Credits, corrections and connected data

A credit note still needs an application decision

Emma: Welcome to Central Perspectives, from ITrators. I'm Emma, with Alex. We're AI hosts, and our scripts and voices are AI-generated. Today, a supplier credit note gets a new route into Business Central, but choosing what it reduces still matters. We'll look at a correction that can fix a financial report without changing every ledger behind it, and at what it takes to know your Fabric data has actually arrived. There's also a useful new option for extension developers, and a planning change coming up very soon. Now, let's get started.

Emma: Payables Agent can now handle vendor credit notes. Microsoft lists that as generally available in update twenty-nine point one. That's a specific addition, not a claim that every company has already installed the update. Previously, our discussion was about invoices and purchase-order matching. A credit note turns the question around: the supplier says we owe less. Alex, before we get into the controls, what does the agent actually do with that document?

What is new in Business Central update 29.1

Alex: It reads the document and recognizes whether it is an invoice or a credit note from the content. It isn't simply treating any negative amount as a credit note. For a credit note, it prepares an unposted purchase credit memo. It doesn't run the purchase-order matching route used for invoices, and it doesn't post the result. So the useful change is getting the credit into a reviewable document. The next decision is which posted invoice, if any, that credit should be applied to.

Payables Agent

Emma: That last part is where I'd want the help. Imagine the company receives a credit for damaged goods, and the supplier has included an invoice reference. I don't want someone hunting through a long supplier account when the document already tells us where to look. But I also don't want a plausible-looking reference to become a decision nobody noticed. When does the system have enough information to suggest that original invoice?

Alex: The documented suggestion requires one eligible open posted invoice that matches the reference, for the pay-to vendor and the applicable currency. The important word is one. If several invoices could match, that ambiguity belongs in the review rather than being silently resolved. In your damaged-goods example, the reference is useful evidence, but it has to lead to an eligible invoice in the correct account and currency context. That is a more specific test than finding the same string somewhere in the company's history.

Process incoming purchase credit notes

Emma: So it saves the search when the match is clear, rather than promising to settle every argument in the supplier account. Let me complicate the example in a very ordinary way: the supplier sends one credit note covering two invoices. Both references are printed on the PDF. It would be easy for a reviewer to see the first invoice on screen and assume that the rest has been handled somewhere underneath. Does the workflow make that split for us?

Alex: No, that would be too much to assume. The current incoming-credit-note guidance says that when the PDF contains multiple invoice references, only the first is exposed and the system doesn't automatically select an invoice. That changes the reviewer's job. They need to compare the actual credit document with the proposed application, not treat the visible reference as a complete account of everything the supplier intended. I'd put a multi-invoice credit into an evaluation of this workflow precisely because a clean single-invoice example won't reveal that boundary.

Process incoming purchase credit notes

Emma: That changes how I'd describe the benefit to the accounts team. The agent is preparing the credit, not interpreting an entire settlement arrangement. There's still a gain there, especially if it removes retyping, but the reviewer needs the original document in view. What if the invoice is obvious and the amount isn't? Say the credit is larger than what remains open, or there's already another unposted credit waiting against that invoice. Are those hard stops?

Payables Agent

Alex: They are warnings to investigate, not a guarantee that finalization will be blocked. The documented review can warn about excess credit and other unposted credits against the invoice, but those warnings do not prevent finalization. That matters for your example: seeing a warning and being allowed to continue are not contradictory. The software is showing a condition that needs a business decision. The reviewer still has to decide whether the amount and application make sense, including any other credit already being processed.

Process incoming purchase credit notes

Emma: I can see a tempting shortcut here: if the system lets me finalize, it must have approved the accounting. That would be a very comfortable interpretation, and apparently the wrong one. I'd want the review instructions to say what to do with those warnings, rather than just where the button is. And if the reviewer selects an invoice, is that selection final in the sense that the agent will respect it, or can it pick a different invoice later?

Process incoming purchase credit notes

Alex: Finalization uses the reviewed selection. It doesn't choose another invoice. It does recheck eligibility, so the review isn't a reason to ignore a changed state. But we should keep three things separate: selecting the intended invoice, finalizing the reviewed document, and posting. Payables Agent does not post invoices or credit notes. If the application fields are left blank, the credit remains unapplied; finalization does not secretly find a home for it. That can be a valid deliberate outcome, but it shouldn't happen because someone assumed a later step would decide.

Payables Agent

Process incoming purchase credit notes

Emma: I like leaving room for that deliberate outcome. Not every credit has to be forced onto the first invoice that looks close enough. It also gives us a better question than, ‘How many documents did the agent finish?’ We could ask whether the reviewer understood where each credit would go, and whether ambiguous ones were left for the right person. For a Danish company listening to this, though, is generally available the same as ready for its normal Danish supplier documents?

Alex: No. The current Payables Agent documentation describes English as the supported and validated language. Availability, setup and consumption requirements still apply, including Copilot Credits. So I wouldn’t take a release-table entry as evidence that the company’s whole document mix is supported. Start by checking the actual prerequisites and the language of the documents you want it to handle. Then evaluate the work you really have: a clear reference, an ambiguous reference, a multi-invoice credit, and a warning that genuinely needs review. We haven’t run that evaluation here.

Payables Agent

Emma: Then the useful promise is narrower, and better: less preparation work on supported documents, with a clearer review of what a credit reduces. Not ‘the supplier balance sorts itself out’. I’d keep the success measure close to that distinction. A fast document that somebody has to investigate again later isn’t much of a saving. We’ve put the direct guidance in the notes, including the application rules, so the release announcement and the actual working limits stay together.

Payables Agent

Correct the G/L dimension, understand the reporting boundary

Emma: Our next topic is an established feature, not a new November release: correcting dimensions on posted general-ledger entries. It’s useful when the amount is right but the department or project used for reporting is wrong. That sounds like a small repair, until different reports start telling different stories. Alex, if the finance team corrects the department on those entries, what has actually changed, and what hasn’t?

Correct dimensions on posted G/L entries

Alex: The correction changes dimensions on the selected G/L entries. It does not make the corresponding change in every other ledger. That is the boundary to understand before using it. Imagine the company’s profit-and-loss report puts an expense in the wrong department. Correcting the G/L dimension can address that reporting problem without treating the amount itself as wrong. But a report built from another ledger may still use the original dimension. You need to know which entries the report reads, rather than assuming every report follows the correction.

Correct dimensions on posted G/L entries

Emma: That makes me hesitate over the word ‘fixed’. The controller may have exactly the report they needed, while a colleague is still looking at the old department elsewhere. I wouldn’t want to turn a straightforward correction into a project, but I would want to explain that difference before sending the month-end pack. How would you narrow the job so the team knows what it’s correcting, instead of chasing every number that looks related?

Correct dimensions on posted G/L entries

Alex: I'd start with the reporting question and the actual G/L entry selection. Which department is wrong, for which entries, and which reporting period is being corrected? Agree that scope with the person who owns the report. Then check the proposed entries, rather than beginning with a broad filter and hoping the result looks sensible. There's a practical sequence in the documented correction flow: select the entries before making the proposed dimension changes. Changing the selection later resets those proposals, so it's worth settling the population first.

Correct dimensions on posted G/L entries

Emma: That's a useful piece of friction to know about beforehand. If someone adds another group of entries halfway through, the original proposed changes don't just carry on unnoticed. I'd still want them to notice that reset, rather than think the correction had forgotten their work. Once the entries and new department are in place, what does validation tell us? Is it checking the business meaning, or whether the proposed dimension combination is permitted?

Correct dimensions on posted G/L entries

Alex: Validation checks the applicable dimension rules and restrictions, including value-posting requirements. It isn't deciding that the expense belongs to the department you intended. That still comes from the business context. Permissions matter too: the correction requires the relevant dimension-correction permission set, and dimensions can be blocked from correction. So there are two separate questions. Is the proposed change allowed? And is it the right change for these entries? Passing the first is useful, but it doesn't answer the second.

Correct dimensions on posted G/L entries

Emma: Let's stay with the month-end example. The team has found the entries and validated the change, but the report goes out tomorrow morning. I can imagine two reasonable preferences: correct it now so everyone sees the same result, or schedule the work when the system is quieter. I'm not sure 'do it immediately' is automatically the more responsible choice. Does the feature give the team a way to make that timing decision?

Alex: Yes, the correction can run immediately or be scheduled. For larger selections, the documentation points toward smaller batches or running outside busy hours. I'd connect that choice to when the corrected report is needed and how much work is involved, not make an arbitrary rule that every correction waits overnight. There's also an option to update analysis views. That adds work, but those views may be part of the reporting path. If the team uses them, the handoff should say when they will reflect the correction.

Correct dimensions on posted G/L entries

Emma: So 'the correction has run' and 'the report I'm sending has been refreshed' are still different statements. That's worth making explicit, particularly when someone else owns the analysis view. What would worry me more is discovering the following day that the wrong set of entries was corrected. Can the team get back to the previous position, or is this a one-way change once it has been applied?

Correct dimensions on posted G/L entries

Alex: There is correction history, and the most recent correction can be undone. But that doesn't mean you can freely unwind any older correction in any order. For other cases, the history and Copy to Draft help you prepare a new correction. I'd use that history to establish exactly what happened before deciding the next step. Otherwise, a second broad change can hide the reason for the first one. The feature gives you a traceable correction process, not permission to keep trying different departments until a report happens to look right.

Correct dimensions on posted G/L entries

Emma: And there's one more reporting relationship we shouldn't leave out: cost accounting. You said at the start that the other ledgers don't automatically follow. If the company uses cost accounting, I'd rather find that out while planning the repair than after somebody notices that the cost report and the G/L disagree. Is the right conclusion that dimension correction is unsuitable there, or that the team has another part of the impact to assess?

Correct dimensions on posted G/L entries

Alex: Another part of the impact to assess. Microsoft explicitly warns that cost accounting can become out of sync, and the affected cost-accounting model has to be identified manually. I wouldn't prescribe a blanket cleanup or deletion of cost registers from that warning. The useful action is to involve the owner of that model, identify what the changed dimensions affect, and agree the appropriate correction there. For a company that only needs the selected G/L reporting repair, the scope may be much simpler. The point is to find that out, not assume either extreme.

Correct dimensions on posted G/L entries

Emma: That brings it back to the original job nicely. We wanted the right department in the month-end report, not a mysterious promise that every related record would be rewritten. A defined set of entries, a permitted and meaningful change, and a clear reporting handoff make this a useful standard feature. The history matters if the decision later needs correcting. And if cost accounting is involved, that owner belongs in the conversation before we declare the whole job finished.

Correct dimensions on posted G/L entries

Connected, exported and ingested are different states

Emma: With Fabric, there's a fresh change in the available guidance. When we discussed Business Central mirroring earlier in the series, the story was a private preview with important details still unclear. Microsoft now describes a production-ready preview and documents the operating controls. That is still preview, not general availability. I'd like to skip another tour of possible dashboards and go straight to the practical question: when somebody says the connection is set up, do we actually have data ready to use?

FAQ about Business Central mirroring to Fabric

Alex: Not from that statement alone. There are separate steps: the connection is configured, Business Central exports the selected data, and Fabric ingests it. A successful setup is evidence about the first of those, not all three. For someone preparing a report, that distinction is very practical. They need to know whether the required company and tables have actually exported, and whether the data has arrived on the Fabric side. Otherwise, a reassuring setup screen can coexist with a report that has nothing useful to read yet.

Monitor and troubleshoot Fabric synchronization

Emma: The phrase ‘production-ready preview’ could easily blur that distinction too. It sounds encouraging, but it doesn’t tell the report owner that their tables are there. Before they even get to that check, what has to be in place? I’m especially interested in the setup burden. A useful analytics route can become a rather different proposition if the first instruction is to invent an application identity and start managing another secret.

FAQ about Business Central mirroring to Fabric

Alex: The current setup uses Microsoft’s Fabric Export application identity rather than asking you to create your own app registration and client secret. There is still a tenant service principal and the necessary access to configure. The prerequisites include an active Fabric capacity in the same tenant and access to the workspace. So this is not a no-setup feature, but the documented identity route is more specific than a do-it-yourself integration. You should follow the current guide and check extension compatibility, not copy the assumptions from the earlier private preview.

FAQ about Business Central mirroring to Fabric

Set up Business Central data mirroring to Microsoft Fabric

Emma: That avoids the company having to manage its own registration and secret for this route, even though permissions and capacity don’t disappear. Another earlier expectation was very broad table coverage. If a report owner hears ‘Business Central data in Fabric’, they might reasonably assume that any table or field they can name will be available. Does the newer documentation support that, or do we need to correct the way we describe the scope?

Set up Business Central data mirroring to Microsoft Fabric

Alex: We need to be more precise. The current controls work with eligible normal tables; external, linked and virtual tables are excluded, along with certain sensitive internal tables. And selecting all fields at table level is not a guarantee that every kind of AL field will be mirrored. I would take the actual data requirements for the report and compare them with the eligible selection. That’s a better starting point than designing the report around a broad ‘everything is available’ assumption and discovering the exceptions after setup.

FAQ about Business Central mirroring to Fabric

Manage Business Central synchronization with Microsoft Fabric

Emma: Let’s make the report owner’s question concrete. They need a daily view from one company, the required tables are eligible, and the connection has been configured. Someone says export is enabled. I’m still waiting for the part that tells me whether this morning’s run worked. Where would you look next, especially if the overall screen looks calm and there isn’t an obvious red warning demanding attention?

Monitor and troubleshoot Fabric synchronization

Alex: At the run details, broken down by company and table, and then at ingestion on the Fabric side. The setup and export flags don't establish that every table succeeded. There's a particularly useful monitoring detail: rows with zero changes can be hidden by default, including error rows. So an empty-looking result isn't automatically evidence that there was nothing to process. Before drawing that conclusion, make sure the view includes the rows you need to inspect. Then follow the specific company and table rather than treating the connection as one undifferentiated success.

Monitor and troubleshoot Fabric synchronization

Emma: That's the sort of detail that changes an investigation. 'Nothing changed' and 'the error row is filtered out' would lead people in completely different directions. If the team wants to automate monitoring later, can it simply collect the detail records and attach them to whichever run happened most recently? I'm wary of a monitoring dashboard that looks more precise than the records underneath it really are.

Monitor and troubleshoot Fabric synchronization

Alex: That shortcut needs care. The documented monitoring API detail does not carry a run identifier. In the client, opening Details from the run gives you the run-bound view. So I would use that for a specific-run investigation and avoid presenting an inferred API association as an exact one. It's the same report owner's question, really: which run, which company, which table? A dashboard is useful when it preserves those distinctions, not when it fills in missing relationships with a confident-looking label.

Monitor and troubleshoot Fabric synchronization

Emma: Suppose that investigation finds a problem. My instinct would be to stop exporting while the team works out what happened, not immediately erase the setup and begin again. But interfaces often put a tempting reset action nearby. What is the difference here between stopping, resetting the whole connection, and resetting one table? I'd like that choice made before someone reaches for the most dramatic button.

Alex: Stopping preserves the state. A full reset removes connection state, watermarks, and synchronization overview and detail history, so it's a much broader action than pausing exports. Resetting a table requests a new snapshot for that table without clearing the whole connection. Those are different recovery choices, not three names for refresh. And a reset is not a general schema-repair tool. I'd first identify whether the problem is limited to a table, involves export, or is on the ingestion side. Then choose the action whose scope matches what the evidence actually shows.

Manage Business Central synchronization with Microsoft Fabric

Emma: So my pause-first instinct is reasonable when the aim is to stop more exporting while we investigate, but it isn't a repair by itself. There's another small ownership question I'd put beside that. If a table was selected through a package, and someone deactivates that package, they might assume the table has gone from the export as well. Is package selection the only place that can keep a table included?

Manage Business Central synchronization with Microsoft Fabric

Alex: No. Another package or a manual selection can still keep it included. That means changing one package isn't sufficient evidence that the table is no longer selected. Check the resulting selection, not just the change you made. It's a modest point, but it helps the person who has to explain why a table is still exporting. More broadly, this new documentation gives administrators something concrete to operate. We can now discuss selection, run details and recovery scope instead of treating mirroring as a future architecture diagram.

Manage Business Central synchronization with Microsoft Fabric

Emma: And the report owner gets a better handoff than 'Fabric is connected'. They can ask whether the required data was selected, exported and ingested, with the actual run in view. That doesn't give us a fixed delivery-time promise, and we haven't tested a tenant here. It does turn an encouraging preview announcement into a set of useful questions for a real setup. I'd call that progress, while keeping the preview label firmly attached.

FAQ about Business Central mirroring to Fabric

Monitor and troubleshoot Fabric synchronization

A newly possible key for an actual query

Emma: For developers, Business Central twenty-nine and runtime eighteen introduce a more flexible table-extension key. You can mix fields from the base table with fields defined in your own extension. Earlier in the series we discussed enabling and disabling existing indexes. This is a different question: what index can you define in the first place? Alex, give us the plain-language version before we disappear into property names. What kind of work does that newly possible combination help express?

Table keys in AL

Alex: Think of an extension that adds a processing status to records that already have a date in the base table. A work queue might need records with a particular status within a date range. The query naturally uses both fields, but they come from different parts of the table definition. The new capability lets the extension define a key using base fields together with its own fields. That gives the developer another way to support the actual access pattern, rather than considering those field groups separately.

Table keys in AL

Emma: So the queue doesn't become a new business feature just because a key is possible. The opportunity is underneath it: expressing the combination the queue already asks for. I can see why that matters to an extension owner. But 'can define' is doing some work in that sentence. Would you go looking for every place that combines a base field and an extension field, and add a key wherever you find one?

Table keys in AL

Alex: I wouldn't. I'd start with a query that matters and evidence that its current behavior is a problem. A key has a maintenance cost when data is written, so adding one is a trade-off, not a free improvement. In the queue example, I'd want representative data and the actual filters, then compare the read behavior and the writing workload before and after the change. The new syntax gives us an option to investigate. It doesn't establish that this queue is slow, or that this particular key is the best answer.

Table keys in AL

Emma: I'd be more interested in the queue that people wait for every morning than in a list of places where the new syntax could be demonstrated. Let's keep the example there. Suppose the extension's status identifies work still to do, and the date narrows that work to the relevant period. If another extension also adds a useful classification, can we simply include its field in the same key and finish the picture?

Alex: No. The documented boundary is base-table fields plus fields from your own table extension. Fields defined in another table extension still cannot be used for that key. That means the ownership of the field matters, even when all three fields appear together in the application. Before designing the key, check where each field is defined. Your queue may need a different query or design decision if its important classification belongs to another extension. The new capability doesn't remove that boundary.

Table keys in AL

Emma: That makes the ownership check part of understanding the queue, rather than an unpleasant surprise after writing the definition. What about the other familiar key properties? A developer might hear that mixed fields are allowed and assume all the usual table-key choices have become available in an extension too. Are we talking about a broad removal of extension restrictions, or a specific expansion inside the existing rules?

Alex: A specific expansion. The versioned guidance covers mixing those fields in KeyFields, SumIndexFields, SqlIndex and IncludedFields. Other restrictions remain: Unique and Clustered are not supported on keys in table extensions. I wouldn't turn that property list into a prescription for the queue. It tells the developer where the new field combination is permitted; the design still follows the actual query and workload. The important change is that a previously unavailable combination can now be expressed, not that every kind of key has become valid.

Table keys in AL

Emma: There's a documentation wrinkle worth flagging while people look this up. The table-keys page still contains an older example comment saying mixed keys aren't allowed, while the versioned limitations section describes the new runtime-eighteen behavior. I'd rather mention that than send someone to a page that appears to disagree with us. If they maintain an extension across versions, which part of that tension should guide the change they actually ship?

Table keys in AL

Alex: The explicit version requirement has to be part of the decision. We are describing runtime eighteen and Business Central twenty-nine, not retroactively changing what an older target supports. Check the target and the current versioned guidance, then compile and test the intended definition in that context. We haven't done that compilation or benchmark in this discussion. For an extension supporting older versions, adopting the new capability is also a compatibility decision. It's worth making deliberately rather than copying a definition into every maintained branch.

Table keys in AL

Emma: That's a sensible place to leave the developer with work to do: choose one worthwhile queue or query, check field ownership and target version, and investigate the trade-off on representative data. Not a campaign to add indexes because a new option exists. I do like this kind of change. It doesn't need to look impressive on a screen to give someone a better way to support an everyday task. The benefit still has to be demonstrated in the workload that matters.

Table keys in AL

Keep the planning context before Release Planner retires

Emma: One short planning notice before we finish. Microsoft has announced that Release Planner will retire by November fifteenth, with Dynamics 365 and Power Platform roadmap information moving into the AI at Work roadmap. That date is close, so it's worth checking any planning routine that depends on the old site. The particularly practical point is that personalized saved views won't carry over. Alex, what would you preserve rather than just replacing the bookmark and hoping for the best?

One always-on roadmap: transition milestones

Alex: The purpose of the view: which product areas the team follows, which changes it is considering, and who needs to act on them. A bookmark alone won't preserve that working context. I'd review the saved views while they are still available and keep the information the team needs to reconstruct its planning routine. The announcement does not say that those personalized views automatically migrate. And a roadmap remains a place to discover direction and planned changes, not proof that a feature is enabled in your environment.

One always-on roadmap: transition milestones

Emma: And this is a change in where the planning information lives, not an announcement that Business Central's release schedule has been cancelled or replaced. Microsoft says product release and deployment schedules are unchanged. For the implementation detail, we still need the relevant Learn documentation; for tenant-specific notices, Message Center still matters. That gives the new roadmap a useful role without asking it to answer every operational question. We've linked the official transition announcement in the notes.

One always-on roadmap: transition milestones

Emma: That’s all for this edition of Central Perspectives. The new credit-note workflow is most useful when the reviewer understands the application decision. Correcting a G/L dimension can solve a reporting problem, provided the team knows what remains unchanged elsewhere. And Fabric’s newer operating guidance gives us a much clearer way to ask whether the required data has arrived. Those are practical opportunities, not reasons to skip the details that make them work. Our direct references and notes are available with the episode. Thanks for listening to Central Perspectives, from ITrators.