

THE COMPLETE GUIDE

Master Workflow Design Guide

Four connected best-practice guides - designing the templates, scheduling them, the notes that keep every job standard, and putting your own AI to work on all of it.

One platform, one standard:
build templates that capture how your practice delivers work, schedule them to run themselves, and surface the right client context at the point of work.

Setting up your workflow automation neatly from the start allows you to then connect your own AI assistant, like Claude or ChatGPT via our MCP, to gain real insight over your practice faster.

INSIDE THIS GUIDE

- 01 Workflow Templates
- 02 Template Schedules
- 03 Client notes, topics & link tags
- 04 Ask your workflow

01

GUIDE ONE

Workflow Templates

Here's how to design XBert workflow templates that deliver consistency, accountability, and clarity across every job, every time.

-
- 01 Why templates matter
 - 02 Anatomy of a good template
 - 03 The ten elements, in detail
 - 04 A real example: Payroll Processing
 - 05 Best-practice checklist



Templates aren't just for recurring work

The most common assumption about workflow templates is that they're for work that repeats on a schedule - weekly pay runs, quarterly reviews, annual compliance. And yes, templates are outstanding for that. But limiting templates to recurring jobs leaves a significant gap in your practice's consistency.

Consider the ad hoc jobs that land with no warning: a client calls about a compliance audit, you prepare a one-off grant application, or a new client onboards mid-year and needs a full history review. These jobs may happen only once for a given client, or for the practice once - but the way they should be delivered is entirely predictable. That's where templates earn their keep as a **standardisation tool**, not just a scheduling tool.

A template captures the institutional knowledge of how your practice delivers a particular type of work. A newer team member handles an unfamiliar job to the same standard as your most experienced one - and nothing is forgotten because someone was busy.

KEY PRINCIPLE

If your team would need to think about the steps from scratch each time a job type is performed, that job type needs a template. Frequency is irrelevant - **standardisation is the goal**.

Anatomy of a good template

Every XBert workflow template is built from several distinct elements. Getting each one right is what separates a template that genuinely drives efficiency from one that just adds administrative overhead. Each is covered in the section that follows.

ELEMENT	WHAT IT CONTROLS	WHY IT MATTERS
01 Name & title	How the job appears in lists, filters and reporting	Findability, clarity, consistency across the team
02 Description	Context, scope and guidance for whoever does the work	Reduces questions, sets expectations, supports training
03 Sub-tasks	The ordered checklist of actions to complete the job	Drives accountability; ensures nothing is missed
Sub-task prefixes	Identifies the role responsible for each sub-task	Communicates ownership; supports mid-workflow reassignment
04 Note link tag	Links the template or sub-tasks to XBert's note system	Surfaces client-specific notes at the point of work
05 Process	The parent process category the template belongs to	Enables reporting, filtering and workload management
06 Status	Starting state of the job when created	Enters the right workflow stage; can auto-reassign on update
07 Estimated time	Expected hours and minutes to complete	Supports capacity planning and workload visibility
08 Step number	Sequences a template within a multi-template process	Used only when one process is split - generally avoided
09 Group name	Classifies how the template is intended to be used	Distinguishes recurring (Schedule) from on-demand (ad hoc)
10 Attachments	Supporting documents, workpaper templates, or reference files	Puts the right resources at the point of work

Make it findable

The **template name** is the first thing your team sees in a job list, a filter or a report. It must identify the job type immediately, distinguish it from similar templates, and sort logically.

WHAT IT CONTROLS

How the job appears in lists, filters and reporting.

WHY IT MATTERS

Findability, clarity and consistency across the whole team.

BEST PRACTICE

- ▼ Use one consistent naming convention across all templates - decide on it once, then enforce it.
- ▼ Make the name clear in plain language, so the whole team understands it - not just the person who built it.
- ▼ Run a separate template per accounting product only where the workflow genuinely differs; otherwise handle product-specific variation via **sub-tasks** or **note link tags**.
- ▼ Use template tokens to auto-populate dynamic information - frequency, start and end dates - into the title. Type [in the title field to access them.

YOUR ASSISTANT READS THE SAME LIST

When you ask your assistant to create a month-end task from your template, it finds that template by name. A plain-language, consistent name is what lets a plain-English question land on the right template.

FROM OUR TEMPLATE

The example template in this guide is named **Process Payroll - With Timesheets - [dueday] - [ScheduleFrequency]**. The [dueday] and [ScheduleFrequency] tokens auto-fill the relevant period whenever a job is created - every job is immediately identifiable, with no manual renaming.



Brief the next person

The **description** field should contain a clear summary of the task for whoever picks up this job - especially ad hoc tasks as they may not have been done before, or newer team members.

WHAT IT CONTROLS

Context, scope and guidance for whoever does the work.

WHY IT MATTERS

Reduces questions, sets expectations and supports training.

BEST PRACTICE

- ▼ State the purpose of the job in one sentence.
- ▼ If this template is to be scheduled on one recurring period only, you may want to specify the frequency or trigger condition - e.g. "used for clients on a weekly, fortnightly or monthly pay cycle". Note though you do not need different templates for different recurring periods - that is what schedules are for.
- ▼ Note any prerequisites - what needs to be done, or confirmed, before this job can begin.
- ▼ Reference key documents or reports, and recommend a review of client notes - they often carry requirements that affect delivery.
- ▼ Keep it concise - two to four sentences. This is a brief, not a manual.

FOR AD HOC TEMPLATES

The **description** matters even more. Because these jobs don't happen on a regular cycle, whoever creates them may not remember the nuances. Use the **description** to capture the why and when clearly.



The heart of the template

Each sub-task should represent a discrete, completable action. The order matters - **sub-tasks** should flow logically from data gathering through to sign-off and lodgement.

WHAT IT CONTROLS

The ordered checklist of actions to complete the job.

WHY IT MATTERS

Drives accountability and ensures nothing is ever missed.

BEST PRACTICE

- Write each sub-task as an action - start with a verb.
- Make them granular enough to be checkable, but not so granular they become noise. One sub-task, one meaningful action.
- For micro-steps, use the Add Instruction element inside the sub-task - detail without cluttering the main list.
- Sequence so each step creates the conditions for the next, and include client-facing actions - sending, approving, lodging.
- If you decide to use prefixes, then all **sub-tasks** should have them, or none - it makes it clear who is doing which element if your template is for multiple staff.



Role prefixes - when & how

Prefixes make it clear who is responsible for each sub-task where a workflow involves more than one role, and support reassigning the work at a defined handover point.

BK **Bookkeeper**

Sub-tasks completed by the bookkeeper, where the workflow involves more than one role.

P **Partner**

Steps needing partner involvement - review, sign-off, lodgement, or client-facing actions.

MR **Manager**

Steps completed by the manager: review, payslip distribution, or client-facing actions.

NOT EVERY TEMPLATE NEEDS PREFIXES

Prefixes are used only where more than one person is involved and the work is reassigned at a defined point. If a job is completed entirely by one person, they add no value.



Example sub-task sequence

From the payroll template - checks and approvals first, then processing pay and distributing payslips, then filing and closing off with a client update.

01 BK Check for any XBerts and resolve

Pre-work check: clears flagged issues so payroll runs on clean data. You can also do this step by asking your assistant: what fired on this client this week, and what exactly fixes it?

02 BK Approve Timesheets

Timesheets must be approved before pay can be calculated accurately.

03 BK Process Pay

Bookkeeper runs the pay calculation for the period once timesheets are confirmed.

04 BK Update Status for Manager Review

Use **status** auto-assign to move the work to the manager's workboard for review and action.

05 MR Confirm payrun approved - re-assign to BK for finalisation work

Use a **status** change to auto-assign back to the bookkeeper for finalisation.

06 MR Send Payslips to employees

Manager distributes payslips so employees have their records for the period.

07 BK Submit payment file

Generates the bank payment file and submits it so wages can be paid.

08 BK Email the client

Confirms the pay run is finished and flags anything needing their attention.

09 BK Mark Task completed

Marks the pay run complete and updates the job **status**.

COMMON MISTAKE

Avoid bundling multiple actions into one sub-task, like "close off task and email client". It makes progress impossible to track - and when something goes wrong, you can't identify where in the **process** the issue occurred. Use **Add Instruction** if a step needs extra guidance without being split.



Surface the right note, in context

One of the most powerful elements. Note link tags connect your workflow directly to XBert's note system, surfacing client-specific information right at the point of work.

WHAT IT CONTROLS

Links the template or individual subtasks within a template to XBert's notes system

WHY IT MATTERS

Surfaces client-specific notes at the moment they're needed.

BEST PRACTICE

- Give every template a **note link tag**. If you can't identify one, that's a signal to review your note configuration.
- Be specific - **Payroll Dates** is far more useful than a generic tag.
- Note link tags can be applied at two levels, and both can add value:
 - **Template level** - the note surfaces when the job is opened, giving the team member and immediate view of any relevant client notes before work begins.
 - **Subtask level** - the note surfaces at the specific point in the workflow where it's most relevant, ensuring the right information is visible at the right step.
- Align tag naming with your note categories - consistency between the two systems is what makes the connection meaningful.

XBERT SUPPORT

Understanding Note Link Tags in Workflow Templates



Categorise for reporting

The **Process** field categorises your template within XBert's **process** hierarchy. It's what allows you to filter, report on, and manage workloads by job type across your whole practice.

WHAT IT CONTROLS

The parent **process** category the template belongs to.

WHY IT MATTERS

Enables reporting, filtering and workload management.

BEST PRACTICE

- ▶ Map each template to the correct parent **process** - in the payroll example, this is Payroll.
 - ▶ Don't create new **process** categories without a clear reason. Proliferating **process** types fragments your reporting and makes workload management harder.
 - ▶ Review your **process** structure when adding templates for new service lines - does a new category genuinely add clarity, or can an existing one be used?
-

RULE OF THUMB

Fewer, well-scoped **process** categories beat many narrow ones. The goal is clean filtering and reporting - not a category for every minor variation.



Set the starting state

The **Status** field sets the starting state of every job created from this template. All **statuses** are defined by your practice - you build the list that reflects your workflow, and templates draw from it.

WHAT IT CONTROLS

The starting state of the job when it is created.

WHY IT MATTERS

Enters the right stage; can auto-reassign on update.

BEST PRACTICE

- ▶ Select the **status** that reflects where this job type enters your workflow - usually your practice's equivalent of "not started".
- ▶ Design your **status** list with automation in mind - a well-structured progression creates a workflow that moves itself forward.
- ▶ Choose a colour scheme for your **statuses**. Start with white for "Not started", then deepen the shade through each in-progress stage toward completion - you might want to reserve reds for anything urgent or sent back for rework.

THE REAL POWER OF STATUS

When a **status** is updated on a job, it can trigger automatic team-member reassignment - the right person is notified and the job lands in their queue at the right moment, with no manual intervention.



Plan your capacity

The **estimated time** fields - hours and minutes - are your capacity-planning tool. They feed directly into XBert's workload visibility, letting managers see whether the team can take on more, and where bottlenecks are building.

WHAT IT CONTROLS

The expected hours and minutes to complete the job.

WHY IT MATTERS

Supports capacity planning and workload visibility.

BEST PRACTICE

- ▶ Set a realistic estimate based on actual experience, not an aspirational one - optimistic estimates undermine your capacity planning.
- ▶ Review estimates periodically, especially for frequently-used templates. As your team gets more efficient, the estimates should reflect that.
- ▶ Where time varies at a client level, you don't need separate templates - refine the figure on the client schedule and keep a representative baseline on the template.

FOR NEW JOB TYPES

Set an initial estimate and flag it for review after the first few jobs are completed - then lock in the figure your real data supports. You can't always estimate an average time for some templates - that's OK, as these are refined further when you apply schedules. Don't leave it blank though: put an estimate in, as something is better than nothing.

ALSO FEEDS YOUR WORKFLOW REVIEW

Estimates are also what you can ask your AI assistant measures - budget vs actual against - see Guide Four. An optimistic estimate produces a misleading capacity view and a misleading review finding. Keep it realistic.



The exception, not the rule

The **step number** sequences a template within a larger, multi-template **process** - where more than one template is used to cover a single end-to-end workflow.

WHAT IT CONTROLS

Sequences a template within a multi-template **process**.

WHY IT MATTERS

Used only when one **process** is split - generally avoided.

BEST PRACTICE

- ▶ In most cases you should not need **step numbers** at all. A well-constructed template captures the complete job from start to finish, on a single template.
- ▶ Only use them when a **process** is genuinely too large or structurally distinct to sit on one template, and the sequence matters for management or reporting.
- ▶ If you find yourself regularly needing **step numbers**, treat it as a signal to simplify your template structure.

DEFAULT POSITION

When building a new template, start from the assumption that everything belongs on one. Only reach for **step numbers** if there's a clear structural reason a single template genuinely can't accommodate.

Schedule or ad hoc

The **group name** classifies how a template is intended to be used - making it immediately clear whether it's for planned, recurring work or for on-demand standardised delivery.

WHAT IT CONTROLS

How the template is intended to be used.

WHY IT MATTERS

Distinguishes recurring (Schedule) from on-demand (ad hoc).

BEST PRACTICE

- ▼ Use Schedule for templates that create work items on a recurring basis.
 - ▼ Use ad hoc for templates that create a work item on demand.
 - ▼ Apply the **group name** consistently - it helps the team grasp a template's context at a glance, and distinguishes committed recurring work from variable demand.
-

WHY THIS MATTERS FOR AD HOC WORK

Tagging a template ad hoc reinforces the core principle of this guide - templates aren't just for recurring work. It signals that this template exists precisely to standardise work that doesn't follow a predictable pattern.



Resources at the point of work

The **attachments** field lets you attach reference documents or working papers directly to the template - so they're available to whoever is working the job, without anyone needing to search for them.

WHAT IT CONTROLS

Supporting documents, workpaper templates, or reference files.

WHY IT MATTERS

Puts the right resources at the point of work.

BEST PRACTICE

- ▶ Attach documents that are genuinely used every time this job type is performed - not documents that might occasionally be relevant.
 - ▶ Keep **attachments** current. An outdated document creates false confidence and can lead to errors.
 - ▶ For ad hoc job types, **attachments** are particularly valuable - they give a first-time team member exactly the reference material they need.
-

KEEP IT LEAN

A short set of current, always-relevant documents beats a long list of "just in case" files. Review **attachments** whenever the underlying working papers change.

A real example: payroll processing

A strong example of a well-constructed XBERT template - it demonstrates most of the best practices in this guide.

●
●
●

XBert - Template: Process Payroll - With Timesheets - [dueday] - [ScheduleFrequency]

Ready to Begin
Payroll
⊕
📄
🗑️
✕

6
5

Test Equipment Pty Ltd

x Task Created Automatically

Group

T

Step 1

1 Process Payroll - With Timesheets - Thursday - Weekly

OPEN LINK

⚡

Due: 7/16/2026

Assigned to
Mr(s) Owner

🕒 Estimated time
1h 30m
7

XPM Link
📄

Summary

Subtasks

☐ Check for any XBerts and resolve

☐ Approve Timesheets

☐ Reassign to manager for review

☐ BK: Process Pay

☐ MR: Send Payslips to employees

☐ File ABA

☐ File away

☐ Close off task

☐ Email the client

Additional Subtasks
 +

Attachments (Drag or paste files here)
 10

Recurring: Weekly Recurring - Week 3 July
 Email: Generate Email Address
 Created: 6/16/2026, 12:09:06 AM by the system

Fig 1. The Process Payroll - With Timesheets template in XBERT, showing the full sub-task sequence, **note link tag**, **process** and estimated-time fields.

What this template does well

- ▶ **Clear naming with tokens** - the [dueday] and [ScheduleFrequency] tokens name each job for the correct cycle automatically.
- ▶ **Logical sub-task sequencing** - data quality checks, and approvals first, then process pay and payslips, update client then filing.
- ▶ **Appropriate note link tag** - Payroll Dates surfaces client context in the job.
- ▶ **Scoped process category** - mapped to Payroll for clean reporting.
- ▶ **Realistic estimated time** - 40 minutes populated, feeding capacity planning.
- ▶ **Works beyond the cycle** - equally useful ad hoc - a one-off termination pay or off-cycle pay run gets the same rigour.



Best-practice checklist

Run through this before implementing any new template to your team.

✓ Name	Follows your convention? Variant marked? Sorts logically in a list?
✓ Description	Would a new team member know when and why to use this template?
✓ Sub-tasks	One discrete action each? Prefixes at the correct handover points?
✓ Note link tag	Is a tag assigned, and the most specific one for this job type?
✓ Process	Is the category correct, or does an existing one fit instead?
✓ Status	Does the starting status match where this job enters your workflow?
✓ Estimated time	Is the estimate realistic and based on actual experience?
✓ Step number	Genuinely needed, or can it all sit on one template? Default: none.
✓ Group name	Tagged ad hoc or Schedule, matching how it will be used?
✓ Attachments	Are attached documents current and relevant every time?
✓ Ad hoc use	Could this also cover an ad hoc version of this job?
✓ Assistant-ready	Do the name, description and process read clearly cold?

Once your library is built, run a Workflow Review to catch duplicates and stale templates - see Guide Four.

GUIDE TWO

Mastering XBert Schedules

A template defines how a job is done. A schedule makes it happen - connecting that template to your clients and triggering work automatically, for the right person, every time.

-
- 01 What a schedule is
 - 02 The naming convention
 - 03 Assignment, dates & frequency
 - 04 Capacity, clients & time
 - 05 A real example: Payroll
 - 06 Before-you-activate checklist

What a schedule is, and what it does

A schedule sits on a workflow template. It is the mechanism that connects a workflow template to one or more clients and triggers the automatic generation of work items on a recurring basis. Without a schedule, a workflow template is a blueprint that never activates. With a schedule - or several - that blueprint becomes a machine that runs itself.

One workflow template can carry multiple schedules. This is deliberate. It is how a single payroll workflow template handles clients on weekly, fortnightly, and monthly cycles - across different days of the week and different pay fortnights - without duplicating the workflow template. The variation lives in the schedules, not in the workflow templates.

KEY PRINCIPLE

Best practice is to tag recurring workflow templates with Group Name: Schedule - this signals clearly that the workflow template is intended to run on a repeating cycle and will have schedules attached to it. Ad hoc workflow templates are triggered manually and do not use schedules. Tagging your workflow template correctly before building schedules keeps your template library clean and your intent visible to the whole team.

Schedule name - the naming convention

The schedule name is the first field you fill in when building a schedule from scratch - and later, it is the primary identifier in the schedule list on a workflow template. The list view shows a summary row for each schedule (name, frequency, due date, assigned role or user, and number of attached clients), but it is the name that tells you instantly what each schedule covers. With a single payroll workflow template potentially carrying 20 or more schedules, a clear, consistent naming convention is the difference between instant clarity and minutes of opening schedules to work out what's what.

BEST PRACTICE CONVENTION

The confirmed convention is: **Frequency · Cycle identifier (if applicable) · Day · Due date qualifier (if applicable)**

SCHEDULE NAME	WHAT IT TELLS YOU
Weekly - Tuesday	Weekly, runs on Tuesday, due same day
Fortnightly - FN1 - Tuesday	First fortnight of the pay cycle, Tuesday, due same day
Fortnightly - FN1 - Thursday (due 1 day)	First fortnight, Thursday start, due the following day
Fortnightly - FN2 - Monday (due 1 day)	Second fortnight, Monday start, due the following day
Monthly - 15th	Monthly on the 15th, due same day

FN1 / FN2, tokens and pitfalls

THE FN1 / FN2 CONVENTION (FORTNIGHTLY)

FN1 identifies clients whose pay run falls in the first fortnight of the cycle. FN2 identifies those in the second fortnight. This matters because a fortnightly payroll practice can have clients whose pay dates stagger across the month - FN1 and FN2 make it immediately clear which group a schedule covers without opening it.

TIP

The tokens **[ScheduleFrequency]** and **[dueday]** in a workflow template title pull their values from the schedule. A well-named schedule means every work item generated from it is immediately identifiable in the work list - no manual renaming, no guessing which client or period a task relates to.

COMMON MISTAKE

Using vague schedule names like 'Payroll 1' or 'Fortnightly A'. When you have 20+ schedules on a single workflow template, you cannot tell what each one is without opening it. Apply the naming convention from the first schedule you create - retrofitting it later is tedious. The same applies when you ask your assistant which schedules cover fortnightly Thursday payroll - it can only answer from the names you gave them.

Activate schedule

When building a schedule, the option to activate it sits at the bottom of the schedule form. This is where you make the schedule live - switching it on so that XBert begins generating work items according to your configuration.

Once a schedule is saved, it appears in the schedule list on the workflow template. In the list view, each schedule shows its current status - **OFF** or **ON** - as a toggle in the first column. OFF means the schedule exists but is not generating work items. ON means it is active. This toggle is visible and editable from the list without opening the individual schedule.

Importantly, deactivating one schedule does not affect other schedules on the same workflow template - each schedule is controlled independently. And turning a schedule OFF applies to all clients on that schedule simultaneously; it cannot be switched off for just one client.

BEST PRACTICE

- ▼ Build all schedules before activating any of them - get the name, frequency, start date, end date, due date, and assigned role correct first; add clients and set any time overrides as the final steps, then activate.
- ▼ Leave new schedules inactive while building mid-cycle - activate only when the schedule is fully configured and ready.
- ▼ When onboarding a large client base, build schedules in batches and activate progressively rather than activating everything at once.
- ▼ Use the OFF toggle in the list view to deactivate a schedule that is no longer active or no longer has any active clients attached to it - the configuration is retained and can be reactivated at any time.

Assigned role vs assigned user

When you add a schedule, you nominate who the generated work items are assigned to. There are two options: a role, or a specific user. Role assignment is the right default for the vast majority of schedules.

ASSIGNMENT	WHEN TO USE	EXAMPLE
Role (recommended)	All recurring work where the responsible person is determined by their role for that client - the default for the vast majority of schedules	Payroll Preparer, Bookkeeper, Tax Agent
User (rare)	Situations where a specific named individual must always receive the work regardless of which client it relates to - most commonly internal practice work, but occasionally applies to specific client-facing arrangements	Professional development time capture; a practice owner task that must not be delegated

WHY ROLE ASSIGNMENT IS THE RIGHT DEFAULT

When you assign a schedule to a role, XBert reads the role board for each client and routes each generated work item to whoever holds that role for that client. The schedule doesn't need to know who the person is - it knows the role, and the role board does the rest.

The practical consequence: if a team member leaves and is replaced, you update the role board once. Every future work item routes correctly from that point, across every client on every schedule that uses that role. Without role assignment, you would need to update every individual schedule manually.



When user assignment applies

User assignment is appropriate when a specific named person must always receive the work - regardless of which client it relates to. Internal practice jobs such as professional development, compliance reviews, or practice administration are the most common scenarios. Occasionally a client-facing arrangement will require the same - for example, where a principal has a standing commitment to personally handle a particular type of work. The key test is whether the assignment is about the individual, not the role they happen to hold.

THE COMMON SCENARIOS

- ▶ **Internal practice work** - professional development, compliance reviews, practice administration - the most common reason to assign to a user.
- ▶ **Standing commitments** - a principal who personally handles a particular type of work for a client, every time, regardless of role.
- ▶ **The key test** - is the assignment about the individual, not the role they happen to hold? If so, assign to a user.

Start date

The start date is when XBert begins generating work items from this schedule. You always nominate a specific date - the first of the month, the end of the month, the beginning of a particular week, or a specific day of the week. The start date sets the rhythm for everything that follows: every subsequent work item generates from that anchor point according to the chosen frequency.

SETTING	WHEN TO USE
Specific current or future date	The standard approach - nominate the date that reflects when the first work item for this cycle should appear. This could be the first Monday of the coming week, the first of next month, or any other appropriate anchor date.
Past date	Backdating - generates tasks retrospectively from that point. Use deliberately; backdating a weekly schedule by several months will generate a large volume of historical tasks.

NOTE ON COPIED SCHEDULES

When you copy an existing schedule, the start date shown will reflect when the next batch of tasks is due to generate from the source schedule - not today's date. This is expected behaviour. Update the start date to the correct anchor date for the new schedule before saving.

Frequency - and month / day on-off

Frequency sets how often the schedule generates work items. The standard options cover the full range of practice needs: daily, weekly, fortnightly, monthly, quarterly, and annually. Two frequency types have additional configuration options worth understanding in detail.

MONTHLY SCHEDULES - MONTH ON/OFF

When you select monthly frequency, you can nominate which specific months the schedule triggers - not just 'every month'. This unlocks patterns that a simple monthly frequency cannot express:

- ▶ **Bi-monthly:** trigger in alternating months only - useful for clients on a two-month reporting or delivery cycle.
- ▶ **Intra-quarter:** trigger in the first and second months of each quarter but not the last - useful where work is needed in the build-up to a quarter-end but not at the close itself.
- ▶ Any custom combination of months the practice or client requires.

TIP - BI-MONTHLY EXAMPLE

Some clients require a management reporting pack every second month. Rather than building a separate workflow template, set the frequency to monthly and use month on/off to select only the months the report is required. The schedule generates work items in those months only - no manual intervention, no missed cycles.

Day on/off & end date

When you select daily frequency, you can nominate which specific days of the week the schedule triggers. This means a 'daily' schedule does not have to run every calendar day - you can configure it to run on Monday, Wednesday, and Friday only, for example, or on any combination of weekdays that matches the client's actual requirement.

COMMON MISTAKE

Selecting daily frequency and assuming it will only trigger on working days. Configure the day on/off settings explicitly to match the actual days required - or enable non-workday scheduling at the practice level to handle weekend date shifts automatically.

END DATE

The end date controls when XBert stops generating work items from a schedule. For most recurring client work, no end date is correct - the schedule should run indefinitely.

SETTING	WHEN TO USE
No end date	Default for all ongoing recurring work. The schedule generates indefinitely until manually deactivated.
Specific end date	Use when the work requirement ends on a known future date. This could be for a short term project or requirement.
Repeat number	Use for finite work - 'do this exactly 10 times then stop' - without needing to remember to return and deactivate the schedule manually.

Due date - the options

The due date determines when the generated work item must be completed. XBert offers significant flexibility here - it is not limited to a fixed number of days after the start date.

Options include:

- ▶ **Days after start date** - the most common setting for payroll and short-cycle work (0 days = due same day, 1 day = due the following day).
- ▶ **Months after start date** - useful for longer engagements where the due date falls weeks or months after the start.
- ▶ **End of week period** - due X number of days or months after the end of the working week in which the work starts.
- ▶ **End of fortnight period** - due X number of days or months after the end of the fortnight in which the work starts.
- ▶ **End of month period** - due X number of days or months after the end of the month in which the work starts.
- ▶ **End of quarter period** - due X number of days or months after the end of the quarter in which the work starts.
- ▶ **End of year period** - due X number of days or months after the end of the financial or calendar year in which the work starts.

REFLECTED IN THE SCHEDULE NAME

The due date qualifier is also reflected directly in the schedule name. The convention includes it where it is anything other than same-day: **Fortnightly - FN1 - Thursday (due 1 day)** tells anyone reading the schedule list that tasks from this schedule are due the day after the Thursday start date.

Due date - best practice

- ▼ Set the due date to reflect the actual delivery expectation - not an aspirational one.
- ▼ For tax and compliance obligations, consider setting the due date a few days before the actual lodgement or payment deadline - this builds in breathing space for review, unexpected delays, or client-side issues without risking a late submission.
- ▼ For longer engagements such as end-of-year work or annual reviews, the end-of-period options give a more meaningful due date than a fixed day count.
- ▼ Include the due qualifier in the schedule name whenever the due date is anything other than same-day - this keeps the schedule list self-documenting.

TIP - COMPLIANCE DEADLINE BUFFER

If a compliance obligation falls on a fixed regulatory deadline, consider setting the schedule due date a few days earlier. The work item lands on the team's board with a built-in buffer - enough to absorb a client delay, a data query, or a review cycle without pushing into late lodgement territory. The regulatory deadline doesn't move; your internal deadline can.

Create tasks in the future

This setting determines how many weeks or periods ahead XBert pre-generates work items. It is one of the most consequential settings for practice capacity planning - and one of the most commonly overlooked.

When future tasks are generated, they appear on the team overview dashboard and count toward each team member's committed hours for those future periods. A practice manager can see not just this week's load, but weeks of committed work ahead across the whole team.

The minimum recommended setting is **two weeks**. This is because XBert's timeflow board view - which gives a visual overview of work across the team - always displays two weeks in advance. If future tasks are not generated at least two weeks out, the timeflow board will show gaps that don't reflect actual committed work, making capacity planning unreliable.

RECOMMENDED SETTINGS

- ▼ **Weekly schedules:** 4–6 weeks. XBert automatically bulk-creates the next batch as the end of the generation window approaches.
 - ▼ **Fortnightly schedules:** 4–6 periods is a practical guide - adjust to the planning horizon your practice works to.
 - ▼ **Monthly and quarterly schedules:** 3–6 months gives meaningful visibility without generating too far ahead for work that may be subject to client changes.
-

WHY THIS MATTERS FOR CAPACITY PLANNING

Without future task generation set to at least two weeks, the timeflow board view cannot show you an accurate picture of upcoming committed work. Generating tasks ahead is what makes XBert's workload visibility genuinely predictive - you can see a bottleneck forming and act before it becomes a problem, rather than managing it after the fact. It is also what gives your assistant something to answer when you ask what is on the team's worklist for the coming fortnight.

Clients

Clients are added to a schedule after the rest of the configuration is in place - frequency, start date, due date, end date, and assigned role should all be set first. Adding clients is one of the last steps when building a schedule, because it is only once you have the schedule structure right that you know which clients belong on it.

Each work item generated by the schedule is created for each client on the list - so a weekly payroll schedule with eight clients generates eight work items every week, each assigned to the relevant team member via the role board. Once a schedule is active, clients can be added or removed at any time without rebuilding the schedule.

BEST PRACTICE

- ▶ When a new client joins: find the existing schedule that matches their frequency and day, and add them to it. Do not create a new schedule for one client if a matching schedule already exists.
- ▶ Group clients by their actual pay day and cycle - that is what drives which schedule they belong to.
- ▶ A client can appear on multiple schedules across the same or different workflow templates - for example, on a weekly payroll schedule and a monthly bookkeeping schedule simultaneously.
- ▶ Use filters when selecting or deselecting clients - this allows bulk selection across large client lists rather than working through them one by one.
- ▶ The same filter approach applies when editing **estimated time** in bulk for a group of clients on the same schedule.

COMMON MISTAKE

Creating a new schedule every time a new client is onboarded, even when an existing schedule already covers that frequency and day. Over time this creates unnecessary schedule proliferation. Check existing schedules first - adding a client to an existing schedule takes seconds. You can also ask your assistant to list the existing schedules on the template before you create anything.

Estimated time per client

The workflow template carries a default estimated time - the baseline for how long this type of work takes on average across your client base. Once clients are added to a schedule, you can override this estimate for individual clients where their actual time requirement differs from that average.

This is the mechanism that makes client-level time variation possible without creating separate workflow templates. One workflow template, one schedule, client-specific time overrides where needed.

HOW IT WORKS

- ▶ The default **estimated time** from the workflow template applies to all clients on the schedule automatically.
- ▶ For clients whose time differs, select those clients within the schedule and update their **estimated time** individually - filters can be used to bulk-select clients with the same time requirement.
- ▶ All other clients retain the workflow template default.
- ▶ The override applies every time a work item is generated for that client from this schedule.

Estimated time - best practice

- ▼ Set the workflow template default as a realistic average - not the fastest possible time, and not the slowest client you have.
- ▼ Use schedule-level overrides for clients whose time is genuinely and consistently different - a larger payroll, a more complex ledger, a client with unusual requirements.
- ▼ Make sure your estimates align with the scope of works you have agreed with the client - this is particularly important for fixed-fee clients, where time is the only variable you can control. Accurate estimates are what make the actual vs **estimated time** reporting in XBert meaningful, and that data is the foundation of confident pricing decisions.
- ▼ Review overrides periodically - a client's employee headcount, transaction volume, or service scope changes over time.
- ▼ Think in terms of the frequency: for a weekly payroll schedule, the **estimated time** is per pay run - not per month. Set it accordingly.

COMMON MISTAKE

Using a single optimistic **estimated time** for all clients on a schedule, then finding that the capacity dashboard significantly understates actual load. Realistic estimates - even if they feel uncomfortable - give you the data to make resourcing and pricing decisions with confidence.

Non-workday scheduling

Non-workday scheduling is a practice-level setting - configured in Settings → Features, not on individual schedules - but it directly affects how every schedule in the practice behaves when a task's date falls on a weekend. When enabled, you nominate what happens to date-based tasks that would otherwise land on a non-working day:

- ▼ **Move to the previous working day** - the task is brought forward to the Friday before a weekend.
- ▼ **Move to the next working day** - the task is pushed to the Monday after a weekend.

BEST PRACTICE

- ▼ Enable this setting for any practice running monthly, quarterly, or annual schedules - these are the frequencies most likely to land on weekends.
- ▼ Decide whether your practice convention is to bring work forward or push it back, and apply consistently.
- ▼ For compliance deadlines (tax lodgement, payroll processing), the choice between forward and back has a real consequence - make a deliberate decision rather than leaving it at default.
- ▼ Once set, the rule applies across all schedules automatically - no schedule-by-schedule configuration needed.

NOTE

Non-workday scheduling currently handles weekends only - it does not factor in public holidays. Tasks falling on a public holiday will not be automatically moved, so plan for these manually where public holidays affect key compliance or payroll dates. This setting also does not affect the actual lodgement or payment deadline set by the relevant tax authority or regulatory body - always confirm your internal task dates align with the regulatory deadline.

A real example: payroll processing

Payroll is the clearest demonstration of the schedule system working as intended. A practice with 20 payroll clients will typically have clients on different frequencies, different days of the week, different fortnights of the cycle, and different processing windows before the due date. All of this variation is handled by one workflow template with multiple schedules - not by multiple workflow templates.

Schedule list - payroll processing workflow template								
Payroll - Process Payroll - With Timesheets - [dueday] - [ScheduleFrequency]								
6 Schedules + Add								
Status	Name	Recurs	Due Date	Attachments	Assigned	Clients	Edit	
OFF ☒	Fortnightly - FN1 - Thursday (due 1 day)	Fortnightly	1 day after start	-	Payroll Officer	00		
OFF ☒	Fortnightly - FN1 - Tuesday	Fortnightly	0 days after start	-	Payroll Officer	00		
OFF ☒	Fortnightly - FN2 - Thursday (Due 1 day)	Fortnightly	1 day after start	-	Payroll Officer	00		
OFF ☒	Monthly - 15th	Monthly	0 days after start	-	Payroll Officer	00		
OFF ☒	Weekly - Monday	Weekly	0 days after start	-	Payroll Officer	02		
OFF ☒	Weekly - Thursday	Weekly	0 days after start	-	Payroll Officer	04		

Fig 1. The schedule list for a payroll workflow template - one workflow template, multiple schedules covering weekly, fortnightly FN1, and fortnightly FN2 variations across the practice's payroll client base.

One template, many schedules

A representative schedule list for a payroll workflow template might look like this:

SCHEDULE NAME	FREQUENCY	DUE
Weekly - Monday	Weekly	0 days after start
Weekly - Tuesday	Weekly	0 days after start
Weekly - Thursday	Weekly	0 days after start
Fortnightly - FN1 - Tuesday	Fortnightly	0 days after start
Fortnightly - FN1 - Thursday (due 1 day)	Fortnightly	1 day after start
Fortnightly - FN2 - Monday (due 1 day)	Fortnightly	1 day after start
Monthly - 15th	Monthly	0 days after start

All assigned to **Payroll Preparer**. Every work item generated from any of these schedules will have a title populated automatically by the tokens in the workflow template title - no manual renaming, no ambiguity about which client or period the work relates to.

BUILDING IT EFFICIENTLY

Create the first schedule from scratch, configure it fully, then use the copy schedule function to create the next one. Change the name, the clients, and the start date - everything else is already set. A full set of payroll schedules for a practice can be built in minutes using this approach.

Before-you-activate checklist

Run through this before activating any new schedule. A schedule that generates work before it is correctly configured creates noise on the work board - and fixing it retroactively means bulk-updating tasks that have already been assigned.

✓ Schedule name	Does the name follow the convention - frequency, cycle identifier, day, due qualifier? Will it be immediately readable in a list of 20+ schedules?
✓ Assigned role/user	Is it assigned to a role (default) or a user? If a user, is there a genuine reason this cannot be a role?
✓ Start date	Is the start date a specific nominated date - not a default? Is a copied schedule's start date updated to the correct anchor date for the new schedule?
✓ Frequency	Is the frequency correct? For monthly: are the correct months selected (month on/off)? For daily: are the correct days selected (day on/off)?
✓ End date	Is it set to no end date for ongoing recurring work? If an end date is needed, is it correct?
✓ Due date	Does the due date reflect the actual delivery expectation? Is a buffer built in before any regulatory deadline? Is it reflected in the schedule name where it is anything other than same-day?
✓ Create tasks in future	Is the future task window set to at least two weeks so that the timeflow board view is populated? Is it set appropriately for this frequency?
✓ Clients	Are the correct clients selected? Does an existing schedule already cover this frequency and day - could these clients be added to that instead?
✓ Check first	Have you asked your assistant, or checked the schedule list, whether an existing schedule already covers this frequency and day?
✓ Estimated time	Have client-level time overrides been reviewed now that clients are on the schedule? Are overrides aligned with the agreed scope of works? Are the estimates realistic?
✓ Non-workday scheduling	Is the practice-level non-workday scheduling setting enabled and configured for the correct direction? Have any public holidays been planned for manually?
✓ Activate schedule	Is the schedule fully configured and ready? Activate only when name, frequency, start date, due date, assigned role/user, clients, and estimated time overrides are all correct.

GUIDE THREE

Client Notes, Topics & Note Link Tags

Client notes, note topics and note link tags - set them up right and the correct client context is one click away, at the exact point in the workflow where it's needed.

-
- 01 Why the system exists
 - 02 Note topics
 - 03 Note link tags
 - 04 Client notes & placement
 - 05 A real example: Payroll
 - 06 Best-practice checklist

Context that lives inside the workflow

Ask any practice principal where client information actually lives and the answer is rarely a good one. It's in a spreadsheet one person maintains. It's in an email chain from three months ago. It's in the head of the team member who's been on that client for two years - and completely invisible to everyone else when that person is on leave. Practices report this costing 20–40 minutes per team member, per day, just locating context that should already be at their fingertips.

The deeper problem is that none of those tools - the spreadsheet, the email, the Slack thread - are linked to the actual client record or to the work being done. There's no single source of truth. There's no mechanism for the right information to surface at the right moment. Information exists somewhere; it just isn't where your team is when they need it.

XBert's notes system solves this by keeping client context inside the workflow itself. Notes sit directly against client records, structured by Note Topics and connected to templates via Note Link Tags. When a work item is created for a client, the relevant notes surface automatically - at the overall job level if they apply broadly, or at a specific sub-task if they relate to one step. Your team doesn't have to know which clients have exceptions. The system tells them at the moment it matters.

KEY PRINCIPLE

Client notes don't change your workflow - they provide the client-specific context that travels with it. One template, configured once, handles every client. Notes handle the exceptions. And because notes live inside XBERT rather than scattered across tools, your whole team has access to a single source of truth. It is also the only version of that context your assistant can reach. Ask it to summarise every note you hold on a client and it answers from XBERT. It cannot answer from a spreadsheet on someone's desktop.

Three layers, one source of truth

This also addresses the institutional knowledge problem directly. When client context lives in XBert rather than in one person's memory, anyone on the team can pick up a job and act on it with the same quality as the account manager who's been on that client for years. That's the system working as intended.

40

minutes a day
lost hunting for
client context

3

layers: topics, link
tags, client notes

1

source of truth,
inside the
workflow

Notes also extend beyond workflow delivery. They can capture client intelligence that any team member needs at hand - the right person to call for payroll sign-off, a note that an employee is paid across two entities, software access credentials kept for reference, the name of a client's business partner. If your practice currently stores this kind of detail in spreadsheets, shared documents, or individual team members' heads, client notes are where it belongs.

THE SHAPE OF THE SYSTEM

Three layers work together: **Note Topics** (the top-level filing cabinets), **Note Link Tags** (the bridge that connects notes to a workflow template), and **Client Notes** (where the client-specific detail lives). Set them up in that order.

Your top-level structure

Before you write a single note, set up your Note Topics. Topics are configured in Settings, and they define the high-level categories your notes will sit under. Think of them as the filing cabinets - broad, clearly labelled, and consistent across the practice. Topics should stay simple. Most practices need four to six. Common examples:

• BAS / VAT Agent Services

• Bookkeeping

• Payroll

• EOFY

• General

BEST PRACTICE

- ▼ Keep Note Topics high-level and stable - these shouldn't change often once set.
- ▼ Align topic names with the way your practice thinks about its service lines and the language your team already uses.
- ▼ Match your Note Topics to your **Process** structure in XBert where possible - consistency between the two systems makes the connection meaningful.
- ▼ Resist the urge to create a topic for every client or every minor variation. If in doubt, use General.

COMMON MISTAKE

Don't create a Note Topic for every service variation. Topics are the top level of the hierarchy - keep them broad. The granularity comes in the Note Link Tags.

KEY PRINCIPLE

"Keep your Note Topics really simple - a handful of broad categories like payroll, bookkeeping and general. Then let the Note Link Tags do the detailed work. That's what keeps the whole system easy to navigate as it grows."

Connecting notes to workflow

Note Link Tags are where the notes system stops being informational and starts being operational. A Note Link Tag is the bridge between a client note and the workflow template - it's what causes the right note to surface at the right point in a work item. Tags are configured in Settings → Notes and Collaboration → Note Link Tags. Each tag belongs to a Note Topic and should follow a consistent naming convention: use the Note Topic as the prefix, then add the specific area of the workflow the tag relates to. Tags should progress from broad to specific.

NOTE TOPIC	NOTE LINK TAG	WHAT IT SIGNALS
Payroll	Payroll - Pay Run	Broad - applies to the pay run workflow overall
Payroll	Payroll - Pay Run - Salary Sacrifice	Specific - applies at the sub-task for salary sacrifice
Payroll	Payroll - Termination Pays	Scoped to termination pay work items
Payroll	Payroll - General Processing	Broad - overall payroll processing workflow
Payroll	Payroll - General Processing - Leave	Specific - surfaces at the leave-related sub-task
Payroll	Payroll - Timesheets	Applies specifically to the timesheet approval workflow

NAMING CONVENTION

Note Topic name → Note Link Tag name. The topic is the filing cabinet. The Note Link Tag is the drawer - and the tag name tells the team immediately what area of work this note relates to, and how specific it is.

Building tags well

BEST PRACTICE

- ▶ Always prefix the tag name with its Note Topic - this makes it immediately clear what category of note will appear.
- ▶ Build tags progressively: start broad (topic → service area) then add specific tags as you identify the exact **sub-tasks** where client nuances apply.
- ▶ You don't need to build every possible tag up front. Add tags to templates as and when you discover a client deviation worth capturing.
- ▶ A tag on its own does nothing - it only becomes active when a client note references it. An unused tag is fine; it just means no client has a specific note for that area yet.

HOW IT APPEARS ON A WORK ITEM

When a work item is generated, a **note link tag** widget appears at the relevant point - either alongside the relevant sub-task, or higher up in the work item for a template-level note. If a client note exists for that tag, the widget is active and noticeable; clicking it turns it blue and opens the note. If no note exists for that client, the widget is greyed out. Your team only needs to act on active widgets - and they know exactly what kind of information to expect.

Where the detail lives

Once your Note Topics and Note Link Tags are set up, you can start writing client notes. Notes are added at the client level - they're not part of the template, which is what keeps the template standard. Each note is linked to a client, a Note Topic, and a Note Link Tag. A good client note explains three things clearly:

- ▼ **What's different** for this client compared to the standard workflow.
- ▼ **Why it's different** - the context that helps a team member understand, not just follow.
- ▼ **What the team should do** as a result of that difference.

COMMON MISTAKE

Creating a note but leaving it as a draft, or forgetting to set it to Published. Unpublished notes will not appear on work items - they exist in the system but deliver no value to the team actually doing the work.

Writing good client notes

BEST PRACTICE

- Write notes for things that are true every time this job type runs for this client - not one-off observations. If it's a recurring deviation, it belongs in a client note.
- Give the note a meaningful title. The title is the first thing a team member sees when they click the widget - it should tell them immediately what the note covers, not just label it with the client name or topic. "Salary Sacrifice - J. Smith" is useful; "Payroll Note" is not. This matters more again when your assistant summarises a client's notes - the title is what it leads with.
- Keep notes concise and actionable. This is an instruction, not a history. If the team member needs to act on it, write it as an action.
- Use formatting to aid readability - notes support text formatting including bold, bullet points, and headings. For notes covering multiple points or steps, formatting makes the content significantly easier to scan at the point of work. Notes can be set up as secondary checklists - a great example is listing all bank accounts that need to be reconciled.
- Set notes to Published and Public by default - notes that aren't published won't surface on work items.
- Use Private **status** selectively and deliberately - for sensitive information such as access credentials or information not appropriate for all team members. Always apply your practice's cybersecurity protocols first.
- Notes can include **attachments**. If a client has a unique working paper - a grouping arrangement document, a specific authority form - attach it to the relevant note so it's available at the point of work.
- Notes can be edited at any time. Review them when a client's circumstances change, not just when something goes wrong.
- Private stays private. Your assistant inherits your XBERT permissions, so it sees exactly the notes you see and nothing more.

Placing tags: template vs sub-task

Note Link Tags can be applied in two places on a workflow template: at the overall template level, and at individual sub-task level. Both are valid, both add value - and using them together gives your team the right level of signal at the right point.

T **Template level**

Surfaces when the work item is opened. Use for broad client context that applies to the whole job - grouping arrangements, specific contacts, pre-work considerations.

S **Sub-task level**

Surfaces at that specific step. Use for precise instructions that apply to one action - who receives the authority to lodge, where to file a specific report, a client-specific payment threshold.

A single template can carry both. A broad tag at template level says: there's something to know about this client before you start. A specific tag on a sub-task says: right here, do it differently for this client. Your team doesn't have to hunt for either - the notes surface at the point they're needed.

BEST PRACTICE

When building templates, consider which sub-tasks are most likely to have client-specific variations. Lodgement contacts, payment methods, approval chains, and client-specific file locations are common candidates. Add a note link tag at those sub-tasks now - even if you don't have a client note to add yet. You can populate the notes later as you discover deviations.

A real example: payroll processing

The payroll processing workflow is a useful example because almost every practice has one, and almost every practice has at least a few clients whose payroll deviates from the standard in some way.

Payroll - Payroll Processing - with timesheets [ScheduleFrequency]

Subtasks

Review client specific instructions and action as required

i

Ensure all timesheets are completed and approved

i

Payroll - General Processing - Timesheets

x

Review all leave to confirm it has been processed correctly,

i

Process Payroll

i

Ensure Payroll Processing is peer reviewed prior to sending to client

i

Prepare and send to client for payment

i

Payroll - General Processing - Payments

x

Send payslips to employees

i

File STP

i

Process Payday Superannuation

i

+ Add

Process *

Payroll

Status *

Not Started

Note Link Tag [Learn more.](#)

Payroll - General Processing

x

Fig 1. The **Payroll Processing - with timesheets** template - the full **sub-task** sequence with its **process**, **status** and **note link tag**.

How the notes carry the exceptions

Here is how the notes system works in practice for this template:

WHAT'S BEING CAPTURED	NOTE LINK TAG	PLACED AT
Employee is paid across two related entities - confirm the cost split before the pay run begins	Payroll - Pay Run	Template level
Timesheets must be confirmed approved by the client before the pay run starts	Payroll - Pay Run	Template level
Authority to release payment must go to the business owner, not the operations manager	Payroll - Pay Run - Authority to Pay	Sub-task level
Employee has a salary-sacrifice arrangement - check the super split before finalising	Payroll - Pay Run - Salary Sacrifice	Sub-task level

None of this required a separate template. The Payroll Processing template stays unchanged - it's the same standard workflow used across all clients. The notes carry the exceptions, and they surface where and when they're relevant. When a work item is created for a client with active notes, the team member opening it sees a visible widget at the relevant point; clicking it turns it blue and opens the client note directly. They don't need to know which clients have exceptions - XBert tells them at the moment it matters.

Beyond the workflow: client intelligence

Client notes don't have to be about the mechanics of delivering work. They can capture any intelligence that helps a team member act confidently on behalf of a client:

- ▼ Who the key contacts are and their roles.
- ▼ Relevant entity structures.
- ▼ What software the client uses and who has access.
- ▼ Relationship context useful for client-facing interactions.

ONE SOURCE OF TRUTH

If your practice currently keeps this kind of detail in spreadsheets, email threads, or Slack messages - none of it linked to the actual client record - Notes with a General topic is where it belongs in XBert. One source of truth, inside the system your team is already working in.

Creating notes via email

Notes don't have to be created manually inside XBert. Each XBert organisation has a dedicated inbound email address. Send or forward an email to that address with **#note** included in the subject line, and XBert will create a note from the email content.

The created note will not be automatically assigned to a client - it will need to be manually finalised, including assigning the correct client, Note Topic, and Note Link Tag. Think of it as a fast capture mechanism: the content arrives in XBert immediately, and you complete the linking as part of your regular review routine.

This is particularly useful for capturing client instructions that arrive by email: a specific lodgement preference confirmed over email, a change to payroll arrangements, a new contact who should receive correspondence. Rather than copying and pasting into a note manually, forwarding the original email brings the content straight into XBert with the original communication preserved.

TIP

Remember to include **#note** in the subject line - without it, the email will not create a note. Build the habit of reviewing and finalising email-created notes - assigning the client, Topic, and Note Link Tag - as part of your daily or weekly routine so they connect correctly to your templates.

ALSO VIA YOUR ASSISTANT

Notes can be captured through your assistant too. Tell it what to record against a client and it proposes the note for your approval - nothing is saved until you say so. Finalise it the same way as an email-created note: check the client, Note Topic and Note Link Tag.

Version history and restore

Every time a client note is edited, XBert saves a version of the previous content. The full version history is accessible from within the note, showing who made each change and when. If a note is edited incorrectly - information removed by mistake, instructions overwritten - any previous version can be restored with a single action. Your assistant can read this history too when it summarises a note.

This matters in practice for two reasons. First, it removes the risk of permanent information loss when notes are updated, which makes the team more confident about keeping notes current rather than leaving outdated content in place to avoid losing something. Second, it provides an audit trail for client instructions - if there's ever a question about what was documented and when, the version history is there.

GOOD PRACTICE

Version history works best when notes are updated promptly as client circumstances change - not just when something goes wrong. A note updated frequently with a clean history is far more useful than one left untouched for months because nobody wanted to overwrite what was there.



Best-practice checklist

Run through this before rolling out your notes structure to the team.

✓ Note Topics	Are topics high-level, stable, and aligned with your practice's service lines? Do they reflect the language your team already uses?
✓ Note Link Tags	Does each tag follow the naming convention (Topic prefix → service area → sub-area)? Is the tag name specific enough that a team member knows what kind of note to expect?
✓ Tag placement	Have you considered both template-level tags (broad context) and sub-task-level tags (specific steps)? Are tags placed at the points where client deviations most commonly occur?
✓ Client Notes	Is each note Published and set to the correct visibility (Public or Private)? Does the note explain what's different, why, and what the team should do?
✓ Attachments	If a client has a unique working paper or reference document relevant to this note, has it been attached?
✓ Maintenance	Is there a process for reviewing and updating notes when a client's circumstances change? Outdated notes create the same risk as no notes at all.
✓ Structure first	Have Note Topics and Note Link Tags been fully configured before notes were written? Notes added without the right structure in place are harder to find, manage, and retire.
✓ Reachable	Is this context in a client note, where your team and your assistant can both reach it - or still in a spreadsheet where neither can?

GUIDE FOUR

Ask Your Workflow

Templates, schedules and notes give your work a standard shape. The XBert MCP lets whichever AI assistant your firm already uses, like Claude, ChatGPT and Copilot, read that shape and answer from it - in plain English, under your login.

-
- 01 Why it exists
 - 02 Connected in three steps
 - 03 Questions worth asking
 - 04 The Workflow Review plugin
 - 05 Governance
 - 06 Before-you-connect checklist

Your data and your AI, in the same room

Most practices made the same move in the last year: someone opened Claude or ChatGPT, pasted in a report, and got a useful answer. The problem is what happens next. The paste goes stale the moment it lands, the export is partial, and nobody can trace where the answer came from. So every question that touches real numbers means leaving the assistant, opening XBert or the ledger, exporting, and pasting back - or the question doesn't get asked.

The XBert MCP closes that gap. MCP - the Model Context Protocol - is a secure, standard way for Claude, ChatGPT or any MCP-aware assistant to connect to the software your practice already uses and work with it on your behalf. Point your assistant at XBert and it can read your clients' books, your practice's worklist, your schedules and your client notes, and answer the way a senior would. It signs in as you, sees what you see, and proposes changes rather than making them.

KEY PRINCIPLE

AI is only as good as the data it reads. Everything in Guides One to Three - consistent template names, self-documenting schedules, published client notes - is what makes your workflow readable to your assistant. Good design is what turns a plain-English question into a reliable answer.

Ensuring all your clients are connected on a data quality plan ensures that you get the full and trusted picture across your practice. Do not trust answers on financials of clients unless they have been connected.

Connected in three steps

There is no code, no integration project and nothing to install in XBert. The whole setup takes about five minutes.

01 Point your assistant at XBert

In your assistant's connector settings, add one address: **<https://mcp-gateway.xbert.io/mcp>**. That is the whole technical bit.

02 Sign in to approve

A browser window opens and you sign in with your XBert login. The assistant gets your permissions - no more, no less - and your credentials stay with XBert.

03 Start asking

Type the question the way you would ask a colleague.

TIP

Locked-down workspace? Some firms' Claude or ChatGPT workspaces block new connectors by default. If the address will not add, your workspace admin needs to allow it. This is the most common hard stop, so check it first.

XBERT SUPPORT

How to add XBert as a custom MCP connector in Claude



Questions worth asking

The examples below map to the three guides you have just read. Any question that would change data is proposed for your approval - you always get the final say.

AREA	ASK IT	PROPOSED - YOU APPROVE
Templates and worklist	"What's on my worklist today, across every client?"	-
Templates and worklist	"Create a month-end task for [client] from our template and assign it to [name]."	Yes
Templates and worklist	"Snooze everything on [client] until Monday."	Yes
Schedules	"Which schedules on our payroll template run fortnightly, and which clients are on each?"	-
Client notes	"Summarise every note we hold on [client], newest first."	-
Client notes	"Find where we discussed the vehicle purchase with the client."	-
Alerts	"What fired across my clients this week? Group it by type."	-
Alerts	"For this duplicate-payment alert, what exactly needs fixing, and how?"	-
Capacity (firms on Xero Practice Manager)	"Who has capacity next week, and who's already over?"	-

TIP

Ask in the order a senior would: find the client or job, get the totals first, then drill into the detail, then act. Numbers first, detail second.

Ask one Question - "How well have I setup up my XBert workflow?"

The checklists in Guides One to Three are how you build a clean workflow. When using your own AI Assistant, it can pull your clean workflow configuration from XBert, runs a fixed set of analyses, and hands back information, in a Word document if you like.

WHAT IT COVERS

- ▼ Template consolidation - duplicate and near-duplicate templates worth merging.
- ▼ Budget accuracy - estimated vs actual time, ranked by capacity impact.
- ▼ Workload balance - concentration risk across the team.
- ▼ Setup hygiene - orphaned assignments and stale templates.

Every finding carries a confidence label - Direct, Likely or Needs review - so you know how much weight to put on it. The calculations are fixed rules, not AI judgement.

WHEN TO RUN IT

Once, after you have built your template library and activated your schedules. Then monthly. Treat it the way you treat a period-end review: a fixed rhythm, not a reaction to something going wrong.

WHAT YOU NEED

An active XBert account with configured workflows and at least one client on a schedule. Claude Code or Claude Desktop.

You stay the reviewer

Connecting an assistant to client data is a governance decision, not a technical one. Three principles hold it in place.

1 Your role, exactly

The assistant signs in as you. It sees exactly what your XBert role allows, client by client, and nothing more. There is no firm-wide super-user mode.

2 Proposed, not posted

Anything that would change data is validated, checked for duplicates and put to you for approval before it happens. The assistant cannot post to a ledger, create a task or save a note on its own.

3 Audit trail on every action

Every question and every proposed change is logged: who asked, against which client, and what happened. The assistant and the XBert app write to the same audit log.

BEST PRACTICE

- ▶ Start read-only. Spend the first week asking questions and approving nothing. Most firms stay there for a month.
- ▶ Decide who connects. The assistant only ever holds the permissions of the person signed in, so review your role board before rolling out.
- ▶ Brief the team on proposed, not posted - so nobody waits for a change that needs their approval.
- ▶ Disconnecting is one click in your assistant's settings.

Before-you-connect checklist

Run through this before pointing your assistant at XBert. Most of it is the work Guides One to Three already asked you to do.

✓ Template names	Consistent, plain-language, following your convention? The assistant finds templates by name.
✓ Schedule names	Frequency, cycle, day and due qualifier in every name? Vague names are unreadable to the assistant too.
✓ Client notes	Published, correctly tagged, with meaningful titles? Unpublished notes do not surface for the team or the assistant.
✓ Estimated time	Realistic on templates and overridden where clients differ? This is what the Workflow Review measures against.
✓ Role board	Current? The assistant inherits each person's permissions from it.
✓ Workspace	Does your Claude or ChatGPT workspace allow new connectors? Check before the rollout, not during it.
✓ Team briefing	Does everyone know that changes are proposed, not posted, and that they approve them?
✓ First month	Agreed to start read-only and review who used it before approving any writes?

See it working on your practice

Templates, schedules and notes give your practice one standard. The XBert MCP lets whichever AI you use work from it. Book a 30-minute session and our team will walk you through a real XBert workflow - then show you what it looks like to ask your own assistant about it.

[Book a meeting →](#)[Start Free Trial →](#)

Already on XBert? Connect your assistant
at intelligence.xbert.io