

What Leading UK VARs and MSPs Still Handle Manually – and Where Off-the-Shelf Tools Fall Short

Six questions to score your own operation against the market



Where the Manual Work Concentrates

The manual processes in a service operation are easy to name. Which ones can be changed (and what changing them would actually cost) is the harder question, and most have been reviewed at least once.

The obvious automation is long since delivered: ticketing, monitoring, provisioning, alerting, billing runs. What remains is still manual for good reasons, not out of negligence.


Providers manage a range of these processes. Two seams recur often enough across the market to be worth examining on their own, and they sit at opposite ends of the same pipeline.

01

Information arrives in a shape the system cannot read.

A vendor's rebate statement as a PDF laid out for a human reader.


Or a supplier confirmation with the detail in the body of the email. The data is present and correct; it is not in a form anything downstream can take, so somebody puts it there.



02

Information leaves in a shape the client cannot use.

A figure on a dashboard that is accurate and unreadable without context. A report that establishes what happened without addressing what it means or what should follow. Again, the data is present and correct; the client cannot act on it, so somebody explains it.



Neither is a capability gap, and neither has been neglected. Both have been addressed, and the market's answers are consistent enough to lay out plainly, which is what the next page does.

How the Market Covers Both Ends Today

○ What Arrives

Vendors and suppliers send information in whatever form suits them: an attached PDF laid out for a human reader, a table pasted into the body of an email, a portal export that handles one account at a time, or a scan of something printed. Those are examples rather than a complete set — the format is chosen at the other end, and there are as many formats as there are senders.

The first cost is hours.

None of this is hard to read; it is simply read by a person, every time. The work does not accumulate into anything — last month's effort makes this month's no shorter — and the hours come from people employed to do something else: a finance analyst at month end, a service co-ordinator, a delivery lead. It is a recurring operational cost.

The second cost is that formats do not stay the same, and that is why the first one persists.

A vendor who redesigns their statement has broken nothing of their own; the breakage happens downstream, and it tends to surface later as a figure that does not reconcile. It is also why a one-off fix does not hold: automate against today's layouts and the hours return as the layouts change.

○ What Leaves

Four answers recur across the providers reviewed. They are not the only ones, and the first is the most interesting thing in this paper.

01 The explanation is sold as a service.

Not absorbed, not automated — priced. Among the providers reviewed, a few sell interpretation as a named line: specialist help to get more out of your reports, a bookable coordinator session, a quarterly service review, a written recommendation that accompanies the data. This is worth pausing on. It is not a theory that clients would value explanation — it is a service they are already paying for.

02 Reporting by exception, outside the system.

A person decides each period what matters and sends it. The portal holds everything; the human sends the part that counts. It works, and it scales exactly as far as that person's week.

03 Another login.

Where a gap appears, the answer is frequently one more destination. The phrase one way in is common in the market's own copy and rarely true underneath it: among the providers reviewed, clients routinely hold four, five, and in some cases six separate entrances into the same supplier.

04 The AI points inward.

This market has a lot of AI, and most of it faces the provider. Assistants that summarise conversations, surface similar cases, recommend the next step, flag a risk before it lands — built for engineers and service teams, and they work. The client, at the other end of the same organisation, still receives the raw view and a phone number.

What These Two Have in Common

Both are covered by people, and those people are good at it. Nothing here is broken; it is staffed. That is a reasonable answer to either problem, until the volume, client count, or vendor count changes, at which point it becomes a headcount question rather than a systems one.

Where Off-the-Shelf Tools Fall Short

The reasonable response to everything on the previous page is: there is software for this, and we are looking at it. Sometimes that is right. It is worth being precise about when.

01 Where the Requirement is Specific

Extraction is available off the shelf, and it is good. Any major cloud platform can read a document, find the fields, and return structured output. That is the part that looks hard and is not.

The part that is actually yours is what to extract, and where it goes. Which of eleven near-identical figures on a vendor statement is the one your finance team reconciles against? Which of your systems owns that field? What happens when the same vendor sends two documents that disagree? A product general enough to be sold to every industry is general in exactly the place where your requirement is specific.

02 Where the Process is Too Expensive to Change

A tool is a set of assumptions about how work should be done. Adopting one means accepting those assumptions and reshaping the work to match.

For a firm of forty people, that is usually a good trade. The process is young, few people hold it, and the tool's assumptions are probably better than the ones that grew by accident.

For a firm of several hundred or thousand, it is the opposite trade, and this is the part that gets underestimated. At that size, the process is not an inefficiency wrapped around the software — it is the thing that works. It encodes years of exceptions: the clients who are invoiced differently, the vendor whose claims arrive late, the approval that must exist for audit. The cost of bending it around a product is not a licence fee. It is retrained people, rewritten procedures, renegotiated internal SLAs, and a period of months in which the old way and the new way both run.

Against that, a narrow application built to sit on the process you already have is the smaller piece of work. **Past a certain size, it is cheaper to fit an application to your processes than to fit your processes to an application.**

03 Where the Documents Cannot Leave

The third reason is the one most often cited: firms decline to send their own commercial documents — vendor pricing, client contracts, rebate claims — to an external service. Assurances are offered, and they are generally not believed. So the work is done inside the firm's own infrastructure, and the question becomes what to build there.

There is also a cost argument hiding behind the security one. An enterprise licence for a general product, plus the customisation needed to make it fit, frequently comes out above the cost of building the narrow thing.

Score Your Own Operation

Six questions and five minutes of your time. Score each from **0 to 3** and add them up.

#	Question	0	1	2	3
1	How much of the information you receive from partners, suppliers, and other senders can your systems process without human intervention?	☐ None — everything is opened by hand	☐ A little, for the largest few	☐ Most of the volume, by exception	☐ Most, and the exceptions are known and listed
2	How many hours a month go into re-keying data from emails, PDFs or external portals into your own systems?	☐ Nobody has counted	☐ Roughly estimated once	☐ Counted	☐ Counted, owned, and acted on
3	When a sender changes their format, what happens?	☐ Somebody notices later, when a figure is wrong	☐ The person who handles it spots it	☐ It is caught in the same week	☐ It is caught on arrival and handled as routine
4	How many separate logins does your largest client hold with you?	☐ Four or more	☐ Three	☐ Two	☐ One, and it is genuinely one
5	When a client asks what a number on their dashboard means, who answers?	☐ A person, on a call, every time	☐ A person, but from a standard explanation	☐ The system answers some of it	☐ The system answers most of it, in context
6	What share of client questions do clients resolve without contacting you?	☐ Unknown	☐ Estimated, not tracked	☐ Measured	☐ Measured, and the remaining questions are known

WHERE YOU LAND

Score

14-18	The seam is covered	Both ends are systems rather than people. This is rare, and usually the result of deliberate work rather than good fortune
8-13	One end is built, the other is staffed	The most common result. In this market it is normally the client-facing end that is still staffed
0-7	Both ends run on people	And they run well, which is the reason it has never reached a roadmap

Want a second opinion on your score?

Share your results with us. We will look at where the manual work sits, and what is realistically worth automating.

DISCUSS YOUR SCORE ↗

What It Looks Like When It Is Built

If questions 1 to 3 were the low ones

The change at this end is **where the reading happens: on arrival, by a system, instead of later, by a person.**

For a UK client, documents arriving in whatever form the sender chose are now classified, extracted, and normalised as they land. The difficulty was never reading them, but deciding which fields mattered, which system owned each, and what to do when two documents disagreed. That was the part no product arrives knowing — and where most of the actual work lay.

95/100 document extraction score	500+ scanned handwritten documents processed	96% faster document processing
--	--	--

[FULL CASE STUDY ↗](#)

If question 4 was the low one

The change here is not a better destination. It is **not needing more than one.**

For another UK client, service access, documents, support and reporting had accumulated across roughly ten internal departments — the way four and five logins appear without anyone ever deciding to build a fourth. They were consolidated into one: a single sign-on behind which sit the core applications, the third-party tools and the vendor portals the client had previously reached separately.

[FULL CASE STUDY ↗](#)

If questions 5 or 6 were the low ones

The move at this end is an **AI assistant built inside the system your clients already sign into** — not added alongside it, and not pointed at your own engineers. Things it has to do:

- ☐ **explain the figure in front of them** — what it means, whether it is good, and what usually follows from it;
- ☐ **answer the question they would otherwise phone about**, at the point it occurs to them rather than in a knowledge base they have to think to open.

Assistants are not scarce, and this market has them already. What a general one cannot hold is **your definition of good** — it sits in your service, your thresholds, and what your own team would say about that number. That is the part that gets built rather than bought off the shelf.

The ordinary pieces still matter alongside it: in the portal above, a knowledge centre searchable by topic and a guided setup for new client teams produced fewer repeat questions and a shorter time to first value.

Know which process you would change first?

Bring us the use case, and we will help you define the right solution.

[EXPLORE THE USE CASE ↗](#)