Microsoft calls this a production-ready preview, which still means preview terms and possible change. It's enough to justify a focused evaluation, not a blanket conclusion that the whole approval requirement is solved.

What's New or Changed in Business Central 2026 release wave 2 - Update 29.1

Review and Approve Expense Reports

A controlled sandbox test window

Emma: Let's move to a preview that could make those evaluations more useful, although it addresses a different kind of question. Administrators can temporarily match a sandbox database configuration to a typical production environment. Microsoft lists this in 29.1, with performance testing and cloud-migration data upgrades among the scenarios. The word 'typical' is doing some work there. If a developer says, 'It was slow in the sandbox, so it will be slow when we go live,' does this new option make that conclusion more trustworthy?

What's New or Changed in Business Central 2026 release wave 2 - Update 29.1

Sandbox environment resource configuration (included section)

Alex: It can make the test environment more representative, but it doesn't turn that sentence into a guarantee. A sandbox normally has a different database configuration. The new action temporarily changes that part of the test conditions toward a typical production setup. That helps when the database resources are relevant to what you're measuring. It doesn't recreate your exact production workload, every integration, every user's connection, or the pattern of people working at the same time. I'd ask the developer what was slow. Opening one page? Posting a batch? Running a data upgrade? Those aren't interchangeable observations. Once we've named the task, we can use this option to reduce one important difference between the test and the intended operation, while keeping the other differences visible.

Sandbox environment resource configuration (included section)

Emma: Let's name one, then. A company wants to know whether its evening batch can finish before the next morning's work. I wouldn't be satisfied with a screenshot showing that one small sample ran faster. The people planning the change need to know whether the whole job fits into the available window. And they need time to investigate if it doesn't. What would you prepare before switching the sandbox configuration?

Alex: I'd prepare the workload and the observation first. For that evening batch, decide which records represent a normal run and which represent a demanding but plausible run. Write down what counts as finished, including any follow-on work that actually belongs inside the window. Then make sure the people investigating can see enough information to explain a slow run, not merely time it. If we change the data volume, the code, and the environment configuration all together, a better result won't tell us which change helped. So I'd keep a small comparison plan: same task, known data, recorded configuration, and a clear description of anything else that changed. That doesn't require turning the exercise into a laboratory. It just makes the result useful when someone asks whether a particular code change helped, or whether we only gave the same work more suitable resources.

Emma: So don't start the clock while somebody is still trying to find the test data. That matters because this isn't an unlimited setting. The matching period is fixed at seventy-two hours, with three matches per tenant per calendar month. The configuration change also restarts the environment. I can imagine two teams independently deciding that Tuesday is a good day for their own test, without realizing they're sharing that allowance or interrupting each other's work.

Sandbox environment resource configuration (included section)

Alex: That's why I'd treat it as a booked test window, not a personal speed switch. One person coordinates the environment change, both teams know the restart time, and the experiments are ready before the window begins. They might share the period if their work doesn't interfere, but overlapping workloads would need to be intentional. Otherwise one team could be investigating a slowdown caused by the other's test. For the overnight batch, I'd reserve time for a baseline, the proposed change, and a repeat if the result needs explaining. I wouldn't spend the entire period on one long run that ends just as everyone needs to inspect it. The feature gives you a more useful place to ask the question. Planning gives you a chance to answer it while that place is still configured the way you intended.

Emma: There's a timing detail I'd put in that booking. The database returns to its previous configuration during the first update window after the seventy-two hours expire, not necessarily at the exact expiry second. If somebody compares Tuesday's result with Friday's, the dates alone won't tell them whether the conditions matched. But I also don't want the lesson to become 'only a perfect copy of production is worth testing.' That would leave a lot of useful checks undone.

Sandbox environment resource configuration (included section)

Alex: It shouldn't. A functional test can tell you whether the batch selects the correct records and produces the intended result even when it doesn't settle production performance. A controlled timing comparison can tell you whether a change is promising under the tested conditions. A larger rehearsal can give the team better evidence about the operational window. Those are all worthwhile, provided we don't rename one as another. For your evening batch, I'd report the result in plain language: this amount of work completed in this environment, with these conditions, and these differences from production remain. Then the owner can decide whether another test is needed. That is much more helpful than either 'the sandbox is useless' or 'it passed here, so we're guaranteed.' This preview narrows a real gap. It doesn't remove the need to say which question the test answered.

Sandbox environment resource configuration (included section)

Emma: There's another difference hidden in that rehearsal. A company might copy production into a sandbox and think it has brought the whole working day with it. But copying the data and matching the database configuration are separate actions. Microsoft's environment guidance describes precautions in a copied environment, including stopped job queue tasks and blocked outbound calls from extensions by default. That sounds sensible for avoiding accidental activity against live services, but it could also leave our evening batch waiting for something that isn't running. How would you keep the test useful without carelessly reconnecting the sandbox to everything?

Sandbox environment resource configuration (included section)

Precautions for sandbox environments with production data

Alex: I'd make those dependencies part of the test design, not restore every switch because a run looks incomplete. First identify what the batch needs from outside itself. If it normally receives a file, sends a message or waits for another task, decide how that interaction will be represented safely in the test. Some checks can use controlled inputs and stop before an external handoff. Others need an explicitly prepared test endpoint or coordinated task. The result should say which of those paths was actually exercised. A fast run with an important integration disabled isn't evidence that the complete overnight process is fast. Equally, a blocked test call isn't proof that the production database is slow. That's where the distinction helps the team investigate instead of arguing over one elapsed-time number. I'd also make sure the copied data is appropriate for the people doing the test. A sandbox isn't permission to expose information more broadly. The resource option changes the database configuration; it doesn't decide access, safe destinations or the meaning of the workload. For the company planning its evening run, the useful outcome is a rehearsal whose boundaries everyone understands, with a clear next step for any part that wasn't exercised.

Sandbox environment resource configuration (included section)

Precautions for sandbox environments with production data

Inventory age without historical assumptions

Emma: Our next item isn't a November release. It's an existing Excel report from the 2026 first release wave: Item Age Composition by Quantity and Value. Microsoft's documentation makes a distinction worth noticing. The report ages the inventory that's still remaining now. Choosing an ending date defines the age buckets; it doesn't recreate the stock balance on that historical date. I think that could change the answer to a very ordinary question. If someone asks for 'the stock that was old at the end of September,' what would you want them to clarify?

Item Age Composition by Quantity and Value Excel (report)

Item Age Composition: 2026 release wave 1 introduction (included section)

Alex: Whether they want a historical position or a decision about stock they can still act on. Those sound similar, but they aren't. Imagine the company received an item in June, still had it at the end of September, and sold it in October. A historical September report might need to include it. A report about today's remaining stock has a different purpose: that sold quantity is no longer sitting there waiting for a decision. This report works with remaining quantities on open item ledger entries, using posting dates to place them in age groups. So it can help the team ask what older stock remains now. It isn't a time machine just because the request page contains a date. Before discussing the numbers, I'd write the question above the report: are we explaining a past balance, or deciding what to do with what's left?

Item Age Composition by Quantity and Value Excel (report)

Emma: For a warehouse meeting, I'd choose the second question. We can't make space by discussing the item that already left. But I can also see how a familiar month-end date would slip onto the report and make everybody assume they're comparing like with like. If we keep a copy this month and run it again next month with that same ending date, should we expect the old view to stay unchanged?

Alex: Not on that basis. The remaining quantities can have changed between runs. Keeping the same bucket date doesn't freeze the underlying remaining stock. A saved output is evidence of what that run showed; rerunning the report is a new observation. For your warehouse meeting, that can be a benefit, because the team sees what still needs attention. But for a month-to-month comparison, I'd retain the run date, filters, and the question each output was intended to answer. Otherwise a fall in the older-stock column could be read as a change in ageing policy, when the relevant stock simply moved out. And if the actual requirement is to reconstruct a prior inventory position, we'd choose and verify a report or analysis suited to that requirement rather than force this one to do it.

Item Age Composition by Quantity and Value Excel (report)

Emma: Let's stay with the meeting about what's left. The report includes quantity and value, and that gives two people a legitimate reason to pick different priorities. The warehouse may point at hundreds of old low-cost pieces. Finance may point at a few expensive items. I don't think either person is necessarily reading it incorrectly. But a single 'old stock' list can make it sound as if there's one obvious order in which to act.

Item Age Composition by Quantity and Value Excel (report)

Alex: They may be optimizing different things. If the immediate problem is tied-up inventory value, the expensive remaining stock deserves attention even if there are only a few units. If the problem is physical space, quantity alone still isn't the complete answer: a hundred tiny parts and a hundred large boxes occupy very different space. The report gives a useful starting point, not a universal priority score. I'd have the meeting agree the decision first. Are we looking for stock to review with sales? Are we questioning future replenishment? Are we investigating a particular location? Then use the relevant view and add the missing context. The quantity and value columns are useful precisely because they stop us treating every old unit as equally important. They open a better discussion rather than settle every trade-off.

Emma: I would also hesitate before calling everything in an older bucket 'unwanted.' Suppose the company deliberately holds a spare part for equipment that rarely fails. Its age might be exactly what we'd expect. A promotion isn't necessarily the right answer. The documentation itself says inventory age doesn't measure demand. So a useful discovery in the report could be a question to take to somebody who knows why the item is held, not an automatic instruction to mark it down.

Item Age Composition by Quantity and Value Excel (report)

Alex: That spare-part example is exactly why I'd separate observation from decision. The observation is that a remaining quantity falls into an older posting-date bucket. The business decision depends on demand, service commitments, replacement availability, and why the stock is there. An old entry doesn't by itself prove the item is obsolete, physically expired, or worth less than its recorded value. Those are different questions with different evidence. Conversely, a newer entry isn't automatically a good purchase. We could have bought too much last week. So I'd use the report to choose a manageable set of items for review, then ask the responsible person what would justify keeping, selling, moving, or changing future orders. That turns the age column into a conversation starter with a purpose, rather than a red label everybody feels compelled to remove.

Item Age Composition by Quantity and Value Excel (report)

Emma: And there's a practical way to keep that conversation manageable: filter by item or location instead of handing everyone the entire company inventory. I like that the Excel layout includes both print-oriented and analysis-oriented worksheets. It gives the meeting a usable starting point without making a new reporting project the first step. But once someone starts changing the workbook, I'd want the title to remain honest about the date and the scope. A polished spreadsheet can look rather more definitive than the question behind it.

Item Age Composition by Quantity and Value Excel (report)

Item Age Composition: available Excel layouts (included section)

Alex: I'd put the scope where the next reader can't miss it: current remaining stock, the run date, the bucket ending date, and the item or location filter. Those are presentation choices, but they protect the meaning of the numbers when the workbook leaves the meeting. The value calculation also deserves attention if someone wants to reconcile an individual item in detail; the documentation calls out a separate calculation for average-cost items. We don't need to turn that into a costing lecture here. We do need to avoid assuming that every value is simply the original receipt cost copied into a bucket. For most listeners, the immediate next step is smaller: open the existing report, choose one relevant location, and check whether it answers the operational question you actually have. If it does, use it. If you're asking about a historical balance, don't let the familiar date field persuade you that you've asked the same question.

Item Age Composition by Quantity and Value Excel (report)

Signed totals for visible payments

Emma: We'll finish with a smaller change that's generally available in 29.1: batch totals in payment journals. Before posting or exporting, you can review Batch Total in local currency at the bottom of the page. The detail that makes it useful is also the detail to remember: it totals the lines currently shown. If you've filtered the page, that's the set you're looking at. Alex, this sounds like a modest improvement worth using without turning it into a major project.

What's New or Changed in Business Central 2026 release wave 2 - Update 29.1

Record payments and refunds in payment journals

Alex: It is. A visible total helps you notice that the payment set on screen isn't the set you expected, before you take the next step. The field sums the signed amounts in local currency, so positive and negative lines offset each other. Think of a simple arithmetic example: displayed amounts of a thousand and minus two hundred produce eight hundred, not twelve hundred of gross movement. That's not a defect; it's what a signed total means. It also means the number shouldn't be described casually as 'all the cash leaving the bank.' If the result surprises you, inspect the lines and the filters before deciding that the calculation is wrong. The improvement is a quick check at the place where you're working, not a substitute for understanding the payment run.

Record payments and refunds in payment journals

Emma: I'd use it for the moment when somebody says, 'These are the payments we're reviewing today,' and then sends a screenshot with a filter still applied. The total can be perfectly correct for that screenshot and still not answer what the other person thinks was approved. That's a useful reminder to include the scope in the handoff, not just the number. And there's one boundary we should keep: posting a payment journal records the payment in Business Central; it isn't, by itself, proof that the money has reached the recipient.

Record payments and refunds in payment journals

Alex: Right. The actual bank transfer still has its own steps. For this feature, I'd keep the adoption equally small: notice the total, understand its sign and local-currency basis, and confirm the visible lines are the ones you intend to review. If a colleague expects the whole batch while you're showing a filtered subset, clear that up before treating the numbers as a disagreement. We don't need to invent another approval mechanism or build a spreadsheet to get value from this. Sometimes a standard feature simply removes a bit of mental arithmetic and makes a mismatch easier to spot. That's a perfectly respectable improvement.

Record payments and refunds in payment journals

Emma: We've looked at three previews that open up useful choices, and two reporting details that can make everyday decisions clearer. Shared data still needs an agreed source, an extra approval needs a distinct purpose, and a more representative sandbox needs a question worth testing. With inventory ageing and payment totals, knowing what the number includes is part of using it well. The direct references and practical notes are in the show notes, including the preview restrictions and the unresolved inconsistency in the master-data documentation. Thanks for listening to Central Perspectives from ITrators. Until next time.