

Small changes, clearer choices

Keep the attachment, move the file

Emma: Welcome to 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 exploring useful standard capabilities alongside selected October changes. Can attachments move out of the database without changing where people work? When does a simpler stock transfer remove duplicate effort, and when does it remove information you need? We'll also look at understanding access and finding a shorter route back to your analysis. Now, let's get started.

Emma: Let's start with something useful that arrived back in April. Business Central can store document attachments externally as a standard capability. Imagine a wholesaler adding supplier PDFs and supporting documents to records every day. Those files accumulate, but people still want to find them where they work. Alex, if that company chooses external storage, what actually changes for the person attaching a document or opening it later?

[Store document attachments in external file storage](#)

[External storage support for document attachments](#)

Alex: For that person, the familiar attachment workflow stays in Business Central. They don't have to start filing supplier PDFs in a separate application just because the storage location changes. Behind that workflow, an administrator sets up an external file account in Business Central, connects it to the storage provider, and assigns it to the Document Attachments scenario. With that enabled, new attachments go to that external storage instead of keeping the file itself in the Business Central database. The supported choices are SharePoint, Azure Blob Storage, and Azure Files. So, for your wholesaler, the useful separation is this: staff still attach and find documents alongside their work, while the company can choose a different home for the files. That's the benefit I'd start with, rather than promising a cheaper or faster system.

[Store document attachments in external file storage](#)

Emma: You said new attachments. Our wholesaler already has years of documents sitting in Business Central. What happens to that pile when they switch this on?

[Store document attachments in external file storage](#)

Alex: That pile stays where it is. Switching the scenario on changes where new attachments go; it doesn't move the old ones. For those, there's a separate action called Storage Sync, where you choose the direction of the move. It can move existing attachments out, and it can bring files back into the database if you decide to stop using external storage. That second direction is useful: this isn't a one-way choice. My practical suggestion would be to try the workflow in a test setup first, with a document you can instantly recognise. Attach it, open it, and check its contents. Then plan the existing-file move as a separate step, rather than bundling everything into the switch-on.

[Store document attachments in external file storage](#)

Emma: Being able to bring them back makes this easier to consider. But if I'm running that wholesaler, the documents still look like part of Business Central to me. When someone says, "We can restore our Business Central backup," what does that now mean for my attachments?

[Store document attachments in external file storage](#)

Alex: It means you need to distinguish the attachment record from the file. Business Central's backup doesn't include or manage those external files, so restoring the database isn't a promise that the documents themselves have been restored. I'd make the recovery test very ordinary: find a supplier document in the restored system, then actually open it. Seeing its name isn't the whole test; the external file must still exist and be accessible. Whoever looks after that storage needs to include those files in the recovery plan. It's a separate responsibility, not a reason to reject the option.

[Store document attachments in external file storage](#)

Emma: A filename can look reassuringly complete, can't it? That brings me back to your test setup. If we copy the company to try this out, do we also get an independent copy of its attachment files?

Alex: Not just by copying the company. The attachment records can come across while still pointing to the original external files. So two company records can lead you to the same file, not two independent copies. There's a Migrate Files action to move attachments from the previous company or environment folder into the current one. But that's a migration operation; it isn't evidence that you've made an isolated backup copy.

[Store document attachments in external file storage](#)

Emma: Then tidying up an attachment in that test company might not be quite as local as it looks. When I press delete in Business Central, what am I actually deleting?

Alex: That depends on a setting: Delete External File on Attachment Delete. With it enabled, deleting the attachment in Business Central also sends a request to delete the external file. With it disabled, Business Central removes the link, but the file stays in external storage. Those are two quite different meanings of "tidy up." I'd choose that behaviour deliberately with whoever owns the documents, rather than leave staff to discover it. Keeping an unlinked file isn't, by itself, a retention policy either. You still need to decide what should happen to it and when.

[Store document attachments in external file storage](#)

Emma: That makes me wonder about the scope of the connection itself. If our wholesaler already uses SharePoint for several departments, does connecting Business Central require access to all those sites, or can they limit it to a specific place for these attachments?

Alex: They can limit the connection to selected SharePoint sites. With the default Microsoft Graph connector, that involves selecting the site-access scope and separately granting the application write access to each site it will use. Choosing the limited scope alone doesn't grant access. So your wholesaler can plan a specific home for these attachments without treating every department's SharePoint site as part of the connection. It does need an app registration and the right authentication setup, not just a folder name pasted into Business Central. But that's a defined setup task; it isn't a requirement to make the connector an administrator of all their sites.

[Store document attachments in external file storage](#)

Emma: We've spent a moment with the delete button, but I don't want that to bury the opportunity. What I'm hearing is a choice about who looks after the files, not a reason to be afraid of external storage. When would you seriously consider this standard option for a company whose current setup is working?

Alex: I'd look seriously at a company adding substantial volumes of attachments, particularly if it already has someone responsible for one of those supported storage services. There's a real decision there about where the growing collection belongs, with an owner who can evaluate it. If a company attaches a handful of small documents and has no storage problem to solve, I wouldn't manufacture a migration project. And I wouldn't start either conversation with "Which add-on do we need?" This capability is already in the standard product. It needs proper account setup and authentication, so it isn't simply a switch for every user. But evaluating that standard option is a sensible first step, not a commitment to moving every file.

[Store document attachments in external file storage](#)

A new field without another proxy extension

Emma: There's another standard capability worth knowing about, this time for companies connecting Business Central with Dataverse. Since April, adding a newly available Dataverse field to an existing mapping needn't mean changing a proxy extension. Imagine a sales team has added a customer classification in Dataverse, and Business Central already has the corresponding field. They want that value to travel between the systems. Alex, what work does this change actually take off their plate?

[Map new Dataverse fields in Business Central](#)

Alex: They can avoid changing an extension just to make the integration recognise that new field. Think of the proxy as the integration layer's description of the fields available in Dataverse. Previously, exposing a custom field could mean updating that description through an extension. Now an administrator can refresh the Dataverse field list, add a supported field to the existing mapping, match it with the Business Central field, choose the synchronisation direction, and enable it. Your classification has a route across without that particular code change or rebuilding the mapping. It doesn't create the Business Central field for you, though. In your example, that destination already exists, which is an important part of why this is simpler.

[Mapping the tables and fields to synchronize](#)

[Map new Dataverse fields in Business Central](#)

Emma: The code change goes away, but that choice of direction still needs an answer. Suppose sales owns the classification in Dataverse and finance reads it in Business Central. “Both directions” sounds convenient. Is it still a sensible choice if people on either side might edit the value?

[Mapping the tables and fields to synchronize](#)

[Map new Dataverse fields in Business Central](#)

Alex: For the arrangement you’ve described, I’d start with one direction: from Dataverse into Business Central. Say sales changes a customer’s classification from standard to strategic. Finance needs that new value, not a second place to maintain a competing classification. Two-way synchronisation can make sense when both teams genuinely need to edit it. But selecting both directions only defines how values can travel; it doesn’t settle whose judgement should win when the teams disagree. I’d resolve that ownership question first. Otherwise, a convenient setting leaves the business decision sitting there, unanswered.

[Mapping the tables and fields to synchronize](#)

Emma: Let’s follow that strategic customer across, then. We’ve matched the classification fields, but each system has its own customer record. What tells the integration which customer in Business Central should receive that new value?

Alex: That’s the coupling: the link saying this Dataverse record corresponds to this Business Central customer. By default, the synchronisation jobs work only with coupled records. Mapping the classification field tells the integration where that value belongs on the linked customer; it doesn’t magically pair up every customer in the two systems. So an already-integrated customer can gain another useful piece of information. Connecting previously unlinked customer records is a separate question.

[Mapping the tables and fields to synchronize](#)

[Map new Dataverse fields in Business Central](#)

Emma: Suppose the existing customer updates perfectly, but a brand-new customer from sales never appears in Business Central. I can imagine “only coupled records” looking like the obstacle. Would unticking that be a fix, or are we changing more than that?

[Mapping the tables and fields to synchronize](#)

Alex: You’re changing more. Clearing that setting allows the synchronisation to create destination records for source rows that aren’t coupled yet. That could be exactly what the business wants, but it’s a different job from updating a classification on an existing customer. Suppose the sales system contains both approved customers and prospects who aren’t ready to become customers in Business Central. You’d want to define which of those should cross. Table filters can limit the Business Central records, and integration table filters can limit the Dataverse rows. With the intended direction and scope, record creation can be useful. Unticking one box, though, isn’t a promise that every missing customer will appear.

[Mapping the tables and fields to synchronize](#)

Emma: That covers which records we bring across. What about the other end of their life? Suppose a linked record is deleted in one application but still exists in the other. What can the owner choose should happen then, rather than leaving the two teams to guess?

Alex: The table mapping has a deletion-conflict strategy for that. You can choose to restore the deleted record, remove the coupling between the records, or do nothing and resolve the conflict manually. Those choices represent different intentions. If a record was deleted by mistake, restoring it might be exactly what the team wants. If the intention is to end the relationship between those records, removing the coupling addresses that link instead. And if a deletion needs someone to understand the circumstances, manual resolution leaves that decision with them. I wouldn't pick one simply because it sounds the most automatic. The useful part is being able to express what a deletion should mean for this particular integration.

[Mapping the tables and fields to synchronize](#)

Emma: A record reappearing could actually be the agreed behaviour, not something broken. What about a less final disagreement: both teams change the customer's classification before the next sync? We've already chosen which way that field can travel. How does an update-conflict strategy fit alongside that direction setting?

[Mapping the tables and fields to synchronize](#)

Alex: Direction says where the value can travel. The conflict strategy says how the integration should resolve competing updates when they arise. At the table mapping, you can choose to send an update to the integration table when synchronising towards Dataverse, or get an update from it when synchronising towards Business Central. Or you can leave the conflict for manual resolution. In plain terms, you're choosing whether the relevant synchronisation direction should settle that clash, or whether somebody needs to decide. These are table-mapping strategies, not a new conflict policy invented by adding the classification field. So the existing strategy deserves attention when that new field joins the conversation between the systems.

[Mapping the tables and fields to synchronize](#)

Emma: There's a useful detail in the documentation here: new field mappings added to an existing table mapping start disabled. I like that separation. Someone can prepare the classification mapping without that new mapping immediately taking part in synchronisation. Creating it and deciding to use it are distinct steps.

[Mapping the tables and fields to synchronize](#)

Alex: And you can make the check specific: in a test setup, enable that mapping, change the classification on a known, coupled customer, and check the value at the other end. Not just "the customer synchronised," but "this field did what we intended." There's an earlier detail worth catching, too: if the field isn't in the lookup, refreshing again may not help. The guide supports particular field types; FlowFields and FlowFilters aren't eligible. A missing field can be a capability boundary, not a stale list.

[Mapping the tables and fields to synchronize](#)

Emma: Before we lose anyone at "FlowField," what does that exclusion mean for a business request? Can you give us one example of the distinction, beyond the field's technical label?

[Mapping the tables and fields to synchronize](#)

Alex: Take our classification: it's a stored value someone maintains. Now imagine the request is instead, "Send this calculated total to the sales team." If that total is a FlowField, Business Central calculates its result rather than storing it as an ordinary field value. Both can look like a number or a label on a page, but they're different things for the integration. The mapping guide excludes FlowFields. That doesn't prove the business request is impossible; it means this particular setup route doesn't cover it. A developer may need to assess another approach, including when the total should be calculated and how current sales needs it to be.

[Mapping the tables and fields to synchronize](#)

[FlowFields overview](#)

Emma: And there's a release distinction worth keeping straight. The new-field capability became generally available in April. The broader New Table Mapping assisted setup is still labelled preview in the operational guide. Those aren't interchangeable claims. Our customer-classification example is the narrower opportunity: adding a supported field to a mapping that already exists, without that proxy-extension change.

[Mapping the tables and fields to synchronize](#)

[Map new Dataverse fields in Business Central](#)

Alex: That narrower improvement is worth having. An existing integration team can stop treating every newly exposed field as an extension-maintenance task. I like that much more than a sweeping "integration without code" promise. There's still a difference between getting a value across and getting a usable value across. Two classification fields might use different codes for the same meaning; enabling the mapping doesn't decide the transformation you need. And conflicting updates still need a chosen resolution strategy. Removing the repetitive plumbing leaves more room for those decisions, which are the part the business actually cares about.

[Mapping the tables and fields to synchronize](#)

[Map new Dataverse fields in Business Central](#)

Direct transfer is a process choice

Emma: Now, from moving customer information to moving stock. October's Business Central 29 release adds the ability to post direct transfer orders from warehouse-enabled locations. Direct transfers themselves aren't new. Imagine a managed central warehouse replenishing a maintenance stockroom, set up as a separate Business Central location without a receiving warehouse process. Alex, what practical step can this change remove for the people moving those supplies?

[Update 29.0 for Business Central 2026 release wave 2](#)

[Transfer items between locations](#)

Alex: It can remove a separate receiving posting at the stockroom. With a supported setup, the outbound warehouse process can complete the direct transfer with shipment and receipt posted together. The warehouse still does its work: if picking is required, the pick is created and registered before the warehouse shipment is posted. We haven't made the supplies travel by clicking a button. For that maintenance stockroom, I'd see the opportunity where the physical handover is simple and the posting can accurately represent stock reaching its destination. You retain the outbound handling you need without asking someone at the other end to perform a separate receipt posting just to finish the transfer.

[Transfer items between locations](#)

Emma: "Posted together" makes me pause there. What if that simple stockroom is at another site, and the supplies still need time to get there? Is having no receiving warehouse process enough reason to choose direct posting?

Alex: No, the time between leaving and arriving still matters. If you need to show that the supplies have left the warehouse but haven't reached the other site, a normal in-transit transfer does useful work. It tracks that quantity in transit, with shipment and receipt posted separately. That receiving step is then a record of something that actually happened, not just administrative repetition. Direct posting isn't inherently a bad choice; it suits a different process. I wouldn't let the absence of warehouse receiving at the destination decide how we represent the journey.

[Transfer items between locations](#)

Emma: There's a detail in the documentation that seems to muddy that distinction. With direct transfer enabled and the posting method set to Shipment and Receipt, you can still select an In-Transit Code. Why doesn't that give us the separate departure-and-arrival process you've just described?

[Transfer items between locations](#)

Alex: Because the code changes the location used in the item ledger entries, not the timing of the two postings. In that direct Shipment and Receipt mode, the entries use your selected in-transit location rather than a blank location, but shipment and receipt still post together. The name does invite the other interpretation. You haven't created an open period where stock is waiting to be received. For that, it's the separate shipment-and-receipt process we discussed, not simply a transit code filled into a direct transfer.

[Transfer items between locations](#)

Emma: That field name is doing us no favours. But if both direct methods post the two sides together, what would make a company choose Shipment and Receipt rather than Direct Transfer? What difference would they actually notice afterwards?

Alex: The documents left behind are different, and that can matter to the accounts team. Direct Transfer creates a posted direct-transfer document. Shipment and Receipt creates a posted transfer shipment and a posted transfer receipt, even though they post together. That receipt gives you somewhere to assign an item charge; the Direct Transfer method doesn't offer that assignment. Imagine a transport charge that the company needs to allocate to the transferred goods. The receipt can therefore be more than another piece of paperwork. I'd bring that accounting requirement into the choice, rather than assume the method with fewer documents is always the better fit.

[Transfer items between locations](#)

Emma: That extra document earns its place, then. Back at the warehouse, suppose only part of the maintenance team's requested supplies is available now, with the rest to follow. Can a direct transfer handle that first portion without treating the entire request as finished?

Alex: With the Shipment and Receipt method, yes, subject to an important distinction. You can post part of the order, but the quantity to ship must equal the quantity to receive. So the available portion can move and be received together, while the rest remains for later. That isn't the same as saying, "We sent this much, but only some of it arrived." Combined direct posting doesn't give you that separate receipt decision. The method actually named Direct Transfer doesn't allow partial posting. So the maintenance team's perfectly ordinary request does affect which method fits.

[Transfer items between locations](#)

Emma: While we're talking quantities, suppose the warehouse spots a wrong quantity just after posting. With the two sides completed together, do they still have the same Undo Shipment option they might use on a shipment that hasn't been received yet?

Alex: Not that same option. With a normal transfer, Undo Shipment is available while the shipment hasn't yet been received. The comparison explicitly lists it as unavailable for both direct posting methods. Practically, you've completed both sides rather than left an unreceived shipment waiting there. That doesn't mean a mistake can never be corrected; it means we shouldn't promise this particular undo route. For the warehouse process, I'd make the quantity confirmation part of completing the direct movement. The useful question is whether the team knows the quantity is right at that point, instead of relying on a later receiving step to establish it.

[Transfer items between locations](#)

Emma: I'd still want to lose the duplicate entry if the team already confirms the quantity properly. How would you try this with the warehouse team's actual work, so they can demonstrate that the simpler posting fits rather than just show that the button works?

Alex: I'd have them walk one ordinary replenishment through a test environment in the order they actually do the work. Who confirms what was picked? Who records that the supplies have been handed over? Put the proposed posting at its intended point in that sequence, then look at the quantities on both locations and the documents produced by the chosen method. Ask the people doing the work whether those results describe what has happened. That shows whether you've removed duplicate entry or merely moved an important confirmation out of sight. It's the team demonstrating its process, not just somebody demonstrating a posting button.

[Transfer items between locations](#)

Emma: We've deliberately kept the receiving end simple, though. The headline says warehouse-enabled locations, which sounds quite broad. If that maintenance stockroom grows and starts requiring warehouse receipts, or introduces inventory put-away, how does that change the options we've been discussing?

[Update 29.0 for Business Central 2026 release wave 2](#)

Alex: Those two changes lead to different answers. If the destination requires warehouse receipts, neither direct method supports it. The same exclusion applies to a destination using directed put-away and pick, the more advanced warehouse setup. Inventory put-away is different. The direct Shipment and Receipt method doesn't support that destination, but the Direct Transfer method supports combinations that do. Microsoft gives an example with inventory pick at the sending location and inventory put-away at the receiving location. You can post the inventory put-away and create both transfer transactions. So receiving-side handling isn't universally ruled out. But that example still completes both sides; it doesn't introduce a separate in-transit waiting stage. The exact combination of the two locations and the posting method matters.

[Transfer items between locations](#)

Emma: There are several workable routes here; we'll leave the full comparison in the notes. For the warehouse manager who hears this headline and thinks, "Could we stop entering the same movement twice?", what's the first question you'd ask about their process?

Alex: What does the second posting tell you that you didn't know at the first? If the answer is, "Nothing; the supplies are handed over in one operation, and someone later copies the same information into a receipt," there may be a useful simplification here. If the answer is, "It confirms what arrived after another team checked it," that's a separate piece of information worth preserving. I'd start there, before looking for a setting to change. The useful improvement is removing repeated administration, not removing the moment when someone establishes what actually happened.

Where did that access come from?

Emma: Here's a different question for the person looking after Business Central: why can someone do something you didn't expect? Permissions Overview, available since May, is a useful discovery for that investigation. Imagine a colleague can edit customer information, and the system owner wants to understand where that access comes from. Nobody has decided it's wrong yet. Alex, where does this overview help you start, before you touch anyone's permissions?

[Define granular permissions](#)

[Audit user and group permissions across apps](#)

Alex: It gives you a route from the object you're investigating to the permission sets that include permissions for it, across the installed apps and extensions. For your example, I could start with the relevant customer page or table, rather than guessing from permission-set names. The overview can be filtered by object, permission set, or extension. Then the Users and Security Groups FactBoxes show where a selected set is assigned. That helps turn "someone can edit this" into a more specific investigation: which sets are involved, and who has them? A page permission alone doesn't explain every possible edit, but now you have something concrete to trace instead of a list of reassuring-sounding role names.

[Define granular permissions](#)

Emma: A reassuring name is a fairly thin explanation when you're looking at an unexpected edit. Finding an actual permission in an assigned set gives us something firmer. But does that tell me everything this particular colleague can do, or only where one part of their access comes from?

Alex: Only part of the picture. Finding that assignment is useful, but for this particular colleague I'd also inspect Effective Permissions to investigate the rights they actually have. That's a user-level question, rather than just "what does this set contain?" There's a specific access requirement for that investigation: viewing another user's effective permissions requires SECURITY or SUPER. That applies to the person doing the inspection. It is absolutely not a suggestion to give the colleague SUPER so you can work out their access.

[Define granular permissions](#)

Emma: Suppose that investigation leads to a permission set assigned through a group. Our colleague has moved to a different job, while the rest of the team still does the same work. If the owner edits that set, are they solving just this person's access issue?

Alex: No, that changes the shared permission set for everyone who has it. If the rest of the team still needs to edit customers, removing that permission from their common set targets the wrong problem. In your example, I'd review how the colleague's new job should be reflected in their assignments or group membership, rather than immediately rewrite what the whole team's role allows. The Users and Security Groups FactBoxes are helpful here because they show where that set is used. You can see the reach of a proposed change before confusing one person's move with a change to everybody's work.

[Define granular permissions](#)

Emma: You also mentioned filtering by extension. Could a system owner use that before anything unexpected happens? Say they've installed an app and want to understand the permission sets it provides, rather than waiting for someone to ask why a colleague can edit something.

[Define granular permissions](#)

[Audit user and group permissions across apps](#)

Alex: Yes, and that's a useful conversation to have while the app is being introduced. Filter the overview to the extension, look at the permissions in its sets, then inspect where those sets are assigned. You can discuss something concrete with the person implementing it: this set covers these objects; which jobs is it intended to support? Then compare that intention with the assignments you can see. A permission being defined in an app isn't the same as a particular colleague having it. Keeping those two views separate makes the discussion clearer. You're learning how the app fits the team without changing anybody's access just to investigate.

[Define granular permissions](#)

Emma: Our surprised system owner might discover that the colleague's access is exactly what their job requires. Nothing needs removing after all. Would you still count that as a useful outcome from this overview, even though no permission has changed?

Alex: Absolutely. “We know why this access exists, and the person responsible for the role agrees it’s needed” is a much better answer than “It’s probably fine.” The useful outcome is an explanation you can stand behind, not a changed setting for its own sake. It also gives you a clearer starting point when that person’s work changes later. Sometimes the most valuable thing an investigation does is let you leave a working setup alone, for a reason.

The same key label, a refreshed file

Emma: A short notice now for administrators of on-premises Business Central, including installations using Dual Use Rights, rather than people simply using the online service. From October first, the licence-key version for Business Central 28 and later is labelled “28 plus.” Earlier releases keep version-specific keys. That could sound like one key covering every upgrade from now on. Alex, what does that plus sign actually change when the next upgrade comes around?

[Dynamics 365 Business Central license key version update](#)

Alex: The version label becomes reusable; the downloaded file doesn’t become a licence file you can keep using unchanged for every future upgrade. Microsoft says the download date determines which objects the 28-plus file unlocks. When upgrading to a new version released on or after October first, you still need a refreshed key. So I’d keep obtaining that current key in the upgrade preparation, rather than looking at an old file, seeing “28 plus,” and assuming the job is done. The useful distinction is the label versus the contents of the file you actually downloaded. Same label doesn’t necessarily mean the same unlocked objects.

[Dynamics 365 Business Central license key version update](#)

Emma: An upgrade checklist that only checks the label could miss that. Just to separate the two things: is this changing how administrators prepare a licence file, or does it also change when they’re expected to upgrade?

[Dynamics 365 Business Central license key version update](#)

Alex: The announcement explicitly says it doesn’t change the release cycle for major or minor versions. So I wouldn’t read the plus sign as a new upgrade timetable. Choosing and scheduling the upgrade is one decision; preparing a refreshed licence file for that upgrade is another. For the administrator listening, the useful action is small: make sure the preparation checks the key’s download date and suitability for the target release, not just whether its version label says 28 plus.

[Dynamics 365 Business Central license key version update](#)

Open the analysis you actually use

Emma: Back to an everyday convenience now, and this one is new in the October wave: bookmarking a specific analysis view on your Role Center. Imagine you’ve arranged an analysis of item entries to review quantities by location. You want to return to that particular arrangement, not just find the page again. Alex, what opens when you use this new bookmark?

[Analyze list page and query data using data analysis](#)

Alex: You land on the page or query already in analysis mode, with that particular view selected. So your quantities-by-location arrangement is the destination, rather than something you have to find after opening the ordinary page. It's a small difference, but if that's how you start a regular stock review, it removes the little detour every time. And it's a shortcut back to the analysis, not a saved picture of the figures. You're returning to the view you want to work in, not opening a report frozen when you made the bookmark.

[Analyze list page and query data using data analysis](#)

Emma: If I have several views of that same page, I'd want to tell them apart in the menu. What name does the bookmark get, and does it follow along if I rename the view?

Alex: It combines the page or query name with the analysis-view name, so you can distinguish those different routes into the same page. And yes, rename the view and the bookmark's caption updates too. Delete the view and its bookmark goes away. I like that: you're maintaining one view with a way to reach it, not creating a second independent copy that gradually gets out of step. A clear view name does useful work in the menu as well as inside the analysis.

[Analyze list page and query data using data analysis](#)

Emma: One practical snag before someone tries it: they have their analysis view, but can't see "Bookmark view" in its menu. What prerequisite should they check?

Alex: Check whether personalisation is enabled. The Bookmark view action isn't available when it isn't. So before concluding that the feature is missing or broken, ask the administrator how personalisation is intended to work in your setup. Once available, it's a handy way to put a frequently used view within easy reach.

[Analyze list page and query data using data analysis](#)

Emma: That's a good place to finish: less work getting to the thing you actually need. Today's standard capabilities can remove a coding chore, a navigation detour, or a genuinely duplicate posting. They don't remove the decisions underneath. Moving attachments still needs someone to look after the files; simplifying a transfer shouldn't erase a real arrival check. The useful question is which bit of effort your team can lose without losing something it relies on. Thanks for listening to Business Central Podcast from ITrators. Goodbye for now.