LTB · PRIVATE REPORT

LTB Improvements

What I found, and a plan to make more books with less of your time

Nothing has been changed. This was a look-only review. Every step in the plan waits for your yes, one point at a time.

What was looked at

- Every order chat since 17 October, read from the first line to the last: 86 chats, about 10,000 Claude steps.
- The cost records of 170 orders, added up line by line.
- The Pi and this PC, measured.
- Your own typed messages (366) and button presses (327).

Then the draft plan was given to separate reviewers whose job was to knock it down: against your rules, against the real records, and for what is missing. Their corrections are in this version.

18 November 2026 · about a 15-minute read · numbers are from the records unless marked “estimate”.

1. The verdict first

Is this worth doing? Yes. But be clear what it buys.

- **It will not sell one more book.** Making a book is already quick and cheap. The limit on books per day is orders being paid for, then your own day.
- **It buys back your time:** fewer typed yeses, one place to look, less waiting while a chat works out how to do a job it has done before.
- **It makes slips with families less likely** (section 6).
- **It should cut Claude use on orders by roughly 40 to 50%** (estimate until tried).
- **The cheaper plan** would save £61.10 a month (£79.10 against £18.00, from your own bills). Honest odds: not before the end of January (section 11).

There is no rush. You are not short of Claude allowance today. So nothing risky needs building in the Christmas weeks. Before about 15 December I suggest wording changes and look-only tools only.

One thing is urgent, and it is not tokens. The Pi's memory card is filling up: 4.8 GB free today, 14 GB on 18 September. At this week's pace it is full between about 3 and 13 December; at the slower monthly pace, mid-December. Either way that is inside the Christmas orders. It is the first decision I need from you (section 12).

2. Your questions, short answers

Faster repeated jobs (pages live, emails)?
They are slow because each chat works the job out again. Putting changed pages live took up to 45 steps. A ready tool that does it in one go exists and was used 0 times in 85 chats, because chats are never told about it. Fix: a rule-and-tool card now, one-command tools in December, skills on top.

Token use?
Three quarters of last week's Claude use was the order system. A quarter of everything was the "postman" chat that only passes news along. Most of the rest is long chats re-reading themselves (section 5).

Skills, so Opus just knows?
Yes. Not one order chat has ever used a skill. But skills are about a third of the answer (section 9).

Downgrade to Pro?
Not today. In the first 30 hours of this week you used 13% of the Max week. On Pro that is 65% of the whole week gone by Wednesday breakfast. Decide in January (section 11).

You handle everything, I check and decide?
Possible, and you have already started it (handing one family's whole conversation to its chat on 18 November). The missing piece: your decisions should reach you on one list, with the picture or the exact words, one tap each. Today they are spread over a chat per order (section 10).

Move from the Pi to this PC? More memory?
Stay on the Pi. Buy no memory. Do not move. Protect its memory card (section 7).

3. The trend, the money, the clock

Week starting	Orders in	Paid
20 Oct	62	7
27 Oct	28	4
3 Nov	23	8
10 Nov	41	7
17 Nov (one day)	7	7

Seven payments on 17 November alone. In 45 days: 170 orders, 37 paid so far.

£1.55 to paint a normal paid book (middle value). From the cost logs; the bank may have paid about 20% more (looks like VAT, not checked yet).

£2.39 of free-sample painting behind every paid order. That is the cost of winning the order, not a leak.

£27.26 direct profit on an average book (£39.99 hardback; print and post £9.37 on the mainland).

6.5 min to paint the samples; 26 minutes to paint a full book. Samples reach a family about 19 minutes after they order (middle value); overnight ones wait for your release.

2.7 hours from payment to the print order (middle value).

9.2 days from print order to the door (18 books delivered so far). The print firm and the post are over 90% of the wait.

> **What slows a book is change rounds, not the machine.** 27 of 37 paid books were changed after the first full paint. With no changes a book went to print 1.7 hours after payment; with any change, 13.1 hours.

What this means: the workshop is not what limits books per day. If orders tripled, the first things to break would be your day and the change rounds.

4. Where your time goes

327 button presses in 45 days. About 7 a day, every day of the week. 163 were sample releases, 125 of them for orders that have not paid.

366 typed messages in a month, in 86 chats.

1 in 3 of what you type in an order chat is a plain yes. Nearly half of those are “Send it”.

3 yeses for one fix, often: yes fix it, yes put it in, yes release. This happened 18 times.

2 min of waiting after a typed yes (middle value), because the yes wakes a very long chat that re-reads itself.

10 chats typed in on 17 November, switching chat 21 times. Things waiting for you sit in four places: site emails, desk cards, a chat per order, texts.

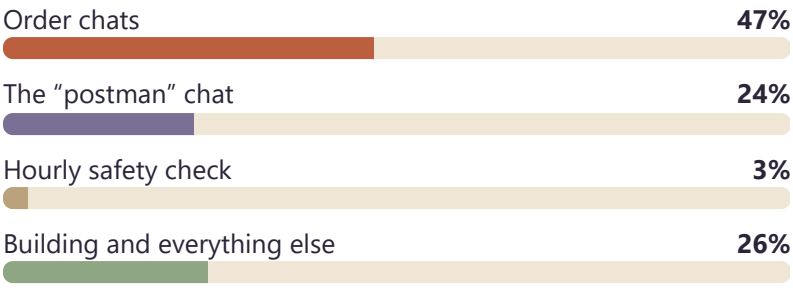
104 of your 294 order-chat messages went to orders that had not paid. One unpaid order took 14 messages and three free rounds, and did not buy.

Your eye stays. 31 of your messages were picture faults or painting direction that Claude had not caught: a seat painted as clothing, two grandads painted as one, pages that looked like photos. You also changed 1 in 5 drafts before sending. Claude caught things first too (a three-eyed bear, a brand name on a cover). Both looks are needed. The plan removes no check. It makes each check one tap with the evidence in front of you.

A gap. You approve on the review page, but Claude's warnings are written in the chat. At least six times the warning arrived after your Approve. One of those became a change round and four emails.

5. Where Claude use goes

Last full week, 10 to 16 November. Shares of all Claude use on this PC.



The four leaks

1. The postman. It only passes news to order chats. But every wake re-reads its whole long conversation, six steps per piece of news. 85% of its wakes are pure routine. Design talk typed into the same chat makes every later wake dearer.

2. Long chats. Every step re-reads the whole chat. Two thirds of all re-reading happens after a chat has grown past 200,000 tokens (over half even without the one giant chat). That one was a bespoke order in October: four days in one chat, grown to 966,000.

3. Working it out again. Chats read the big notes file and the tools' own code to find out "how do I…". They wrote 623 helper files from scratch. A helper to reach the Pi was written again in 9 chats.

4. Waking for nothing. Over half of the news messages were answered "noted, nothing needs you". 34 of 51 orders were opened twice. 137 of 148 hourly checks said "all fine".

A paid order that reaches printing takes about 145 Claude steps, 25 to 45 minutes of Claude session time and 4 typed messages from you (middle values of 10 orders).

6. Slips that reached families

Found while reading the chats. No names here. Most are already fixed one by one; the list shows the pattern.

- An owner-only link went into a family's email. The family started a 4-page repaint that nobody had checked.

- Claude pressed the family's Approve for them, right after a long chat had been shortened and the rule fell out of it.
- A family waited 18 hours because a home-made watcher froze.
- An email asked a family for a needed-by date.
- A typo was flagged five times in a chat, and printed anyway. The print was cancelled and ordered again.
- A name went to print in lower case after being raised three times in a chat.
- An email left from the wrong address; a mail connector re-attached a family's own photos.
- A "you have not ordered yet" text went out two seconds after a family pressed Buy.

The pattern. These are not carelessness in the moment. They happen when a rule lives in a chat's memory, or in a note the chat has to go and find, instead of on a short card and inside the tool that does the job. A tool can refuse to send an owner link to a family. A chat can only try to remember.

7. Pi or PC? More memory?

Verdict: stay on the Pi. Buy no memory. Do not move the workshop, least of all before Christmas.

1.4 of 3.8 GB memory in use on the Pi. It has not run out once in 154 days.

95.6% of the time since June the Pi's processor has been idle.

70 sec of a finished book is the Pi's own work. The rest is waiting for the outside painting services to answer.

45 sec a book is all a move to this PC would save (it is 2.9 times faster at that one minute of work).

What a move would cost: 42 tools and about 30 timed jobs rewritten for Windows, in the peak weeks. And the PC is the less steady machine: Windows restarted it by itself on two evenings (12 August, 12 September), it stopped unexpectedly once (8 August), and it has about 6.5 of 32 GB memory free with all the chats open. It has run without a break since 17 October only. The workshop was moved off this PC in July for these reasons.

The real risk is the memory card, not the memory. 4.8 GB free, and the desk's own alarm is 0.3 GB away. One background job also rewrites a 100 MB file onto it every 10 minutes, which wears a card out. The booking texts for your other client run from the same card.

What really slows painting: outside calls that seem to go silent for 3 to 4 minutes before the painter gives up and tries again. About 180 to 210 times since 2 October, some 600 to 700 minutes, a third or more of live calling time. This is read from gaps in the cost log, so it is an inference, and the cause is not proven. Measure first, change after Christmas.

8. The plan

How it runs. One point at a time. You accept a point, then it is built from a written brief by an Opus build session (desk jobs by the desk's own session), tried on a pretend order, with an off switch. You see the result before it is used.

A Now: your decisions and wording. No code.

A1. Give the Pi's memory card room
Clear three download and browser caches (2.2 GB). Nothing installed is removed. On 18 September you chose to leave them, with 14 GB free then. This asks you to change that.

Me: 10 minutes, plus a daily free-space alert. You: one yes. Old test work (1.5 GB) only after it is checked for copies of family photos.

A2. A one-page card for order chats
The rules an order chat really meets, plus the list of ready tools. It replaces "read the top of the big notes file", which today is mostly other orders' news; the rules the brief points to are not there. A chat still searches the notes for its own order. The card is kept as a file, so it is still there after a long chat is shortened. You see its exact words first.

Me: about 2 hours. You: 10 minutes to read it.

A3. Three working rules

One question per fix, when the picture and the exact words for the family are in the same message. Say nothing when nothing needs you. A permanent fix becomes a numbered brief; it is not built inside an order chat. Your one-free-round rule for unpaid orders goes on the card with its agreed reply.

Me: 30 minutes of wording. A paid painting still needs your go first. One page first stays.

A4. Your daylight test of the desk buttons

Parked since 14 October. Later steps lean on those buttons.

You: about 30 minutes on a Friday or Saturday at 10:00, on your test order.

A5. One invoice from each painting service

Top-ups are 20% or more above the logged spend. It may be VAT, or credit not used yet. If it is VAT, every painting figure is a fifth too low.

You: 10 minutes.

A6. Renew the Pi's own Claude sign-in

It runs out on Monday 7 December at 11:17, in the busiest week. It is behind the "Ask Claude" and "Polish" buttons.

You: 5 minutes on your phone, any day before then.

B Small builds, look-only tools first

B1 and B2 cannot touch a family, so they can be built now. My default for the rest: after 15 December. If you want the saving sooner, B3 to B5 can go in during the week of 23 November, one a day. From 1 to 15 December nothing is changed in any case.

B1. One permanent helper to reach the Pi

Every chat writes its own today. This also removes the Pi login that sits written inside a tool file and in 63 old scratch files. After that, you change the Pi password.

Build: about 1 hour. You: 15 minutes for the password. Removes about a third of all tool errors.

B2. Five look-only commands

Status of an order. Photos and pages on one sheet. What the family's page shows. Real cost. Wait for a job.

Build: 4 to 6 hours. Catching up on an order drops from about 18 steps to about 5.

B3. A thin postman

Your "type here" chat stays yours: you type there as now and it keeps filing new orders. The news-passing moves to a new small chat that you open once and never type in. Fewer steps per wake, no task-list ceremony.

Build: 2 to 3 hours. You: 15 minutes to open it. About 70% less per wake (estimate).

B4. Your list of what wakes a chat

At once: a family's words, a payment, your check needed, any problem. Print and courier steps, chat renames, and echoes of your own taps wait for the next real news. The desk still shows them at once. You choose the list.

Build: 2 to 3 hours.

B5. Open each order chat once, cheaply

Today most orders get a second chat that reads the first one and the notes again: 13 steps against 3.

Build: 1 hour.

B6. Keep order chats small
A chat shortens itself at 200,000 tokens and keeps a six-line order card: what you approved, what was sent, what is open. Order chats get their own folder so this does not touch your other chats. Tried on a pretend order first.
Build: 1 to 1.5 hours, then one order watched for a week.

B7. Five short skills
Check samples. Reply to a family. Fix a page. Print check. Order wrap-up (section 9).
Build: 4 to 6 hours. You: 20 minutes to read them.

B8. "Fix to the side" and "install" tools with guard rails
The page that goes live must be the very file you saw. The tool refuses a book the family has approved, and never changes a status that emails a family.
Build: a day to a day and a half. A page fix drops from 40 to 90 steps to about 5 for a known kind of fix.

B9. Drafts also on the desk, with your Send button
The draft still appears in the chat as now. The desk card is extra. Your tap under those exact words is the yes. A chat can place a draft there but can never send.
Build: 3 to 4 hours by the desk's own session, after your daylight test.

B10. Closed books keep their real cost
Today the desk shows an estimate once a book is closed (one book shows £5.67; the logs say £9.67).
Build: 1.5 hours, with a test on a pretend order.

C January

C1. One “waiting for you” list on the desk
It grows out of the “Needs you” cards you have. One card per decision, with the picture or the exact words. Your button says “Use it”, never “Approve”.

C2. A note under a page starts a cheap, small edit
On your review page only. Made to the side, shown before and after, put in on your tap.

C3. A fix agreed at the samples carries into the paid book
Kept as a picture or a measured colour. Never as words about the child.

C4. The Pi from memory card to SSD
A plan with a real price and the downtime first. Your other client’s booking texts and your wife’s shop key live on the same Pi.

C5. The plan decision
Friday 15 or Saturday 16 January, 10:00, on four weeks of numbers.

Parked Not recommended now

- Painting two or three orders at once. The machine is not what holds books up.
- Shorter time limits on silent calls. Measure first; a shorter limit can mean paying twice.
- The hourly safety check every 3 hours. It saves about half a percent, and a family’s reply could wait 3 hours after an app restart.
- Moving to the PC. More memory.

What I expect (estimates until tried). Order-system Claude use down by roughly 40 to 50% at the same number of orders. Your typing on an order down from about 5 messages to 1 to 3 taps once the list exists. The honest test: compare the next 5 orders after each step with the 5 before.

9. Skills: yes, and what else

Your idea is right. Opus should not have to find out how each time. The numbers back you: no order chat has ever used a skill, and chats spent about 1,100 commands reading notes and tool code to learn the job.

But a skill only tells Opus how. Opus still has to do it, step by step. The savings stack up in this order:

1. Do not wake Claude when no person is needed. A skipped wake saves everything.

2. Keep chats small. Every step re-reads the whole chat.

3. One-command tools. The doing. One step where there were forty.

4. Skills. The how-to for a repeated job, loaded only when that job comes up.

The five skills. Check samples. Reply to a family. Fix a page. Print check. Order wrap-up. Each one says: when to use it, the exact commands, the rules of that job, what to show you, and where to stop.

What must not go into a skill: the rules that hold always (suggest, never act; the family presses Approve; the family's link, never yours). Those go on the card every order chat starts with. A skill only helps when it is triggered, and when a long chat is shortened only the first part of a used skill is kept.

This builds on the bespoke kit of 20 October: 11 ready helpers that today only bespoke chats know about. One small find: your "painting research" skill has a broken first line, so an ordinary request cannot trigger it. Two minutes to fix.

10. The way of working

Today. One long chat per order, a postman chat, and your buttons on the site. You find the right chat, read, type yes, wait.

Target. Claude prepares everything. Your decisions arrive on one list, with the picture or the exact words, one tap each. Nothing from a family is held back. Decisions that can wait are parked for Friday or Saturday 10:00.

Claude alone, already today: read everything, pre-check, one sheet per order, write drafts, work out what a suggestion would cost.

New standing permissions. Each needs your yes, each has a limit:

- A free fix made to the side and shown to you. Never installed.
- A paid fix to the side, up to a pence limit you set. My default: no.
- Suggest a paint hold with one tap when a question is open. A held paid order always texts you.
- Notes and cost lines, only for work you approved.

Always yours: every message to a family, releasing samples or a book, the print button, any money (refund, discount, Express), any likeness judgement, deletes, demo art. Unless you hand one named order to its chat, as you did on 18 November.

Never Claude: the family's Approve.

Why a list and not a chat: a typed yes wakes a chat of about 270,000 tokens. A tap on a card that runs a ready command wakes nothing, or a small fresh helper.

One idea I left out on purpose. Every reader of the chats suggested the same thing: do not open a chat for an order until something needs a person, and use one chat per family, not per order. 22 of 52 orders never needed a typed word from you, and one family's four books meant 8 chats. It would be the biggest single saving. But on 12 November you asked for a chat per order that already knows the order when you open it. So this is your call, not mine (section 12).

11. Pro or Max?

Not today. Pro is one fifth of your Max plan, on both the 5-hour meter and the weekly meter. What the records say:

13%	of the Max week used in its first 30 hours this week. On Pro: 65% of the week.
17 of 53	five-hour stretches in the last two weeks went over what Pro allows in five hours (estimate). Order work alone: 5 of 37.
about 4	Pro five-hour allowances were used in one stretch on the evening of 17 November (seven payments). Order work alone: more than two.
1.7 to 1.9	Pro weeks of use last week (estimate). The order system alone: 1.2 to 1.4.

Also on Pro, the top model (the one that did this review) is not included.

After the plan? If order work really drops by 40 to 50%, order handling alone would still take well over half of a Pro week at today's 40 orders a week, before any building. In October most of your use was building, not orders.

My advice. Keep Max through Christmas and January. Measure four weeks after the changes. Switch only if the Max meters stay under about 14%, for a whole week and on the busiest evening. You already switch plans by season; this would tell you when it is safe. Decision: Friday 15 or Saturday 16 January, 10:00.

How sure: the meter readings are exact. The size of the allowance in tokens is my estimate (a Max week is about 500 to 600 million of my

units; good to about 15%). The app labels your plan “Pro” while the account record says Max 5x; the meters fit Max.

What this review cost. It took the weekly meter from 13% to 27% and used one whole 5-hour allowance, which stopped the work for an hour (the top model counts about double). At 11:35 both meters dropped back, the weekly one to 14%. I cannot see why in the records. On Pro this review could not have been run.

12. Your decisions, one at a time

I will bring these to you one by one. **Only the first needs you now.**

1. The Pi’s memory card. May I clear the three download and browser caches (2.2 GB) and add a daily free-space alert? Nothing installed is removed. This reverses your 18 September choice, made when 14 GB was free.

2. May order chats start from a one-page card instead of reading the top of the big notes file? (You see the card first.)
3. One question per fix, when picture and exact words are in the same message?
4. Build the two look-only tools now (B1, B2)? Then you change the Pi password.
5. What should wake an order chat at once, and what can wait?
6. Watcher changes (B3 to B5): week of 23 November, or after 15 December? My default: after.

7. What may a chat do without asking: a free fix to the side? A paid one up to a few pence?
8. Should Claude's notes on a picture show where you approve? As "things to look at", never as a pass mark. On 9 October you said the check is someone intelligent looking at it.
9. Drafts on the desk with your Send button, as well as in the chat?
10. A chat per order from the start, as you asked on 12 November? Or a chat only when something needs a person, and one chat per family?
11. Should each order keep its own short notes, with only lessons going into the big notes file? My default: no change for now.

How this was checked

14 readers and auditors read the chats and the records. Then six reviewers checked the draft: three re-did the numbers from the raw records, three tested it against your rules, against how the system really works, and for what was missing. Their corrections are in this version. Not verified: exactly how the plan meters weigh different kinds of use; whether the gap on painting spend is VAT; the Pi-against-PC speed test was not run a second time. All savings are estimates until the next orders are measured. Not looked at: building, marketing, Etsy and new-book making, which were most of October's Claude use.

Found in passing: the Pi login is written in one tool file and in old scratch files, and it showed in two reviewers' working output on this PC today (nothing was sent anywhere). B1 fixes the file; then the password should be changed. 859 backup copies of the notes file take 1.4 GB in the project folder. One reviewer left a 3.5 KB temporary file in the Pi's memory-only area; it is harmless and goes at the next restart.