



KEDBYTE

SOFTWARE · AI · AUTOMATION

HOW MONEY MOVES

Volume IV — Moving Money Without Cards

The complete plain-English guide to payments and banking,
followed by the engineer's version of the same thing.

WRITTEN BY

Shikhar Singh

Founder & Chairman

KedByte Technologies Private Limited

Five Volumes • Fifty-Four Chapters • First Edition

HOW MONEY MOVES

Volume IV — Moving Money Without Cards

Author	Shikhar Singh
Designation	Founder & Chairman, KedByte Technologies Private Limited
Published by	KedByte Technologies Private Limited
Imprint	KedByte Press
Edition	First Edition
Subject	Payments, banking systems, financial technology
Extent	Five volumes, fifty-four chapters

Copyright © KedByte Technologies Private Limited. All rights reserved.

No part of this publication may be reproduced, distributed or transmitted in any form or by any means, including photocopying, recording or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other non-commercial uses permitted by copyright law.

Trademarks. All product names, company names, service marks and trademarks referred to in this book are the property of their respective owners. They are used in this book in an editorial and educational capacity only, with no intention of infringement of the trademark. No such use, or the use of any trade name, is intended to convey endorsement or any other affiliation with this book or its publisher.

Accuracy and currency. This book describes technology that changes. Specifications, prices, version numbers, product names and industry figures were checked at the time of writing and are dated in the text wherever they are liable to change. The chapters on artificial intelligence in particular describe a field that moves quickly; readers are told, in the text, where to re-check a figure before relying on it. Where experts disagree, the book says so rather than choosing a side.

Disclaimer. The information in this book is provided on an “as is” basis, for education. The security chapter is written for defence: it explains how classes of attack work so that systems can be built to resist them, and it deliberately contains no working exploit code. Neither the author nor KedByte Technologies Private Limited shall have any liability to any person or entity with respect to any loss or damage caused, or alleged to have been caused, directly or indirectly by the information contained in this book.

Worked examples. Parts F and G are built around one real network fault and one real version-control session, reproduced exactly as they were observed, including the addresses, the traceroute output and the error messages. Nothing in those traces has been altered or invented.

From the Founder & Chairman

You tap a card against a machine and a shop lets you leave with the coffee. Somewhere behind that gesture, money you cannot see is supposed to move from a place you have never visited to another place you have never visited, and almost nobody involved can tell you how. Not the person who served you. Often not the person who installed the terminal. Frequently not the people who work at the bank.

That collective not-knowing is expensive. It produces businesses that cannot reconcile their own takings, engineers who store money in floating point, support teams who ask customers for the one reference number that cannot possibly help, and founders who discover the difference between a payment being authorised and a payment being settled on the day it costs them something. None of these are stupidity. They are the predictable result of an industry that documents itself for people who are already inside it.

This book is the outsider's way in, written to be useful to the insider. It begins with what money actually is, which is a record of obligation rather than a substance, and it ends with regulation, dispute and the ugly arithmetic of reconciliation. In between: the card in your pocket and the small computer embedded in it, the machine on the counter, the message formats that have carried the world's card traffic for forty years, the difference between a payment clearing and a payment settling, the rails that move money without cards at all, and what happens when any of it fails.

The method is the same one used in *How Computers Work*, and it matters more than the coverage. Every idea is explained twice. First in plain language, with an everyday comparison and a worked example, such that a twelve-year-old could follow it. Then again in the exact professional vocabulary, with the real field numbers, the real timings and the real institution names that a practitioner would use. Between the two sits a section that most writing on this subject omits entirely: an explicit statement of where the simple version stops being true. A useful simplification told without warning is how people end up confidently wrong, and in payments confidently wrong is a category of loss.

One thing about this book differs from its predecessor, and the reader should know it. How a transistor works will not change. Payments will. Scheme rulebooks are revised annually, regulatory caps move, and migrations between message standards run on live deadlines. The volumes are therefore arranged so that decay is contained: the mechanics that will outlive us sit in the early volumes, and the material that is true as of writing is gathered deliberately at the end, marked as

such. Where a figure is time-bound, this book says so rather than pretending to a permanence it does not have.

At KedByte Technologies Private Limited we build systems that move other people's money, which is a privilege that ought to come with an obligation to explain the machinery. If this book leaves one reader able to follow their own money from a card terminal to a bank ledger, and to say honestly which part of that journey they have verified and which part they have merely assumed, it has done its job.

Start at Chapter 1. It begins long before cards, and for good reason.

Shikhar Singh

Founder & Chairman

KedByte Technologies Private Limited

How to read this book

The shape of every section

Every section of every chapter is built from the same six blocks, always in the same order. Once you know the shape, you can navigate two thousand pages without ever being lost.

Block	What it is for
PLAIN in simple words	The idea for a complete beginner. No jargon at all.
PLAIN a picture in your head	One everyday comparison, and then a line telling you exactly where that comparison stops being true.
PLAIN a worked example	Concrete numbers, worked step by step.
PLAIN what is really happening inside	The mechanism, still in easy words, but with nothing hand-waved.
TECHNICAL the engineer's version	The correct professional vocabulary, exact units, real specifications, real figures, standard numbers and the commands that observe it.
WORDS remember these	Each term given twice: the plain meaning, then the technical meaning.

Two ways to read it

1. **The single pass.** Read everything, in order, from Chapter 1. This is the intended route and it works.
2. **The double pass.** Read only the blocks marked PLAIN, front to back. You will understand the whole of computing. Then return to the start and read the blocks marked TECHNICAL to be able to hold your own in an engineering conversation.

Conventions used throughout

1. Almost everything is written as short numbered lines, one idea per line, so nothing hides inside a paragraph.
2. *The honest version* marks the places where a simple explanation is not strictly true. The accurate model follows immediately.

3. The book states whether a thing is a **standard** (written down in a specification), a **convention** (everybody just does it this way), or an **implementation detail** (true of one product, not of the idea).
4. In the artificial intelligence chapters, claims are separated into established fact, active research, and marketing claim.
5. Anything that goes stale is dated. Check it again before you rely on it.
6. Every chapter ends with *Common wrong ideas* and a twenty-line summary. If you are short of time, read those two first.

Where to find things

1. **Chapter 1** is the map of the whole book.
2. **Chapter 52** is the reference desk: the numbers, the command cheat sheets and the two diagnostic decision trees.
3. **Chapter 53** is the master timeline and a glossary of more than seven hundred terms, each given in plain words and then technically.

Contents

Volume IV — Moving Money Without Cards **1**

1	Batch Against Real Time	2
	34.0 What this chapter gives you	2
	The plain version	3
	Where the plain version stops being true	4
	The technical version	6
	34.98 Common wrong ideas	14
	34.99 Chapter summary in 20 lines	15
2	Bacs and the Three-Day Cycle	17
	35.0 What this chapter gives you	17
	The plain version	18
	Where the plain version stops being true	19
	The technical version	20
	35.98 Common wrong ideas	30
	35.99 Chapter summary in 20 lines	30
3	The Mandate	32
	36.0 What this chapter gives you	32
	The plain version	33
	Where the plain version stops being true	35
	The technical version	36
	36.98 Common wrong ideas	44
	36.99 Chapter summary in 20 lines	45
4	Faster Payments	47
	37.0 What this chapter gives you	47
	The plain version	48
	Where the plain version stops being true	50
	The technical version	51
	37.98 Common wrong ideas	58
	37.99 Chapter summary in 20 lines	59
5	CHAPS and RTGS	61
	38.0 What this chapter gives you	61
	The plain version	62
	Where the plain version stops being true	64
	The technical version	66
	38.98 Common wrong ideas	72
	38.99 Chapter summary in 20 lines	73
6	Addressing: Sort Codes, IBANs and Confirmation of Payee	75
	39.0 What this chapter gives you	75
	The plain version	76
	Where the plain version stops being true	78

The technical version	79
39.98 Common wrong ideas	88
39.99 Chapter summary in 20 lines	89
7 SEPA	91
40.0 What this chapter gives you	91
The plain version	92
Where the plain version stops being true	93
The technical version	95
40.98 Common wrong ideas	105
40.99 Chapter summary in 20 lines	105
8 SWIFT and Correspondent Banking	107
41.0 What this chapter gives you	107
The plain version	108
Where the plain version stops being true	110
The technical version	112
41.98 Common wrong ideas	120
41.99 Chapter summary in 20 lines	120
9 UPI	122
42.0 What this chapter gives you	122
The plain version	123
Where the plain version stops being true	124
The technical version	125
42.98 Common wrong ideas	133
42.99 Chapter summary in 20 lines	134
10 Instant Rails Elsewhere	136
43.0 What this chapter gives you	136
The plain version	137
Where the plain version stops being true	139
The technical version	141
43.98 Common wrong ideas	147
43.99 Chapter summary in 20 lines	148
11 Open Banking	150
44.0 What this chapter gives you	150
The plain version	151
Where the plain version stops being true	153
The technical version	155
44.98 Common wrong ideas	163
44.99 Chapter summary in 20 lines	164



VOLUME IV

Moving Money Without Cards

Batch against real time. Bacs, Faster Payments, CHAPS,
SEPA, SWIFT and UPI, and why each exists.

Batch Against Real Time

34.0 What this chapter gives you

1. You will be able to explain why a £2,400 salary lands on Thursday while a £248,000 house completion lands before lunch, without reaching for “one system is old”.
2. You will be able to place any rail on two axes — continuous or cyclical, gross or net — and say which of the four boxes CHAPS, Bacs and Faster Payments occupy.
3. You will be able to separate clearing from settlement, and say precisely what happens in seconds in Faster Payments and what happens three times each business day.
4. You will be able to explain how netting discharged £6.05 trillion of Bacs payments in 2025 with average daily interbank settlement of £5.6 billion.
5. You will be able to say what prefunding is, why it converts credit risk into opportunity cost, and what a Net Sender Cap caps.
6. You will be able to read the Bank of England’s RTGS daily timetable and explain why Bacs settles at 9.30am and the card schemes at 11.15am and 12.15pm rather than the other way round.
7. You will be able to say why the £1.8 million average CHAPS payment matters less than the £4,586 median, and what that median reveals about why people use CHAPS.
8. You will be able to explain Herstatt risk from the events of 26 June 1974, and name the four tools that shrink, collateralise, cap or eliminate the gap it describes.
9. You will be able to argue that batch is not simply legacy, citing the recall window, file-level control and predictable funding that real-time rails cannot offer.
10. You will be able to explain why the Bank of England is extending settlement hours, and why that is an argument for making batch safer rather than for abolishing it.

Two payments leave the same British bank at nine o’clock on a Tuesday morning. One is a salary of £2,400 for a warehouse supervisor in Wolverhampton. The other is £248,000 for the completion of a house purchase in Hove. Both are in sterling, both electronic, both instructed by people at screens inside regulated institutions. The salary lands on Thursday. The house money lands before lunch.

The obvious explanation is that one system is old and one is new, and that given enough investment everything will eventually land before lunch. That explanation is wrong, and believing it will make the rest of this volume incomprehensible. The salary is slow on purpose; the house money is expensive on purpose. They are two answers to a question every payment system in every country must settle before it processes a single instruction: when a payment is made, do you act on it now, or write it down and act on a pile of them later?

That is the divide. It runs underneath Bacs and Faster Payments and CHAPS, underneath SEPA and the American ACH, underneath every card scheme in Volume III, and underneath every promise a

fintech has made about instant settlement. It is not a divide between good and bad, or modern and legacy. It is a trade between cost, risk, liquidity and operational calm, and Britain runs several payment systems side by side because different payments want different points on that trade.

The plain version

Imagine a street with two shops on it. Priya runs the bike shop; Tom runs the café opposite. They buy from each other all day long — Tom buying inner tubes and puncture kits for the bikes he lends to customers, Priya buying coffee and sandwiches for herself and her staff. There is no card machine and no bank in this story, just a cash box in each shop and a road between them.

There are two ways to run this street.

The first is that every time money is owed, somebody walks. Tom buys a £6.50 puncture kit, so he takes £6.50 out of his cash box, crosses the road and puts it in Priya's. Twenty minutes later Priya buys two coffees for £5.40, so she walks back with £5.40. Every purchase is a walk. Nobody ever owes anybody anything, because the moment a debt exists it is already paid.

The second is that they keep a slate — a chalkboard by the till. Tom writes "£6.50 to Priya". Priya writes "£5.40 to Tom". Nobody walks. At six o'clock they compare slates, work out who is ahead, and one person makes one walk with the difference.

Now put real numbers on a real Tuesday. Tom buys from Priya eleven times, totalling £54.30. Priya buys from Tom nine times, totalling £49.80. Under the first system there are twenty walks across the road, moving £104.10 of coins. Under the second there is one walk at six o'clock, moving £4.50.

Four pounds fifty instead of a hundred and four pounds ten; one crossing instead of twenty. The slate made nobody richer or poorer — Tom and Priya end the day exactly where they would have either way. What it did was cancel the payments against each other before any coins moved. That cancelling is the single most useful trick in money, and the whole first half of this book has been building towards it.

■ Why the slate gets better as the street gets longer

Add shops. Now there are twelve businesses on the street and by the end of the day they have made two hundred payments to each other. Under the walking system that is two hundred crossings. Under the slate system, all twelve slates go to the newsagent at six o'clock, who works out that seven of the twelve owe money overall and five are owed it, and that the totals come to £312. Seven people walk to the newsagent with money; five collect from her. Twelve walks, not two hundred.

That is roughly what Bacs does with British payroll, and the numbers are not small. In 2025, Bacs handled 6.86 billion payments worth £6.05 trillion. Because it settles once each working day, the amount that actually had to be moved between banks averaged £5.6 billion a day — somewhere near £1.4 trillion across a working year to discharge £6.05 trillion of payments. Roughly four pounds of payments for every pound that moves. The other three pounds cancelled themselves out on the slate.

■ What the slate costs you

The slate is not free, and a bright twelve-year-old spots the problem before an adult does: what happens if a shop shuts before six o'clock?

Suppose Tom sells the café at four and disappears. Priya has been serving him all day on the slate. Under the walking system she would already have his coins in her box; under the slate system she

has chalk. She spent the day handing over inner tubes in exchange for a number on a board, and the number is now worthless.

That gap — the hours between the goods changing hands and the money changing hands — is the price of the slate. Everybody on the street is trusting everybody else to still be solvent at six o'clock. It usually works. Occasionally, catastrophically, it does not, and that is why some payments do not use a slate at all.

■ The trolley and the runner

The other half of the divide is about timing rather than cancelling, and the school post room is the better picture for it.

A school office has a trolley. Letters posted into the internal tray are taken round the building once a day at four o'clock. This is efficient: one person, one circuit, hundreds of letters. It also means a letter dropped in the tray at 4:05pm sits there for twenty-four hours, not because anybody is being slow but because the trolley has gone. There is a cut-off, and being one minute the wrong side of it costs a whole day.

The alternative is a runner, who takes each letter to its recipient immediately. That is fast for you, expensive per letter, and it means the school must keep a runner sitting available at all times, doing nothing, in case somebody has a letter.

Both exist in British money. The trolley is Bacs, which takes three working days on a fixed timetable and carries most salaries and Direct Debits. The runner is CHAPS, which will move a house purchase across London this afternoon and charge for it. And there is a third thing, Faster Payments, which looks like a runner to you and behaves like a trolley behind the scenes: the message about your money sprints across in seconds, while the coins between the banks go on a trolley three times a day.

■ The numbers that make the point

Here is the whole argument in one comparison. In 2025, Bacs handled 6.86 billion payments; CHAPS handled 53.3 million, about one for every hundred and twenty-nine Bacs payments. But the money that had to be moved between banks to settle Bacs averaged £5.6 billion a day, while CHAPS moved £371.3 billion a day. The system with a hundred and twenty-nine times fewer payments moved sixty-six times more actual money.

That is not a paradox. It is the design. Bacs carries an enormous number of small, predictable, cancellable payments and squeezes them down before anything moves. CHAPS carries a small number of large, urgent, un-cancellable payments and moves each one on its own, immediately, in full.

You should now be able to say the following at dinner and be right. Some payment systems collect instructions and process them together, which is cheap and lets payments cancel each other out, but leaves a gap in which somebody could fail. Others act on each instruction the moment it arrives, closing the gap but requiring everybody to keep a great deal of money sitting ready at all times. Britain runs both, because British payments are not all the same kind of payment.

Where the plain version stops being true

The street with the slate is a good picture of netting and a poor picture of a payment system. Four corrections, and each one matters for what follows.

The first correction is that “batch” and “slow” are different things, and so are “real time” and “instant”. The analogy makes it sound as though the trolley is slow because collecting things takes time.

It is not. A Bacs file of four hundred thousand salary payments is validated, sorted and distributed by the central infrastructure in a tiny fraction of the three days it takes to reach people's accounts. The three days are a rulebook, not a speed limit — a deliberate agreement about which day is input, which is processing and which is entry, so that every bank knows exactly when to expect what. Double the computing power tomorrow and the cycle would still be three days.

The reverse correction is sharper. Faster Payments credits the person you are paying in seconds, twenty-four hours a day. But the banks do not settle with each other in seconds. As of writing, the net obligations arising from Faster Payments settle in the Bank of England's system three times each business day. What happens in seconds is the clearing — the exchange of the instruction and the crediting of the customer. What happens three times a day is the settlement — the movement of central bank money between the banks. Your money is available immediately; the banks square up later. Confusing those two clocks is the commonest error in payments writing.

The second correction is that the slate hides who is carrying the risk, and modern systems do not leave it lying around. In the analogy Priya simply absorbs the loss if Tom disappears. That is not how a designed payment system works. The customer receiving a Faster Payment gets usable funds immediately while the sending bank's obligation is not discharged until the next settlement cycle, so the exposure is real — and it is managed rather than tolerated. Bacs, Faster Payments and the cheque-based Image Clearing System all use prefunding: each settling participant must hold cash in a segregated account at the Bank of England covering the largest net debit it could possibly reach before the next settlement, and if it fails, that ring-fenced cash completes the settlement. The Bank's own description is unambiguous — prefunding eliminates settlement risk in the systems that use it, at the cost of tying up liquidity.

So the honest version of the slate is that Priya may trade with Tom on chalk only up to a limit, and Tom deposited that limit in cash with a trusted third party before the day began. The risk did not vanish; it became a liquidity cost.

The third correction is that there is one cash box in the analogy and two tiers of money in reality. When Bacs nets four hundred thousand salary payments down to a single figure, what finally moves is not "money" in the sense a customer means. It is a balance in a settlement account at the Bank of England — central bank money, the only kind of sterling that can extinguish an obligation between banks with finality. The salary in the supervisor's account is commercial bank money: a liability of his bank to him. Netting compresses obligations denominated in commercial bank money; settlement discharges the residue in central bank money. A system that clears in one form of money and settles in another has a seam in it, and every argument in this volume about risk, finality and failure lives in that seam.

The fourth correction is that "batch versus real time" is not one axis but two, and there are four boxes, not two. The plain version blurs the question of *when* obligations are calculated with the question of *how*. A system can process each payment individually the moment it arrives, or in periodic cycles; independently, it can settle each payment at face value, or settle only the net difference. CHAPS is individual and gross. Bacs is cyclical and net. Faster Payments is individual for clearing and cyclical and net for settlement — a hybrid that did not exist in Britain before 2008 and now carries most of the country's non-card retail payments. Once you see the two axes, the industry's vocabulary stops sounding contradictory.

The technical version

■ Two axes, four quadrants

Clearing is the exchange, validation and reconciliation of payment instructions between participants, including calculation of the resulting obligations. **Settlement** is the discharge of those obligations by the transfer of funds — for systemically important sterling systems, in central bank money across accounts at the Bank of England. Chapter 7 established this distinction; here it becomes the coordinate system.

The first axis is timing: does the system act on each instruction as it arrives, or accumulate instructions and act on a defined cycle? The second is obligation calculation: does each payment settle at face value, or are payments offset so that only a net position settles?

	Gross (each payment at face value)	Net (offset positions only)
Continuous / real time	Real-time gross settlement. CHAPS.	Rare; continuous offsetting exists as a liquidity tool inside RTGS, not a settlement model.
Cyclical / deferred	Deferred gross settlement. Uncommon in sterling retail.	Deferred net settlement. Bacs, Faster Payments, LINK, Image Clearing System, Visa Europe, Mastercard Europe, PEXA.

The bottom-right box is where most payments in most countries live, because the arithmetic is cheapest there. The top-left box is where most *value* lives, because the risk is smallest.

■ What the Bank of England actually operates

As of writing (August 2026), the Bank of England supports four settlement models across its real-time gross settlement service: real-time gross settlement, delivery versus payment, and prefunded and unfunded variants of deferred net settlement.

Real-time gross settlement is used by exactly one system: CHAPS. Delivery versus payment is used by the payment leg embedded in CREST, the securities settlement system operated by Euroclear UK & International, where sterling settlement takes place in high-frequency cycles through the day, cash settling against securities in real time against segregated liquidity.

Deferred net settlement is used by seven sterling systems, all settling their net obligations across accounts in RTGS: Bacs, Faster Payments, the cheque-based Image Clearing System, LINK, Mastercard Europe, PEXA and Visa Europe. Each determines its own access criteria and its own number and duration of settlement cycles. Sterling settlement for CLS, the foreign exchange settlement system, runs through CLS's direct participation in CHAPS rather than as a separate model.

The relative sizes are worth stating precisely, because most people's intuition is wrong by two orders of magnitude. In the second quarter of 2026, average daily settlement across RTGS totalled £902,018 million: CREST delivery-versus-payment £477,841 million, CHAPS £411,562 million, and the seven retail net settlement systems between them £12,615 million — Bacs £6,205 million, Visa Europe £2,356 million, Faster Payments £2,298 million, Mastercard £1,480 million, LINK £238 million, cheque imaging £37 million, PEXA £1 million. Every retail payment made in Britain by every non-card rail and both major card schemes, netted down, comes to about three per cent of what CHAPS settles the same day.

■ The RTGS daily timetable

The clearest single artefact in British payments is the Bank of England's published timetable of settlement events, the skeleton on which everything else in this volume hangs. Accurate at time of writing, the sterling settlement day runs as follows, with the caveat the Bank itself attaches: timings can vary in contingency.

Time	Event
5.15am	Transfers between own accounts and user interface access enabled; Note Circulation Scheme settlement
6am	Start of CHAPS settlement; start of CREST DvP, free-of-payment and delivery-by-value settlement
7am to 11am	Hourly CLS pay-in and pay-out deadlines
7am	Faster Payments settlement (first cycle)
8.45am	PEXA settlement
9.30am	Bacs settlement
10.45am	PEXA settlement
11am	LINK settlement
11.15am	Mastercard settlement
12.15pm	Visa settlement
12.45pm	PEXA settlement
1pm	Faster Payments settlement (second cycle)
2.45pm	PEXA settlement
3.45pm	PEXA settlement; end of equity and gilt DvP settlement
4.30pm	Image Clearing System settlement
4.45pm	PEXA settlement
5pm	Faster Payments settlement (third cycle)
5.30pm	End of delivery-by-value settlement
5.40pm	End of CHAPS settlement for customer payments (pacs.008)
6pm	End of CHAPS settlement for financial institution payments (pacs.009); end of free-of-payment settlement
6.30pm	Note Circulation Scheme settlement
8pm	Latest end of contingency extension for RTGS

Two things deserve underlining. First, the deferred systems are not scattered randomly; their cycles are placed so that money arriving from one is available to fund obligations in the next — Bacs at 9.30am, before LINK at 11am and the card schemes at 11.15am and 12.15pm. Second, CHAPS has two closing times: 5.40pm for customer payments carried in the ISO 20022 pacs.008 message, and 6pm for financial institution payments in pacs.009. The gap gives banks twenty minutes to move liquidity between themselves after the last customer payment.

■ Bacs: the shape of a three-working-day cycle

Bacs is the archetype of deferred net settlement in Britain, running since 1968; over 183.6 billion transactions have been credited or debited to British accounts through it. It carries two products: Direct Debit, where the beneficiary pulls funds under a mandate, and Bacs Direct Credit, where the payer

pushes them. In 2025 it processed 6.86 billion payments worth £6.05 trillion, including 5.0 billion Direct Debits, an annual high.

The cycle has three named days. **Input day (day one)** is the last day on which a service user or a Bacs approved bureau may submit a payment file for that cycle. Files must be transmitted to Bacs between 07.00 and 22.30, normally over Bacstel-IP, and may be submitted up to thirty days ahead of the intended payment date — which is why a payroll bureau can build December's run in November. **Processing day (day two)** is when the central infrastructure sorts the file and delivers the resulting items to each receiving payment service provider, which processes each payment against its own customer accounts. **Entry day (day three)** is when payments are credited to recipients' accounts and debited from the payer's, and it is the day the interbank obligation is settled: Bacs net settlement takes place once each business day in RTGS, at 9.30am.

"Day" here means working day. A file submitted on a Thursday enters accounts the following Monday; a file submitted before a bank holiday weekend waits for the calendar. That is the source of most of the frustration ordinary people feel with Bacs, and of most of its value: because the timetable is fixed and published, a payroll department knows in October exactly which day in December the money becomes irreversible, and can plan its funding to the pound.

There is also a control feature real-time systems cannot offer. Because the file sits in the cycle for two days before entry, payments can be recalled after submission, provided the payer's provider is notified before a specified cut-off. A submitted batch is not yet a payment made.

■ **Faster Payments: real-time clearing, deferred net settlement**

The Faster Payment System, operated by Pay.UK, is the hybrid, and understanding it correctly is the point at which most people's model of British payments either becomes accurate or stays wrong forever.

On the customer-facing side it is a real-time system: payments are cleared and the beneficiary credited within seconds, day and night, three hundred and sixty-five days a year. It supports payments of up to £1 million, and it is worth being exact about that limit. £1 million is the value limit applied by the central infrastructure. Individual banks and building societies apply their own, usually much lower, limits, which vary by channel — mobile app, online banking, telephone and branch may face four different ceilings at the same institution. When a customer says "my bank's Faster Payments limit is £25,000", they are describing their bank's policy, not the rail.

On the interbank side it is a deferred net settlement system, settling three times each business day in RTGS, at 7am, 1pm and 5pm as of writing. In 2025 Faster Payments processed 5.55 billion transactions worth £4.84 trillion, while the net obligations settling in RTGS averaged £2,070 million a day. On a 365-day denominator that is around £13 billion of clearing a day compressed to around £2 billion of settlement, better than six to one. The comparison is indicative rather than exact, because clearing runs every calendar day while settlement runs only on business days, so weekend traffic is carried into Monday's cycles.

The exposure between cycles is handled by prefunding. Each direct settling participant holds funds in a dedicated prefunding account at the Bank of England sufficient to cover the maximum net debit position it could reach before the next settlement. The cap on that position is the participant's **Net Sender Cap**. If a participant fails intra-cycle, the ring-fenced cash completes the settlement and no other participant takes a credit loss.

That is an elegant answer to the problem the plain version raised, and it is not free: prefunding converts credit risk into opportunity cost, because cash in a segregated account is cash not being used

for anything else. Pay.UK's framework for setting these caps has historically been prescriptive, built around a Minimum Net Sender Cap and a Peak Contingency Value that participants had to hold at all times, and Pay.UK announced in 2026 that it is moving to a more flexible model. Treat the detail as live; the principle — a capped, collateralised maximum debit — is stable.

The cost of getting prefunding wrong is not theoretical. The Bank's service records note that on 8 September 2021 the first Faster Payments settlement cycle of the day was delayed by forty-five minutes because a participant had insufficient funds in its settlement account. A system instant to customers can still be held up at the settlement layer by one institution's liquidity management.

■ CHAPS: gross, real time, and what that costs

CHAPS is the sterling high-value system, operated by the Bank of England since November 2017, with more than thirty direct participants and several thousand institutions accessing it indirectly through them. Payment obligations between settlement participants are settled individually, on a gross basis, in RTGS, throughout the business day, and settlement is final and irrevocable between direct participants at the moment it occurs. There is no minimum and no maximum payment value.

The system is usually open from 6am to 6pm, Monday to Friday, excluding bank and public holidays in England and Wales, with customer payments closing at 5.40pm. Until 17 June 2016 the settlement day was 6am to 4.20pm; it was extended by one hour and forty minutes from 20 June 2016, and in the following summer an average of 6.1 per cent of CHAPS daily values and 5.3 per cent of volumes settled in the additional hundred minutes.

This is where the phrase “early-afternoon cut-off” needs handling carefully, because it is repeated everywhere and it is not a system fact. The scheme cut-off for customer payments is 5.40pm. What is early-afternoon is the internal cut-off imposed by individual banks on their own customers, set earlier — often much earlier — to leave room for checks, funding and exception handling. The Bank states this explicitly: customer cut-off times set by providers may be earlier. A solicitor who says “we must instruct by two o'clock” is describing their bank, not CHAPS.

Direct participants may submit payments as **urgent**, settling immediately subject to available liquidity, or as **non-urgent**, accepting some delay in exchange for liquidity efficiency. Non-urgent settlement is usually available from 6am until 5.30pm, after which outstanding items are automatically promoted to urgent. Non-urgent payments run through a liquidity saving mechanism of matching cycles, offsetting payments so that a smaller amount of liquidity discharges a larger amount of value; on average they settle around eight minutes after submission. Even a real-time gross system, then, quietly borrows the netting idea — not to defer settlement but to reduce the liquidity needed to achieve it.

The scale figures for 2025 are these. CHAPS volumes grew 1.1 per cent to a record 53.3 million payments, an average of 210,482 per working day. Total value settled rose 7.3 per cent to £93.9 trillion, an average of £371.3 billion a day. The average payment value was £1.8 million; the median was £4,586. Ninety-four per cent of CHAPS payments were for £1 million or less. Financial institution payments in pacs.009 were 24 per cent of payments but 73 per cent of value. CHAPS represents 0.4 per cent of UK payment volumes but 91 per cent of sterling payment values, excluding flows internalised within providers, and it settles the annual GDP of the United Kingdom every nine working days.

The median matters. If you believe CHAPS is only for enormous wholesale transfers, the £1.8 million average confirms it; the £4,586 median tells the truth. The typical CHAPS payment goes through CHAPS not because it is large but because it must be certain and it must be today.

■ The arithmetic of netting, and the arithmetic of liquidity

Under gross settlement, a participant must fund every outgoing payment at the moment it settles, from liquidity it holds. Its peak requirement over the day depends on the timing and size of its outflows relative to its inflows: pay out £3 billion before receiving anything and you need £3 billion available. CHAPS direct participants manage this with their own reserves, secured intraday liquidity provided by the Bank against eligible collateral, and the liquidity saving mechanism, which reduces the amount required by matching offsetting payments rather than deferring them.

Under deferred net settlement, a participant needs only enough liquidity to cover its net position at the settlement point, and if it is a net receiver, none at all. Against that, it must either accept credit exposure to other participants between cycles or prefund the maximum net debit it could reach, pinning down liquidity for the whole cycle regardless of its actual position. Set the two side by side using published figures.

	Bacs	Faster Payments	CHAPS
Payments, 2025	6.86 billion	5.55 billion	53.3 million
Value cleared, 2025	£6.05 trillion	£4.84 trillion	£93.9 trillion
Average value per payment	about £880	about £870	£1.8 million (median £4,586)
Average daily value settled in RTGS, 2025	£5,604 million	£2,070 million	£371,332 million
Settlement model	Deferred net, prefunded	Deferred net, prefunded	Real-time gross
Settlement frequency	Once per business day	Three times per business day	Continuous

Bacs and Faster Payments carry payments of almost identical average size — £880 against £870 — yet one takes three working days and the other takes seconds. The difference is not size but purpose: Bacs carries predictable, scheduled flows where certainty of date is worth more than speed, and Faster Payments carries unscheduled ones where speed is the point. And CHAPS, carrying 0.8 per cent of the payment count of Bacs, moves sixty-six times as much central bank money each day. Real-time gross settlement is not a tweak to netting; it is a different liquidity regime, affordable only because the payments that use it are few.

■ Settlement risk, and why anybody bothers

The reason the whole industry does not simply net everything and settle once a month is a Thursday in June 1974.

Bankhaus Herstatt was a German bank active in foreign exchange. On 26 June 1974 its banking licence was withdrawn and it was ordered into liquidation during the German banking day. Counterparties had already paid Deutsche Marks to Herstatt in Frankfurt that morning; the corresponding dollar payments, due in New York where the business day had not yet ended, were never made. They had irrevocably delivered one leg and never received the other. The exposure created by settling two legs of a transaction at different times, in different time zones, with a gap in between, has been called **Herstatt risk** ever since, and it is the reason CLS exists — a system settling both legs of a foreign exchange transaction simultaneously, whose sterling settlement runs through direct participation in CHAPS, with pay-in and pay-out deadlines every hour between 7am and 11am on the timetable above.

Herstatt is the general case of Priya and Tom. Any gap between the creation of an obligation and its discharge is a window in which a participant can fail and somebody else discovers they were holding

a promise rather than money. The toolkit for that window is small: shrink it by settling more often, collateralise it by prefunding, cap it with net sender caps, or eliminate it by settling gross, in real time, in central bank money.

Finality is the legal counterpart. A settlement that has occurred in RTGS cannot be unwound by a participant's subsequent insolvency. In the United Kingdom that protection comes from the Financial Markets and Insolvency (Settlement Finality) Regulations 1999, which shield transfer orders in designated systems from insolvency law that would otherwise reverse them, together with the recognition regime for payment system operators under the Banking Act 2009. Without the legal certainty the mechanical certainty would be worthless: a liquidator could reach back and undo it.

The Bank of England's current work on settlement hours is framed in exactly these terms. Its consultation paper of 18 May 2026 notes that longer RTGS hours would reduce settlement risk by enabling more frequent net settlement overnight and at weekends, observing that over an Easter weekend the gap between settlements runs from Thursday evening to Tuesday morning. Four and a half days of retail payments accumulating with no settlement in central bank money is the largest exposure window in the sterling system, and it exists because the calendar has holidays in it.

■ Cost per transaction, and what batch buys you

The cost argument is usually asserted and rarely evidenced. It can be evidenced, starting with how the settlement agent charges. The Bank of England operates RTGS and CHAPS on a full cost recovery basis. The CHAPS tariff is an annual participation fee for each direct participant plus a **per-item fee** on debits. For the retail net settlement systems, the Bank charges an annual fee per settlement participant and **no per-item fee at all**. That is the economics of the divide, written into a published price list: a gross system consumes settlement resource per payment and is priced per payment, while a net system consumes it per cycle and its price does not vary with how many payments the cycle contained.

The same shape recurs at every layer. In a batch rail, the marginal cost of the four-hundred-thousandth record in a payroll file is a few dozen bytes and a fraction of a second on infrastructure that was going to run anyway. In a real-time rail it is a full request-and-response round trip with validation, fraud scoring, sanctions screening, an account credit, a customer notification and a durable audit record, executed inside a latency budget of single-digit seconds, on infrastructure sized for peak rather than average and available continuously.

Operational cost differs in kind rather than degree. A batch system has a natural quiet period and an obvious point at which yesterday is closed and reconciliation can run against a fixed, complete set of records; it can be patched and tested between cycles; and its usual failure mode is a late file, recoverable within the cycle without any customer seeing anything. A twenty-four-hour real-time system has none of that: no natural end of day, so reconciliation is cut artificially against a moving target; no maintenance window, so upgrades happen live; and a failure mode that is immediately visible, because the customer is standing there watching for the money. Chapter 53 argues that reconciliation is where fintechs actually die; the reason is that real time removes the seam reconciliation was designed to use.

Batch has three further properties the rhetoric of the last fifteen years has treated as legacy defects and which are in fact advantages. It gives a recall window: a Bacs submission can be withdrawn before the cut-off, whereas a completed Faster Payment cannot be recalled at all, recovery being a request for goodwill rather than a mechanism — which is why authorised push payment fraud is a Faster Payments problem far more than a Bacs problem. It gives file-level control: a failed file is rejected, corrected and resubmitted as a unit, whereas a systematic error on a real-time rail produces thousands of individually completed wrong payments. And it gives predictability of funding: a treasurer running payroll on Bacs knows the exact date and amount days in advance.

■ Why real time is not simply better

Real time eliminates settlement risk in the gross model and shortens it in the hybrid. It gives customers what they want, and it enables use cases — refunds while the customer is still in the shop, gig-economy payouts on completion — that a three-day cycle cannot serve at all.

Against that, it costs more per payment at every layer from the settlement agent downwards. It demands far more liquidity, as intraday balances for gross settlement or prefunded collateral for the hybrid. It removes the recall window, converting fraud from a controllable operational problem into a customer-harm problem with statutory consequences. It removes the maintenance window and the reconciliation seam — the Bank's own consultation weighs a twenty-two-hour settlement day, preserving a defined window for maintenance, upgrades and reconciliation, against a 23.5-hour day that maximises availability but reduces the opportunity to maintain overnight. And it makes availability a systemic obligation rather than a service target: a batch system down for two hours in the afternoon has usually harmed nobody, whereas a real-time system down for two hours has visibly failed thousands of people.

The Bank's published availability record makes the last point concrete. On 18 July 2024 a 245-minute outage to CHAPS settlement was caused by an issue with Swift's Y-Copy service globally; settlement was extended by around ninety minutes and everything submitted settled successfully. On 31 July 2024 a 91-minute disruption caused by an expired certificate in the Bank's IT infrastructure delayed CHAPS and a number of retail settlements. On 14 August 2023 a technical issue kept RTGS closed from 6am until 12.20pm, everything nonetheless settling before the end of the day. These are the incident records of a well-run system, and they are instructive precisely because the deferred systems absorbed the disruption while the real-time one propagated it.

■ The spectrum, and where the rest of this volume sits

Every rail in the chapters that follow can be placed on the two axes.

Rail	Clearing	Settlement	Frequency	Availability	Value limit
Bacs	Batch, three working days	Deferred net, prefunded	Once per business day, 9.30am	Input 07.00 to 22.30 on working days	No scheme limit
Faster Payments	Real time, seconds	Deferred net, prefunded	Three times per business day: 7am, 1pm, 5pm	24/7/365	£1 million central limit; lower bank limits
CHAPS	Real time	Real-time gross, central bank money	Continuous	6am to 5.40pm customer, 6pm interbank, working days	None
Image Clearing System	Batch, two-day cycle	Deferred net, prefunded	Once per business day, 4.30pm	Working days	Not a scheme limit
Card schemes (Visa Europe, Mastercard Europe)	Real-time authorisation, batch clearing	Deferred net	24-hour cycles; weekend and holiday cycles settle next business day	24/7 authorisation	Scheme and issuer limits

Rail	Clearing	Settlement	Frequency	Availability	Value limit
SEPA Instant Credit Transfer	Real time, nine seconds under the 2025 rulebook	Varies by clearing and settlement mechanism	Varies	24/7/365 with permitted planned maintenance	No scheme maximum since the 2025 rulebook
Open Banking payment initiation	An API layer, not a rail; initiates over Faster Payments	Whatever the underlying rail does	Underlying rail	Underlying rail	Underlying rail

The last row is the one people get wrong most often. Open Banking is an interface standard that lets an authorised third party initiate a payment from your account with your consent. It does not clear, does not net, does not settle, and holds no accounts at the Bank of England. When a payment initiated through an Open Banking journey arrives in seconds, that is Faster Payments arriving in seconds. Chapter 44 develops this; the point to carry forward is that an API is not a settlement model, and no API design changes where a payment sits on the two axes above.

■ What is moving, as of writing

The batch–real-time frontier is where most of the change is happening. The following are live as of August 2026 and should each be checked against the primary source before being relied upon.

The Bank of England migrated RTGS to a new core ledger and settlement engine on 28 April 2025, replacing the Enquiry Link user interface with BERTI. In February 2026 it announced that CHAPS settlement will open at 01:30 rather than 06:00, Monday to Friday, from September 2027, subject to final confirmation of timelines with affected direct participants. Its consultation paper of 18 May 2026 sets out the next steps towards near-24x7 settlement: an additional settlement day at the weekend, most likely Sunday, with certain UK bank holidays, but not before 2029; then twenty-two hours a day across Monday to Friday and one weekend day, not before 2031; and in the longer term either twenty-two hours every day with a daily maintenance window, or near-continuous settlement with a short daily downtime. Today there are twelve hours in which CHAPS does not settle, reducing to seven and a half in September 2027.

Notice what that sequence is. It is not the abolition of batch. It is the systematic shortening of the windows during which deferred systems accumulate unsettled obligations, by making the gross system available at more times. The Bank’s stated reason for prioritising weekend settlement is that it addresses the longest period during which settlement is unavailable and reduces the risk and liquidity costs borne by the net settlement systems. Real time is being extended to make batch safer.

In the euro area, Regulation (EU) 2024/886 — the Instant Payments Regulation — entered into force in April 2024, with obligations taking effect for euro-area credit institutions on 9 January 2025 and 9 October 2025, and later for electronic money institutions, payment institutions and providers in non-euro EEA countries. The European Payments Council’s 2025 SEPA Instant Credit Transfer rulebook, in effect since 5 October 2025, gives the payer’s provider, the payee’s provider and the clearing and settlement mechanism nine seconds in total to make funds available, within the ten-second maximum set by the amended SEPA Regulation, and removes the scheme-level maximum transaction amount.

In the United States, the Nacha rules raised the Same Day ACH per-payment limit to \$1 million from 18 March 2022, across three daily processing windows; entries above the limit submitted into a same-day window are not rejected but processed for next-day settlement in the next available window. Nacha

has announced a further increase to \$10 million effective 17 September 2027. That is a batch system being pushed up the value curve rather than replaced by a real-time one — the clearest evidence available that the two designs are converging rather than one winning.

■ A note for implementers

Three design consequences follow, and they are why this chapter comes before the rail-specific ones.

Value date is not processing date, and conflating them produces wrong statements, wrong interest and wrong reconciliation. In a batch rail the debit, credit and settlement all carry the same value date by rulebook design, and that date is not the date you submitted; in a hybrid rail the customer credit and the interbank settlement fall on different clocks and sometimes different days.

Cut-offs are business logic, not configuration. A submission at 22:31 on input day is not late; it is a submission into the next cycle, with a different entry day, value date and funding requirement. Systems that treat a cut-off as a warning rather than a state transition generate an expensive class of incident.

And real-time rails require idempotency while batch rails require reconciliation, neither substituting for the other. A real-time payment that times out may or may not have happened, and the only safe response is a retry that cannot double-pay. A batch file that was accepted may still contain items rejected downstream, and the only safe response is to reconcile the returned reports against what you submitted, item by item.

Everything from here is a specific answer to the question this chapter posed. Bacs answers with three days and one settlement. Faster Payments answers with seconds for the customer and three settlements for the banks. CHAPS answers by refusing to accumulate anything at all. SEPA, SWIFT correspondent banking, UPI and the instant rails elsewhere each answer differently again, shaped by their own regulators, their own liquidity conditions and their own tolerance for the gap between a promise and a payment.

34.98 Common wrong ideas

Wrong: batch rails are slow because collecting and processing takes time. **Right:** a Bacs file of four hundred thousand salary payments is validated, sorted and distributed in a tiny fraction of the cycle; the three days are a rulebook, and doubling the computing power tomorrow would not change them.

Wrong: Faster Payments settles in seconds. **Right:** the beneficiary is credited in seconds, but the net interbank obligations settle in RTGS three times each business day; clearing and settlement run on different clocks.

Wrong: deferred net settlement leaves participants exposed to each other between cycles. **Right:** Bacs, Faster Payments and the Image Clearing System all prefund the maximum net debit in segregated accounts at the Bank of England, so the exposure becomes a liquidity cost rather than a credit loss.

Wrong: what settles between banks is the same money the customer sees. **Right:** netting compresses obligations denominated in commercial bank money, while settlement discharges the residue in central bank money — two kinds of money and two sets of books.

Wrong: CHAPS has an early-afternoon cut-off. **Right:** the scheme cut-off for customer payments is 5.40pm; the early-afternoon deadline is the individual bank's internal cut-off imposed on its own customers.

Wrong: CHAPS is only for enormous wholesale transfers. **Right:** the average is £1.8 million but the median is £4,586 and 94 per cent of payments are for £1 million or less, so the typical CHAPS payment goes that way because it must be certain and today.

Wrong: a bank quoting a £25,000 Faster Payments limit is describing the rail. **Right:** the central infrastructure limit is £1 million, and bank and channel limits are policy layered on top of it.

Wrong: Open Banking is a payment rail. **Right:** it is an interface standard that initiates over Faster Payments; it does not clear, net or settle, and holds no account at the Bank of England.

Wrong: a real-time gross system never nets anything. **Right:** CHAPS non-urgent payments run through a liquidity saving mechanism of matching cycles that offsets payments so less liquidity discharges more value, settling on average about eight minutes after submission.

Wrong: the direction of travel is the abolition of batch. **Right:** the Bank is extending gross settlement hours to shorten the windows in which deferred systems accumulate unsettled obligations, and Nacha is pushing Same Day ACH up the value curve — the two designs are converging, not competing.

34.99 Chapter summary in 20 lines

1. Every payment system must settle one question before it processes a single instruction: act on each payment now, or write it down and act on a pile of them later.
2. Acting later lets payments cancel against each other, which is the most useful trick in money: twelve businesses making two hundred payments can square up with twelve walks rather than two hundred.
3. Bacs is the industrial version of that trick, discharging £6.05 trillion of payments in 2025 with an average of £5.6 billion a day actually moving between banks.
4. The price of netting is the gap between goods changing hands and money changing hands, during which a participant can fail.
5. The other half of the divide is timing: a trolley that runs a fixed circuit with a cut-off, against a runner who leaves immediately and costs more per letter.
6. Bacs is the trolley, CHAPS is the runner, and Faster Payments looks like a runner to the customer while behaving like a trolley between the banks.
7. Batch is not the same as slow, because the three Bacs days are an agreement about input, processing and entry rather than a limit on computing speed.
8. Real time is not the same as instant, because Faster Payments clears in seconds and settles its net obligations in RTGS three times each business day.
9. Designed systems do not leave the settlement gap lying around: prefunding in segregated Bank of England accounts converts credit risk into an opportunity cost.
10. Clearing happens in commercial bank money and settlement in central bank money, and every argument in this volume about risk, finality and failure lives in that seam.
11. Timing and obligation calculation are two independent axes giving four boxes, and CHAPS is continuous and gross, Bacs is cyclical and net, and Faster Payments is real time for clearing and cyclical and net for settlement.
12. The Bank of England runs real-time gross settlement for CHAPS alone, delivery versus payment for CREST, and prefunded and unfunded deferred net settlement for seven sterling systems.
13. In the second quarter of 2026 those seven retail systems together settled £12,615 million a day against CHAPS's £411,562 million, so every retail rail netted down is about three per cent of what CHAPS settles.

14. The published RTGS timetable is the skeleton of the sterling day, and its cycles are ordered so that money arriving from one settlement is available to fund obligations in the next.
15. Bacs and Faster Payments carry payments of almost identical average size, about £880 against £870, so the difference between them is purpose rather than value: scheduled certainty against unscheduled speed.
16. CHAPS carries 0.4 per cent of UK payment volumes and 91 per cent of sterling payment values, and settles the annual GDP of the United Kingdom every nine working days.
17. Herstatt's failure on 26 June 1974 is the general case of the settlement gap, and the toolkit for it is small: settle more often, prefund it, cap it, or settle gross in central bank money.
18. Legal finality under the Settlement Finality Regulations 1999 is what stops a liquidator undoing what RTGS has done, without which the mechanical certainty would be worthless.
19. Gross settlement is priced per item and net settlement per cycle, and batch also buys a recall window, file-level control and predictable funding, which is why authorised push payment fraud is a Faster Payments problem far more than a Bacs one.
20. The frontier is moving through earlier CHAPS opening, weekend settlement and a \$10 million Same Day ACH limit, but the direction is convergence: real time is being extended to make batch safer, not to replace it.

Sources: Bank of England (payment and settlement pages, summary of RTGS daily timetable, payment and settlement statistics, CHAPS pages, and the consultation paper "Extending RTGS and CHAPS settlement hours", 18 May 2026); Pay.UK and Bacs (scheme pages, Bacs processing cycle, Net Sender Cap publications); European Payments Council (SCT Inst scheme and 2025 rulebook); Nacha and Federal Reserve Financial Services (Same Day ACH rules). All timings, limits and figures accurate at time of writing, August 2026.

Bacs and the Three-Day Cycle

35.0 What this chapter gives you

1. You will be able to say what happens on each of the three Bacs days, and name the 22:30 deadline that decides which cycle a file belongs to.
2. You will be able to work out the latest input day for any payroll date, including the case where a bank holiday makes the answer Thursday rather than Friday.
3. You will be able to explain why five calendar days can pass inside a three-day cycle without anything having gone wrong.
4. You will be able to say where the money is on day one and day two, and why nothing is in transit.
5. You will be able to separate the clearing that happens on day two from the 09:30 settlement on day three, and say which kind of money each one moves.
6. You will be able to read a Standard 18 payment record by counting characters, and say what fields 7, 9, 10 and 11 carry.
7. You will be able to explain why a Bacs file must balance and what the contra record does for reconciliation.
8. You will be able to distinguish Direct Debit from Bacs Direct Credit by direction, authority, return mechanism and transaction code, and say what sending code 19 by accident does to a mandate.
9. You will be able to say what a Service User Number is, why one organisation may hold several, and how a facilities management arrangement differs from a bureau.
10. You will be able to describe what a Bacs Approved Bureau inspection covers, how the volume tier sets its cadence, and what a single Fail does to the result.

There is a system that moved 6.86 billion payments worth £6.05 trillion in a single year, that pays most of the salaries in Britain and collects most of the household bills, and that almost nobody outside the payments industry can name. It has been running since 1968. It has never been fast, it has never tried to be, and its slowness is not a defect.

Bacs is the oldest of the sterling rails in this volume and the least glamorous. Faster Payments arrives in seconds and gets the headlines. CHAPS carries the property transactions and gets the respect. Bacs carries the *routine* — the salary that arrives on the first of the month and the gas bill that leaves on the eighth — on a schedule so predictable that entire industries have been built on the assumption that it will not change.

That schedule is three working days long. Every technical decision in this chapter follows from it, and so does every operational mistake people make with it. The most common of those mistakes is to believe that “three days” means three days.

This chapter explains what happens on each of those days, what physically travels between the parties, what the six digits of a Service User Number authorise, and why a Monday in late August can push a payroll into September. It ends with the approval regime governing anyone who submits Bacs files on somebody else's behalf.

The plain version

Imagine a very large sorting office in the middle of a town.

Every business in the town — the electricity company, the water company, the factory that pays four hundred wages, the council — needs to move money in and out of thousands of ordinary bank accounts. The electricity company wants to take £61.40 from Mrs Okafor and £48.10 from Mr Reid and similar small amounts from two hundred thousand other people. The factory wants to put £1,842.16 into one worker's account and £2,103.44 into another's.

None of them talks to the banks directly. They all send their lists to the sorting office.

Here is the rule of the sorting office. If your list arrives before half past ten on Monday night, then on Tuesday the office sorts it, and on Wednesday morning the money moves. Not before. Not after. Wednesday, for everybody, at once.

That is the three-day cycle. Day one you hand in the list. Day two the office sorts. Day three the money moves.

■ What the sorting office actually does on day two

You might think the sorting job is obvious: split the pile into "letters for Barclays", "letters for Lloyds", "letters for Nationwide", and send each bank its bundle. The office does do that. But it does something cleverer as well, and the cleverer thing is the reason the whole arrangement exists.

Suppose that on Tuesday the office finds it is holding instructions worth £4,200,000 that must move from Barclays customers to Lloyds customers, and instructions worth £3,900,000 that must move from Lloyds customers to Barclays customers. It could send two lorries of cash in opposite directions. Instead it does the subtraction, and on Wednesday morning one payment of £300,000 goes from Barclays to Lloyds and the matter is closed.

Eight million instructions have become a couple of dozen numbers. That is what the office is for. It is not a courier. It is an accountant.

■ A worked example, with real numbers

Barrow Lane Bakery employs fourteen people and pays them monthly. September's pay date is Tuesday 1 September 2026. The total wage bill is £31,420.00.

On Thursday 27 August, some time between seven in the morning and half past ten at night, the bakery's office manager sends one file to the sorting office. It contains fourteen lines. Each has a sort code, an account number, an amount in pence, and a short piece of text that will appear on the employee's bank statement — "BARROW LANE PAY", say. There is also a fifteenth line, which says that £31,420.00 in total is to come *out* of the bakery's own account. That fifteenth line is what makes the file add up to zero.

Nothing happens on Thursday except that the bakery gets a receipt saying the file arrived and looked sensible.

On Friday 28 August the sorting office reads the file, checks it, and works out which bank each of the fourteen employees uses. Fourteen tiny instructions join a national pile of tens of millions. The bakery's money has not moved. The employees have no idea anything is happening.

Then there is a weekend. Then Monday 31 August, the summer bank holiday, and the sorting office is shut.

On Tuesday 1 September, at half past nine in the morning at the very latest and usually a good deal earlier — most people see it before six — the fourteen employees are paid, and £31,420.00 leaves the bakery's account. Both things happen on the same morning. The bakery's money was in the bakery's account, fully usable, right up until then.

Now do the same story backwards. The bakery has a broadband contract at £34.99 a month, and the broadband company collects by Direct Debit. The broadband company sends *its* file on day one; the office sorts on day two; and on day three £34.99 leaves the bakery's account and arrives in the broadband company's. Same three days, same sorting office, same file format. The only difference is which direction the arrow points and who is allowed to draw it.

That is genuinely the whole idea. One rail, two directions. Money out of your account on somebody else's instruction is a Direct Debit. Money into somebody else's account on your instruction is a Direct Credit.

■ Why anyone puts up with three days

Because the alternative costs more and tells you less.

The broadband company has 220,000 customers on £34.99 a month. Collected by instant transfer, that is 220,000 separate real-time transactions to pay for and 220,000 confirmations to chase, and — crucially — it could not *make* the customer pay. Instant payments are pushed by the payer. Direct Debit is pulled by the biller, on an authority the customer signed once.

And a factory paying four hundred wages does not want them arriving at four hundred slightly different moments. It wants them all on the first, so that its statement shows one line of £742,000, its payroll report shows four hundred lines adding to £742,000, and reconciling the two takes a minute rather than a morning.

Predictability is the product. Three days is the price.

Where the plain version stops being true

The three days are not days. They are English bank working days, which excludes Saturdays, Sundays and public holidays, and the definition is *English* regardless of where in the United Kingdom the money is going. A Scottish bank holiday that is not an English one does not stop the cycle; an English one does, even for money moving between two Scottish accounts. The bakery's payroll above spent 27, 28, 29, 30 and 31 August somewhere between submitted and paid — five calendar days for a three-day cycle. That is not a delay. That is the cycle behaving exactly as specified.

Nothing is “in transit” on days one and two, because there is nothing to be in transit. The sorting-office image quietly suggests that your money is somewhere in the middle of the process, in a van, in a vault, in limbo. It is not. On day one and day two your money is in your account and it is entirely yours. What exists between day one and day three is not money, it is a *record of an intention*. The transfer of value happens once, on day three, and it happens in two different places at once: individual customer accounts are debited and credited at the banks, and separately a much smaller set of net amounts

moves between the banks themselves at the Bank of England. Those are two different kinds of money and two different sets of books, and the technical section separates them properly.

Day three is not the end. The plain version ends when the money lands, which is when most people stop thinking about a payment. Bacs does not stop there. A Direct Debit that lands on day three can be returned unpaid afterwards, and the service user finds out through a report that arrives one or two working days later. A credit sent to a closed account bounces back the same way. And a Direct Debit can be reversed under the Guarantee long after everyone has forgotten it, through an indemnity claim that hits the collecting organisation, not the bank. Any system built on the assumption that entry day is final will eventually produce a reconciliation break that nobody can explain.

“Simultaneously debited and credited” is a phrase to be careful with. Bacs’ own published description of the cycle says, for Direct Credits, that on entry day “payments are simultaneously credited to the recipients’ accounts and debited from your account”, and for Direct Debits that they are “simultaneously debited from the recipients’ accounts and credited to your account” — where in the second case the party being debited is the payer, not any sort of recipient. The scheme’s plain-language copy blurs the direction, and so does most vendor documentation copied from it. Whenever you read a Bacs description, stop and ask which party is being debited, which is being credited, and which of them initiated the file. Those three answers, not the word “simultaneously”, are what distinguish the two products.

The technical version

■ Who operates what

Bacs is one of the retail payment systems owned and operated by Pay.UK, the company established as the New Payment System Operator in 2018. Bacs Payment Schemes Limited became a wholly owned subsidiary of the NPSO on 1 May 2018, and the NPSO was renamed Pay.UK on 18 October 2018. The scheme itself dates from 1968; Pay.UK reports that over 183.6 billion transactions have been debited or credited to British bank accounts through it since then.

Pay.UK owns the rules. It does not run the machines. The central infrastructure that physically processes Direct Debit and Bacs Direct Credit and maintains the network is operated by Vocalink, now part of Mastercard, under contract to Pay.UK. When practitioners say “Bacs went down” they mean the infrastructure; when they say “Bacs says you must” they mean the rulebook.

Beneath Pay.UK sit the direct participants: the banks and payment institutions that clear and settle Bacs items in their own right. The published Bacs participant list ran to thirty-two institutions when last checked, Pay.UK’s page having been updated on 7 April 2026, and it is a more interesting list than it once was. Alongside Barclays, HSBC, Lloyds, NatWest, Santander and Nationwide sit ClearBank, Modulr, Monzo, Starling, PayrNet, The Bank of London and the Bank of England itself. Everyone else reaches Bacs indirectly, through an agency arrangement with a direct participant.

And beneath *them* sit the service users: the organisations that actually send files. A service user is sponsored into the scheme by a payment service provider, which is on the hook for the service user’s behaviour. This is the single most important structural fact about Bacs and it has no equivalent on the card rails. You do not sign up to Bacs. You are vouched for.

■ The three days, precisely

Bacs’ own definition, as published, runs as follows.

Day	Name	What happens
1	Input day	The latest day a service user or bureau may submit a payment file for that cycle. Files must be transmitted between 07:00 and 22:30.
2	Processing day	Files are delivered to the recipient PSPs, which then process each payment. Interbank totals are calculated.
3	Entry day	Accounts are debited and credited. Settlement of the net interbank positions takes place.

Several precisions matter.

The input deadline is 22:30 on input day. A transmission completing at 22:31 belongs to the next cycle, whose entry day is a day later. Practitioners run payroll cut-offs hours before this, not minutes.

Input day is the *latest* day of submission, not the only one. Files may be sent up to 30 calendar days ahead of the processing day and held — the scheme term is *warehousing* — until the relevant input day arrives. Most large service users warehouse, because it converts a hard nightly deadline into a soft monthly one. Bacs' glossary consequently describes the processing cycle as having four stages rather than three: arrival, input, processing and entry. Arrival is when your file physically reaches Bacs; input day is when the cycle formally starts for it. For a warehoused file those are different dates.

On input day the service user receives an input report confirming that the file was received and describing what was found in it. Collecting it is not optional in any practical sense: a file that failed validation and was never chased is a payroll that does not run.

On entry day, paying PSPs apply the debits and credits to customer accounts no later than 09:30. For Bacs Direct Credit specifically, Pay.UK states that “it is not possible to guarantee the exact time the payment will reach the recipient's account, although at least 90% of Direct Credit payments are credited to beneficiary account by 06:00am”. That 90 per cent figure is the closest thing the scheme offers to a service-level promise on timing, and it is a statistical statement about the population, not a guarantee about your file.

■ What settles, and in what money

Bacs is a deferred net settlement system. The distinction between clearing and settlement, laid out in Volume I, is load-bearing here and gets skipped constantly.

Clearing is the exchange and sorting of the payment instructions themselves, which happens on day two. Settlement is the movement of central bank money between the direct participants to discharge the obligations that the clearing created, and it happens on day three. Pay.UK's own assessment against the international principles for financial market infrastructures states it precisely: “There is one Bacs settlement cycle each working day at 09:30 and this settles the values calculated at the end of day two.” Final settlement occurs at the Bank of England, over the Real-Time Gross Settlement infrastructure, in central bank money.

The netting is multilateral, not bilateral. Each direct participant ends day two with a single number: the net amount it owes to, or is owed by, the system as a whole. At 09:30 on day three those positions are

posted across settlement accounts at the Bank of England, and a system that handled tens of millions of instructions discharges them with a few dozen entries.

The credit risk that netting would otherwise create — the risk that a participant with a large net debit fails between day two and settlement — is removed by prefunding. Pay.UK records that all direct participants hold a cash deposit sufficient to cover their own net transactions in a segregated account at the Bank of England, and that this removes credit risk between direct participants from the settlement process. Bacs is deferred, but it does not have to wonder whether it will be paid.

There is a consequence for the reader who is building something. On entry day your customers' balances change in *commercial* bank money — an increase in what their bank owes them — while the banks square up in *central* bank money at the Bank of England. The two events are related by the rules, not by the same transaction. A payment can be applied to a customer account and appear on a statement before the interbank settlement that funds it has posted. This is normal, and it is one of the reasons prefunding exists.

■ Direct Debit and Bacs Direct Credit

Two products, one rail, one file format, opposite directions.

	Direct Debit	Bacs Direct Credit
Who submits the file	The collecting organisation (service user)	The paying organisation (service user)
Direction on entry day	Payer's account debited, service user's account credited	Service user's account debited, beneficiaries' accounts credited
Standing authority required	Yes: a Direct Debit Instruction from the payer	No
Typical use	Utilities, insurance, subscriptions, council tax, loan repayments	Payroll, pensions, benefits, supplier payments, expenses, refunds, dividends
Amount	May vary between collections	Set by the payer each time
Consumer protection	The Direct Debit Guarantee	None specific to the scheme
Return mechanism	ARUDD, returned unpaid	ARUCS, returned unapplied
Change advices	ADDACS	AWACS
Transaction codes	01, 17, 18, 19 (and 0N, 0S, 0C for instructions)	99 for credits, with a contra debit

Pay.UK describes a Direct Debit as “an occasional or regular payment instruction set up by the end user in advance, authorising an organisation to collect a pre-agreed amount of money from their account, on a prearranged date”, and Bacs Direct Credit as “a regular payment, normally B2C, that pays the beneficiary on a regular date (i.e. first of the month) and is mainly used for paying wages and salaries”.

The transaction codes are worth memorising because they carry meaning the amount field does not. Within Direct Debit, code **01** marks the first collection under an instruction; **17** marks a subsequent collection; **18** marks a collection represented after a failure; and **19** marks the final collection under the instruction. Code 19 is not decoration: it tells the paying PSP to delete the instruction from its records, and it may cause a “final collection” message to appear on the payer's statement. Sending 19 by accident cancels the mandate.

Separately, the AUDDIS service carries the instructions themselves rather than money: **0N** lodges a new Direct Debit Instruction electronically at the payer's PSP, **0S** converts an existing manually lodged

instruction to AUDDIS, and 0C cancels one. AUDDIS is mandatory for all new service users that submit direct to Bacs. It travels on the same three-day cycle in the same file format, which is why a new customer's instruction and their first collection cannot sensibly be sent in the same submission.

The advance notice rule is the other thing to know. A Direct Debit service user must tell the payer, before collecting, the date and amount to be debited. The default notice period recorded against the service user is 10 working days plus postal time, and it can be varied downwards by agreement with the sponsoring PSP. Shorter notice periods are common for utilities and rare for anything unusual. The mandate, the Guarantee and the indemnity mechanism are the subject of the next chapter; here it is enough to say that they attach to Direct Debit and have no counterpart for Direct Credit.

■ Why payroll and utility billing run on this

Four reasons, in decreasing order of how often they get mentioned and increasing order of how much they matter.

The cheap reason is unit cost. A Bacs item is priced by the sponsoring PSP as a bulk item and a Faster Payment is priced as a transaction. Over 220,000 collections a month the difference is the whole business case.

The better reason is direction. Direct Debit is the only sterling mechanism in common use that lets the person owed money initiate the transfer under a standing authority. Everything else — Faster Payments, CHAPS, Open Banking payment initiation — is a push, which requires the payer to act, which means a proportion of payers will not act. For a utility with millions of accounts, the difference between pull and push is the difference between a collections department and a call centre.

The third reason is that the value date is knowable in advance and identical for everyone in the file. Payroll is not only a payments problem, it is a *fairness* problem: everybody must be paid on the same day, and a system where 3 per cent of staff are paid four hours late is worse than one where everybody is paid at 09:00. Over 150,000 organisations use Bacs Direct Credit, and the reason they do not migrate to instant rails is not inertia; it is that a batch with a fixed value date is the correct shape for the problem.

The fourth reason is reconciliation, and it is the one that gets fintechs killed. A Bacs file is self-balancing. The contra record forces the sum of the payment records to equal the single amount that hits the service user's own bank account, so the bank statement shows one figure that must reconcile to one file. A thousand individual Faster Payments show a thousand statement lines with a thousand chances of a mismatch. Batch is not an old-fashioned way of doing real time. It is a different, and for these purposes better, control structure.

Payroll carries one extra thread, because it explains a field in the file format. Under Real Time Information, HMRC for some years asked employers paying by Bacs to include a hash cross-reference on the Full Payment Submission, built partly from a four-character random string inserted into field 7 of the Bacs payment record, so that HMRC could match reported payroll data to payments actually made. HMRC withdrew the requirement for that field, data item 118, from the 2023/24 tax year onwards. The Service User Number remains reportable on the FPS. The field in the file is still there and payroll software still fills it, which is a small monument to the rule that payment formats accumulate requirements faster than they shed them.

■ What a Bacs file physically is

The format is called Standard 18, and Bacs' glossary defines it in one line as "the Bacs standard file/record format used by the Direct Debit and Bacs Direct Credit schemes". It is a flat, fixed-width, fixed-length text file. There is no XML, no JSON, no delimiter, no escaping. Every field is found by

counting characters from the start of the line, exactly as it was when these files travelled on magnetic tape.

Two record lengths coexist. The label records that wrap the file are 80 characters. The payment records inside it are 100 characters, and there is a longer 106-character variant that appends a processing date to each record.

A submission is a sandwich:

Record	Length	Purpose
VOL1	80	Volume header. Identifies the submitting party; carries the Service User Number.
HDR1	80	First file header. File identifier, creation date, expiration date.
HDR2	80	Second file header. Declares record format F, block length 02000, record length 00100.
UHL1	80	User header label. Carries the Bacs processing day and the work code.
Payment records	100 (or 106)	One per payment, plus one contra.
UTL1	80	User trailer label. Debit and credit totals and counts.
EOF1	80	End of file label, mirroring HDR1.
EOF2	80	End of file label, mirroring HDR2.

The label records are declarative rather than interesting, with the exception of UHL1, which carries the processing date as five Julian digits preceded by a space, in the form YYDDD. The work code sits in the same record and reads 1 DAILY. Bacs' glossary defines the Julian convention exactly: "Julian dates run serially throughout the year. 1 January is always 001 and 31 December is 365 (366 if a leap year)." Thursday 27 August 2026 is Julian 239. Tuesday 1 September 2026 is Julian 244.

The payment record is where the work is. Counting characters from position 1:

Field	Positions	Length	Contents
1	1–6	6	Destination sorting code
2	7–14	8	Destination account number
3	15	1	Destination account type
4	16–17	2	Transaction code
5	18–23	6	Originating sorting code
6	24–31	8	Originating account number
7	32–35	4	Free format
8	36–46	11	Amount, in pence
9	47–64	18	Originating account name
10	65–82	18	Reference
11	83–100	18	Destination account name

Bacs' glossary names fields 9, 10 and 11 individually, which tells you which three the scheme considers most likely to be got wrong. Field 9 is the short name of the service user, up to 18 characters, and is

what may appear on the payer's statement. Field 10 carries the reference, and the glossary calls it "a critical field for both credits and debits allowing the recipient to identify what the payment relates to". Field 11 is the destination account name, and for Direct Debits it "must be the name of the person who is paying the Direct Debit and has authorised the Direct Debit Instruction". Field 7, the four-character free-format field, is the one HMRC borrowed for the RTI hash.

Three properties of this layout deserve emphasis.

The amount has no decimal point and no currency symbol. It is eleven digits of pence, right-justified and zero-filled, so £34.99 is written 00000003499 and £31,420.00 is written 00003142000. Eleven digits will express any amount below one thousand million pounds; the limits that actually bite are the account limits set by the sponsoring PSP, not the width of the field.

The names and references are eighteen characters and no more. Not eighteen and a continuation, not eighteen displayed characters of a longer stored value. Eighteen. Every organisation that has wondered why its customers' statements say THAMESIDE PROPERTY rather than the full company name is looking at field 9 doing exactly what it was built to do in 1968.

And the file must balance. Alongside the payment records sits a contra record, the mirror entry against the service user's own account. In a Bacs Direct Credit file the individual credits carry transaction code 99 and the contra carries a debit code; in a Direct Debit collection file the individual debits carry 01, 17, 18 or 19 and the contra carries 99. The UTL1 trailer then states total debits, total credits and the counts of each, right-justified and zero-filled. Bacs' glossary defines a file, tersely and correctly, as "a self balancing set of records from a service user". If it does not balance it is not a file, it is a rejection.

The definitive specification is the Service User's Guide and Rules to Bacstel-IP, which the glossary describes as "the technical specification for submission to Bacs". It is issued to sponsored service users by their PSP and is not public, so no publicly available scheme document sets out the internal offsets of the reserved and system fields inside the eighty-character labels. Two things fill the gap, and neither is the scheme speaking. The labels are not a Bacs invention at all: they are the American national standard for the file structure and labelling of magnetic tapes, ANSI X3.27, which is why VOL1 and HDR1 carry serial numbers, generation numbers, block counts and buffer offsets that no payment system has ever had a use for. And several sponsoring banks publish their own Standard 18 message implementation guides, HSBC's February 2022 edition among them, which do give the character position of every reserved field in VOL1, HDR1 and HDR2. Those guides agree with one another on the payment record and very largely on the labels, but they are the bank's document rather than the scheme's, and a sponsor is free to impose requirements of its own or to read a field differently. Take the offsets from the guide your own sponsor hands you, not from any book, including this one.

Pay.UK also publishes an ISO 20022 to Bacs translation guide and a translation specification with XSDs covering all Bacs Standard 18 messages, so that organisations speaking ISO 20022 internally can generate compliant submissions. As of writing, the format on the wire is still Standard 18; ISO 20022 is a translation layer over it, not a replacement.

■ Service User Numbers

A Service User Number is, in Bacs' own words, "the unique six digit number allocated to organisations authorised to use the Bacs service". Six digits. That is the entire identifier, and it is doing more work than its size suggests.

The SUN is not an account number and it is not a customer reference. It is the scheme's identity for a sponsored originator, and it is what a paying PSP matches against when a Direct Debit arrives claiming authority to take money from one of its customers. The Direct Debit Originator database

holds the record behind it: contact details, advance notice period, dormancy period, AUDDIS status. When a payer telephones their bank to dispute a collection, the SUN is how the bank knows who to point at.

One organisation may hold several SUNs. Large groups routinely run a separate SUN per operating company, per brand, or per application — one for payroll credits, one for subscription collections — because the SUN is the unit against which scheme configuration, reporting and, eventually, blame are recorded. In-house group bureaux, Bacs notes explicitly, may cover “a single organisation that has different Service User Numbers for different applications”.

An organisation that cannot obtain its own SUN is not shut out. Under a facilities management arrangement, a registered FM provider acts as the service user, using its own SUN, and collects Direct Debits on behalf of an FM client that has not been sponsored into the scheme. The client’s customers see the provider’s short name in field 9. Pay.UK maintains a directory of sponsored FM providers, and the arrangement is fundamentally different from a bureau: a bureau transmits your file under *your* SUN, an FM provider collects under *its own*.

The SUN also appears inside the file, in the owner identification area of the VOL1 label. A file whose VOL1 SUN does not match the submitter’s credentials does not get processed.

■ Getting the file to Bacs: Bacstel-IP, smartcards and the Payment Services Website

Bacstel-IP is the channel. The glossary describes it as “a service providing a highly secure access channel into Bacs” using internet technologies and public key infrastructure to allow access to payment services, and notes that it “carries out some online validation of submissions”. That last clause is the one practitioners live by: some errors are caught at submission, and the rest are caught later and reported.

Access to Bacstel-IP requires software that Bacs has approved for the purpose, and it requires cryptographic credentials. Direct submitters are issued smartcards, which authenticate approved personnel to the Payment Services Website and sign and send the payment files themselves. Bacs’ guidance is strict: cards cannot be transferred between people, a new card must be applied for through the sponsoring PSP when staff change, and the old card must be returned to the PSP to be destroyed. Other staff, and organisations that submit through a bureau, may instead be given a password and ID granting access to the Payment Services Website for reports only, without the ability to sign a submission.

The Payment Services Website is separate from the public bacs.co.uk site, a distinction Bacs itself makes in capitals because people confuse the two. It is where reports and advices are collected. Bacs’ own submission checklist reduces to this: submit the file; collect the input report, available on the day of submission; monitor regularly for ADDACS and AWACS advices, ideally daily; and collect any ARUDD and ARUCS reports, which are available within one to two PSP working days after the payment or collection date.

Acting on those advices is a rule, not a courtesy. Service users must action them within three working days. An ADDACS advice saying a customer has cancelled their instruction, ignored for a month, produces a stream of collections that will fail and may produce indemnity claims. If a bureau or software supplier collects the reports for you, you are still the party required to act on them.

■ Why bank holidays extend the cycle

Bacs defines working days, for scheme purposes, as “English PSP working days excluding Saturdays, Sundays and Public Holidays”. The three-day cycle counts in those days and only those days. It follows mechanically that a public holiday does not delay a Bacs payment by one day; it removes one

of the three slots, and whether that costs you one day or four depends entirely on where in the week the holiday falls.

Take the summer bank holiday in England and Wales in 2026, which falls on Monday 31 August.

Wanted entry date	Latest input day (by 22:30)	Processing day	Entry day
Friday 28 August	Wednesday 26 August	Thursday 27 August	Friday 28 August
Tuesday 1 September	Thursday 27 August	Friday 28 August	Tuesday 1 September
Wednesday 2 September	Friday 28 August	Tuesday 1 September	Wednesday 2 September

Read the middle row carefully, because it is the one that catches people. A payroll with a value date of Tuesday 1 September 2026 must be submitted by 22:30 on **Thursday 27 August**, not Friday. The Friday is the processing day. An organisation that has internalised “submit two days before” as “submit on the Friday for the Tuesday” will find its staff unpaid, and will find out on Tuesday morning, which is the worst possible moment to find out.

Note also that there is no entry day at all on Monday 31 August. Pay.UK’s rule for Bacs Direct Credit is that payments can be sent Monday to Friday excluding bank holidays, and “if the prearranged date falls on a weekend or bank holiday, the payment is made on the next working day”. A salary contractually payable on the last day of August 2026 is paid on 1 September.

Christmas compresses this further. In England and Wales in 2026 both Christmas Day, Friday 25 December, and the Boxing Day substitute, Monday 28 December, are holidays, and New Year’s Day 2027 falls on the Friday. A payment dated Thursday 31 December 2026 requires input on Tuesday 29 December, and the next available entry day after that is Monday 4 January 2027.

For exactly this reason Pay.UK publishes an annual Bacs processing calendar, in landscape and portrait, marking non-processing days and giving the Julian date for every day of the year. The 2026 calendar’s legend distinguishes non-input and non-processing days from days that are a banking holiday in Northern Ireland only, and carries a standing instruction to consult your payment service provider for entry date and processing day arrangements for accounts held in Northern Ireland. That footnote is the honest admission that “English bank working days” is a simplification of a United Kingdom with four sets of holidays, and that where it matters you ask your sponsor rather than a calendar.

If you build software that schedules Bacs submissions, hard-code nothing. Load the processing calendar as data, refresh it annually, and treat “is this a Bacs processing day” as a lookup rather than a computation. Every engineer who has tried to derive English bank holidays algorithmically has eventually met the substitute day rule, the royal proclamation, or a one-off national holiday, and has lost.

■ Bacs approved bureaux and Bacs approved software

Most organisations that use these schemes never submit anything themselves. Pay.UK states that over half of those making Direct Debit, Bacs Direct Credit and Faster Payment submissions do so through approved bureaux rather than directly, and that there are approximately 700 approved bureaux worldwide.

A bureau is defined by the scheme as an organisation that sends payments to Bacs on behalf of another organisation. Pay.UK distinguishes commercial bureaux, which submit on behalf of independent third parties; PSP bureaux, offering participant services direct to service users; in-house commercial group

bureaux, working solely within their own group; and a residual “other” category. The distinction that matters for approval is between in-house and commercial: an in-house bureau submitting only for entities within the same legal organisation need not register, while a commercial bureau submitting for third-party service users that are separate legal entities must achieve Bacs Approved Bureau status. Pay.UK must approve any organisation that submits Direct Debits, Bacs Direct Credits or Faster Payments transactions on behalf of third parties.

The route in runs through a sponsor. An applicant must first be authorised by its sponsoring participant; Pay.UK’s Bureau Inspection team then formalises the process; a one-off registration fee is payable; a questionnaire, supporting documents and a signed agreement follow; an inspection must be passed; and only then does Pay.UK issue a certificate and the sponsor enable the live service. An application not completed within 12 months expires.

Approved bureaux are tiered by annual transaction volume, and the tier sets the inspection cadence.

Tier	Annual transaction volume	Cadence over four years
A	Over 5 million	Full (on-site), interim (on-site), full (on-site), interim (on-site)
B	250,000 to 5 million	Full (on-site) in years 1 and 3; Annual Compliance Statement in years 2 and 4
C	25,000 to 250,000	Full (remote) in year 1; ACS in years 2 and 3; full (remote) in year 4
D	Under 25,000	As tier C

An inspection assesses end-to-end payment processing, from the receipt of Bacs-related information through to the point of transmission to the Bacs clearing system. The questionnaire is structured around ISO/IEC 27001 and covers eight areas: organisation and policy; professional services and commercial arrangements; physical security; network environment; systems management; logical access control; business continuity and disaster recovery; and Bacs processing and data controls. Two annexes attach where relevant: Annex A for third-party risk management where outsourcing is material, and Annex B for hardware security module and cryptographic key management.

An on-site visit typically runs five to seven hours: a business overview, a tour of processing and server areas, a walk through the questionnaire responses with the staff who own them, observation of operations, and review of supporting documents — security policies, process guides, customer agreements, business continuity plans.

Each category is rated Pass, Pass with requirements, Fail, or Not Applicable. A Fail in any single category fails the whole inspection, and a failed bureau must be re-inspected within a maximum of six months. Requirements are mandatory and time-bound; recommendations are best practice, and one ignored may be restated as a requirement next time. The report follows within five working days. Cancelling an inspection with less than ten working days’ notice attracts a fee. Annual fees are levied in April at bureau number level, based on the transactions submitted during the previous calendar year, with new bureaux exempt for their joining year; invoices are payable within 30 calendar days. Terminating BAB status requires three months’ notice.

Approved bureaux receive the BAB logo. Pay.UK is careful about what it means: bureaux awarded it “have been inspected and have met appropriate standards of business, security, technical and operational competence at the time of the inspection”, status “is based upon point in time inspections”, and

Pay.UK “does not make any representation in respect of the suitability, or otherwise, of organisations approved under the scheme for any purpose”. It is an attestation about a date, not a warranty.

Running in parallel is the Bacs Approved Software Service. A Bacs approved software supplier is a developer and supplier providing software, hardware and consultancy services to Bacs Payment System customers to enable them to use the service. Testing under BASS examines authentication of parties communicating with Bacstel-IP or Secure-IP; secure communication, meaning confidentiality and integrity, between those channels and the approved software; correct formatting of submissions and reports; handling of errors and warnings; and trademark usage. Approval requires an application, mandatory testing, training for named staff on the Bacs Payment System and the relevant interfaces, an approval fee and an ongoing annual maintenance fee. Full business and technical specifications, including a sorting code directory, are released to applicants.

The two schemes answer different questions. BASS asks whether a piece of software constructs and transmits submissions correctly. The Bacs Approved Bureau Scheme asks whether an organisation can be trusted to operate the process around that software — who holds the smartcards, what happens when the building is unavailable, how third-party clients are onboarded and offboarded, whether someone is reading the ADDACS advices. A vendor can hold approved software status and not be a bureau; a bureau can be approved while running somebody else’s approved software. An organisation intending to do both applies to both, and should expect the bureau inspection to look hardest at the questionnaire areas that have nothing to do with code.

■ The scale of it, as of writing

Measure	Figure	Period
Bacs payments processed	6.86 billion	2025
Value processed	£6.05 trillion	2025
Direct Debits	5.0 billion, an annual high	2025
Transactions since inception	183.6 billion	1968 to date
Direct Debits over five decades	Over 113 billion	to date
Organisations using Bacs Direct Credit	Over 150,000	as published
Direct participants	32	list updated 7 April 2026
Approved bureaux	Approximately 700	as published

Set that against the rails on either side. Faster Payments settles in seconds, runs 24 hours a day, and carries a maximum of £1 million per transaction under the central scheme limit, with individual institutions imposing lower limits of their own. CHAPS settles same-day, irrevocably, gross, one payment at a time across the Bank of England’s RTGS system, with an early-afternoon cut-off for the same-day guarantee, and carries the large majority of sterling value on a tiny fraction of the volume. Bacs sits between them and does something neither can: it moves enormous numbers of small, scheduled, two-directional payments at a known value date, cheaply, in a format that reconciles itself.

Three days is what that costs. It has been three days since 1968, and the reason it is still three days is not that nobody has thought about it. It is that the three days buy the netting, the netting buys the cost, and the cost is why your salary and your gas bill both arrive without anybody thinking about them at all.

35.98 Common wrong ideas

Wrong: three days means three days. **Right:** it means three English bank working days, so a file submitted on Thursday 27 August 2026 enters accounts on Tuesday 1 September — five calendar days later, and exactly as specified.

Wrong: a Scottish or Northern Irish bank holiday moves the cycle. **Right:** the definition is English working days regardless of where in the United Kingdom the money is going, and Pay.UK's calendar tells you to ask your sponsor about accounts held in Northern Ireland.

Wrong: the money is somewhere in the middle of the process on days one and two. **Right:** nothing is in transit, because what exists between input and entry is a record of an intention, and value moves once, on day three.

Wrong: "submit two days before" is a safe rule of thumb. **Right:** for entry on Tuesday 1 September 2026 the file must be in by 22:30 on Thursday 27 August, because the Friday is the processing day and the Monday is a holiday.

Wrong: entry day is the end of the payment. **Right:** Direct Debits come back through ARUDD and credits through ARUCS one or two working days later, and a collection can be reversed by indemnity claim long after everyone has forgotten it.

Wrong: input day is the day you send the file. **Right:** input day is the latest day you may send it, files may be warehoused up to thirty calendar days ahead, and the scheme's own glossary therefore counts four stages rather than three.

Wrong: Pay.UK runs the Bacs machines. **Right:** Pay.UK owns the rulebook and Vocalink, now part of Mastercard, operates the central infrastructure under contract, so "Bacs went down" and "Bacs says you must" are statements about different organisations.

Wrong: you sign up to Bacs. **Right:** you are vouched for, because a service user is sponsored into the scheme by a payment service provider that is on the hook for its behaviour.

Wrong: the 90 per cent by 06:00 figure is a promise about your file. **Right:** it is a statistical statement about the population of Direct Credit payments, and the only commitment is that paying PSPs apply entries no later than 09:30.

Wrong: the Standard 18 offsets in a vendor or bank guide are authoritative. **Right:** the definitive specification is the Service User's Guide and Rules to Bacstel-IP, issued by your sponsor and not public, and a sponsor may impose requirements of its own.

35.99 Chapter summary in 20 lines

1. Bacs has run since 1968, moved 6.86 billion payments worth £6.05 trillion in 2025, and pays most of the salaries and collects most of the household bills in Britain.
2. It behaves like a sorting office: lists arrive by 22:30 on day one, are sorted on day two, and money moves on day three for everybody at once.
3. The clever part of day two is not sorting but subtraction, because millions of instructions are netted down to a few dozen interbank figures.
4. One rail carries two products in opposite directions: Direct Debit pulls under a standing authority, and Bacs Direct Credit pushes.
5. The three days are English bank working days, so a payroll dated Tuesday 1 September 2026 must be submitted by 22:30 on Thursday 27 August.

6. Nothing is in transit during the cycle, because the payer's money stays in the payer's account and fully usable until entry day.
7. On entry day two different things happen in two kinds of money: customer accounts move in commercial bank money, and the netted interbank positions settle at 09:30 in central bank money at the Bank of England.
8. Prefunding removes the credit risk that deferral would otherwise create, because every direct participant holds a segregated cash deposit covering its own net position.
9. Entry day is not the end, since ARUDD, ARUCS, ADDACS and AWACS all arrive afterwards and an indemnity claim can reverse a collection much later still.
10. Pay.UK owns the rules, Vocalink runs the infrastructure, thirty-two direct participants clear and settle, and everyone else reaches Bacs through an agency arrangement or a sponsor.
11. Input day is the latest day of submission rather than the only one, and warehousing up to thirty calendar days ahead is why the glossary describes arrival, input, processing and entry.
12. Standard 18 is a flat, fixed-width text file inherited from magnetic tape, with 80-character labels wrapped around 100-character payment records and no delimiters of any kind.
13. Every field is found by counting characters: amounts are eleven digits of pence with no decimal point, and names and references are eighteen characters and no more.
14. The file must balance, because a contra record mirrors the payment records against the service user's own account, which is why one bank statement line reconciles to one file.
15. Transaction codes carry meaning the amount field does not — 01, 17, 18 and 19 for collections, 99 for credits, and 0N, 0S and 0C for AUDDIS instructions.
16. A Service User Number is six digits, is the scheme's identity for a sponsored originator, and is what a paying bank matches against and eventually attributes blame to.
17. Files travel over Bacstel-IP using approved software and non-transferable smartcards, while reports are collected from the Payment Services Website and must be actioned within three working days.
18. Bank holidays do not delay the cycle by a day; they remove a slot, and where the holiday falls decides whether that costs one day or four.
19. Anyone submitting for third parties must hold Bacs Approved Bureau status, inspected on a cadence set by volume tier, and a Fail in any single category fails the whole inspection.
20. Three days is what the netting costs, the netting is what makes the payments cheap, and the cheapness is why a salary and a gas bill both arrive without anybody thinking about them.

Sources: Pay.UK and Bacs published scheme material (bacs.co.uk processing cycle, glossary, processing calendar 2026 and The Little Bacs Book; wearepay.uk Bacs Payment System pages, participant list, Bacs Approved Bureau Scheme supporting guidelines and Bacs Approved Software Service), Pay.UK's PFMI self-assessment report, the Bank of England's payment and settlement pages, GOV.UK bank holidays, and CIPP guidance on the withdrawal of the RTI Bacs hash. Standard 18 record layouts cross-checked against independent implementation documentation; the authoritative specification is the Service User's Guide and Rules to Bacstel-IP, issued by sponsoring PSPs. All timings, limits and figures accurate at time of writing.

The Mandate

36.0 What this chapter gives you

1. You will be able to say what a Direct Debit Instruction legally is, and why it is an authority given to the service user rather than a contract with the payer's bank.
2. You will be able to state what the Direct Debit Guarantee actually promises, who makes the promise, and who ends up paying for it.
3. You will be able to trace the money in a disputed collection from the payer's immediate refund, through the paying bank, to the indemnity claim that debits the service user's own account.
4. You will be able to count an advance notice period correctly and say what the notice must contain if it is to win a challenge.
5. You will be able to choose between paper, AUDDIS and Paperless Direct Debit for setting up a mandate, and say what obligations each route carries.
6. You will be able to read a DDICA reason code and say whether the claim is about notice, cancellation, authority or the name on the statement.
7. You will be able to work out the challenge deadline from the day a claim is reported, and say what happens at the end of working day nine and working day eleven.
8. You will be able to explain why the Guarantee is faster and broader than the Payment Services Regulations 2017, and where regulation 79(3) does and does not apply.
9. You will be able to explain why cancelling a Direct Debit cancels nothing else, and why a refunded payer may still owe the money.
10. You will be able to name the design choices that keep a collections product alive: storing the notice document, actioning ADDACS within three working days, treating day nine as code rather than a runbook, and putting the trading name in field 9.

Every other payment in this book is something you *do*. You tap a card, you press send, you hand over a note. A Direct Debit is something you *stop doing*. You do it once, and thereafter the money leaves without you touching anything, on dates you may not remember, in amounts you did not choose, for as long as neither side gets round to stopping it.

That is an extraordinary thing to hand to a stranger, and roughly five billion times a year in the United Kingdom people hand it over without a second thought. In 2025, Bacs processed 5.03 billion Direct Debits — a record for the system — as part of 6.86 billion Bacs payments worth £6.05 trillion in total. Nine in ten UK adults with a bank account have at least one Direct Debit running.

They do it because of a single short document that almost nobody reads: the Direct Debit Guarantee, five bullet points long, printed on the tear-off half of every mandate form and reproduced on every biller's website. It is the most generous consumer protection attached to any sterling payment rail.

It is more generous than the statutory rights that sit underneath it. And it is generous in a way that costs the person collecting the money, not the person paying it, which is why it works and why it is expensive to abuse.

The previous chapter covered the three-day cycle that carries the collection. This chapter covers the authority that permits the collection in the first place: what a Direct Debit Instruction legally is, how one is created on paper, by file, or over the telephone, what notice must be given before money moves, what happens when the payer objects, and — the question that decides whether a collections business survives — who ends up out of pocket.

The plain version

Imagine you have an account at the corner shop.

The old way of settling it was that you walked in every month and paid what you owed. That works, but it means the shopkeeper has to chase you when you forget, and you have to remember. So you try something else. You write a short note and give it to the shopkeeper. The note says: *take what I owe you, out of my money-box, on the first of every month*. You sign it. The shopkeeper keeps it.

Two things about that note matter more than anything else in this chapter.

The first is that the note does not say an amount. You have not said “take £30”. You have said “take what I owe you”. The shopkeeper decides how much, and when, within whatever you agreed about the goods. That is what makes the arrangement useful — your bill genuinely does change month to month — and it is also what makes it frightening.

The second is who the note is addressed to. It looks as though you are giving permission to the shopkeeper. What you are actually doing is telling your *bank* that the shopkeeper is allowed to reach in. The bank is the one holding the money-box. It will honour requests from anybody you have named on a note like this, and it will not ring you up first.

■ The promise that makes it safe

If that were the whole arrangement, nobody sensible would sign the note. So there is a second document, and it is the reason Direct Debit is trusted in Britain in a way that no equivalent arrangement is trusted in most other countries.

Your bank makes you a promise. It is short enough to say in full, and it is worth saying in full, because the exact words are the whole point:

If there are any changes to the amount, date or frequency of your Direct Debit, the organisation will notify you — normally ten working days — in advance of your account being debited, or as otherwise agreed. If an error is made in the payment of your Direct Debit, by the organisation or by your bank or building society, you are entitled to a full and immediate refund of the amount paid from your bank or building society. If you receive a refund you are not entitled to, you must pay it back when the organisation asks you to. You can cancel a Direct Debit at any time by simply contacting your bank or building society.

Read the second sentence again, slowly. *You are entitled to a full and immediate refund from your bank*. Not from the shopkeeper. Not after an investigation. Not if the bank agrees with you. You go to your own bank, you say an error has been made, and the money comes back.

There is no minimum amount. There is no stated deadline in the promise itself. There is no requirement to argue with the shopkeeper first. The bank refunds you and then goes and sorts it out with the shopkeeper afterwards, at its own inconvenience.

■ Where the money actually comes back from

Here is the part people miss. The bank refunds you out of its own till, but it does not stay out of pocket. It turns round and takes the money back off the shopkeeper — directly out of the shopkeeper's own bank account, without asking the shopkeeper first.

So the sequence is: shopkeeper takes £39.99 from you; you object; your bank immediately gives you £39.99; your bank then debits £39.99 from the shopkeeper. The shopkeeper can dispute that, with evidence, within a short window. If the shopkeeper does nothing, the money simply goes.

That is the entire economics of the Direct Debit Guarantee in one paragraph. The person collecting the money carries the risk of being wrong. Not the payer, and not, in the end, the bank.

■ A worked example, with real numbers

Northgate Leisure runs a gym. It collects £34.99 a month by Direct Debit from 4,100 members on the first working day of each month.

In September 2026 it decides to raise the price to £39.99 with effect from the October collection, which falls on Thursday 1 October 2026.

The Guarantee's default notice period is ten working days. Counting back ten working days from Thursday 1 October gives Thursday 17 September, and the notice has to *arrive* by then, so in practice Northgate posts and emails its members in the second week of September. The notice says the new amount, £39.99, and the date, 1 October, and it says who is collecting it in the name members will actually recognise on their statements.

On 1 October, 4,100 collections of £39.99 land. Total: £163,959.

One member, Mr Hale, does not remember any letter. On Monday 5 October he telephones his bank and says the gym has taken the wrong amount. His bank does not ring the gym. It does not ask him for the letter. Under the Guarantee it credits £39.99 back to his account that day.

His bank is now £39.99 down, so it raises an indemnity claim against Northgate Leisure. The claim reaches Northgate through a report on Tuesday 6 October. That day counts as working day one. Northgate has until the end of working day nine — Friday 16 October 2026 — to challenge it, and the only thing that will win the challenge is a copy of the advance notice showing the amount, the date and the address it was sent to.

If Northgate can produce it, the claim is cancelled and Northgate keeps the £39.99. If Northgate cannot produce it, or simply does not act in time, £39.99 comes out of Northgate's bank account and Mr Hale keeps his refund.

And here is the sting in the tail: Mr Hale still owes the gym for October. The refund undid the *payment*. It did not undo the *membership*.

■ Why anybody collects money this way

Given all that, why would Northgate use Direct Debit at all?

Because the alternative is worse. Cards expire, get reissued and get declined, and every failure needs a human being to chase it. With Direct Debit, Northgate knows on a fixed date that £163,959 will arrive

in one lump, it reconciles only the exceptions, and the bill collects itself for years without either party doing anything. The Guarantee is the price of that, and for a business genuinely giving correct notice and taking correct amounts the price is low, because most claims never happen: most collections are right.

The Guarantee is not a tax on collecting money. It is a tax on collecting money carelessly.

Where the plain version stops being true

The note is not addressed to your bank, and the Guarantee is not law. The plain version says you are telling your bank that the shopkeeper may reach in. Legally, a Direct Debit Instruction is an authority given by the payer *to the service user* — Bacs' own definition is "an authority (paper, AUDDIS or PDD) from the customer allowing the service user to collect Direct Debit payments from their account" — and the Bacs guidance for anyone producing a mandate states flatly that "the DDI is not evidence of any contract between you and your customer's payment service provider". A copy is lodged with the paying bank so that it knows the collection is expected, but the mandate is not an agreement with that bank. Nor is the Guarantee a statute. It is a scheme rule: a promise that Pay.UK's Bacs rulebook obliges every participating bank and building society to offer, enforced through contracts between scheme participants rather than through an Act of Parliament. Underneath it sit genuine statutory rights, in the Payment Services Regulations 2017, and those rights are in several respects *narrower* than the Guarantee. The technical section sets them side by side.

"Immediate" describes one side of the transaction only. The payer's refund is immediate. Everything on the other side runs on working-day clocks: the claim reaches the collecting organisation through a batch report, the challenge window is counted in English bank working days, and the paying bank's response has its own cut-off. A payer refunded on a Monday afternoon may be looking at a business whose money does not actually move for another fortnight. Cash-flow models that treat indemnity claims as instantaneous are wrong in both directions.

Almost no modern mandate is signed. The corner-shop image has a signed note, kept by the shopkeeper. In 2026 most new Direct Debits are set up over a website or a telephone with no signature anywhere in the process — Paperless Direct Debit — and the mandate travels to the payer's bank as an electronic record through AUDDIS. That shifts a large burden. When there was a signature, the paying bank checked it. When there is not, the collecting organisation is responsible for establishing that the person giving the account details is entitled to give them, and it carries the loss if they were not. Modulus checking, which most systems do and which the scheme recommends, confirms only that a sort code and account number *could* exist together. It says nothing whatever about who owns the account.

Cancelling the Direct Debit does not cancel anything else. Telling your bank to stop paying the gym stops the payments and nothing more. The membership contract survives, the debt keeps accruing, and the gym is entitled to pursue it. The Guarantee is explicit that it does not affect the contract between payer and biller: you can claim the money back and still owe it. In the other direction, a mandate that nobody uses does not last forever either — paying banks remove dormant instructions after a set period, so a service user that stops collecting for a year or more may find the authority has quietly disappeared.

The Guarantee covers the mechanics, not the merits. It protects you if the amount, date or set-up was wrong. It does not protect you if the gym is terrible, if the goods never arrived, or if the company goes into administration owing you nine months of prepaid membership. Those are contract disputes, and the route for them is the biller, then the biller's complaints process, then the courts — not the bank.

Section 75 of the Consumer Credit Act and card chargebacks, which do reach into the merits, do not apply to Direct Debit at all.

The technical version

■ The parties, and what each of them signs

Five roles matter, and confusing any two of them produces most of the bad documentation on this subject.

The **payer** is the account holder whose account is debited. The **service user** is the organisation collecting, identified by a six-digit **Service User Number (SUN)** allocated by its bank. The **paying payment service provider (paying PSP)** holds the payer's account, lodges the instruction on its records, applies the collection and, critically, is the party that owes the Guarantee. The **sponsoring PSP** is the bank that authorises the service user to submit to Bacs at all, vets it, sets its limits and stands behind it. And **Pay.UK** is the scheme operator and rule-maker, with Vocalink as the infrastructure provider that physically processes the files.

The service user's own signature appears on a document the payer never sees: the **indemnity**. Bacs defines it as "a document under the Direct Debit Scheme that contains a legally binding undertaking to make payment to any paying PSP in response to an indemnity claim", and adds that "an indemnity, in standard form, is an essential requirement of the Direct Debit Scheme". That sentence is the legal engine of everything in this chapter. The Guarantee is a promise the paying bank makes to the payer; the indemnity is a promise the service user makes to every paying bank in the country, in advance, that it will make good whatever the Guarantee costs them.

Against the SUN, the scheme maintains a record on the **Direct Debit Originator database (DDO)**: contact details, AUDDIS status, dormancy period, and the service user's agreed advance notice period. That last field is not decorative. It is the number a paying bank looks at when deciding whether a claim about insufficient notice is well founded.

■ The mandate document

The Direct Debit Instruction is a controlled document. Bacs publishes ten standard templates in A4, A5 and small formats, with and without an "official use" box, and the rule is that all DDIs must follow all of the wording and the order of the information in those templates. The design must be approved in writing by the sponsoring PSP before use, and again after any change. The Direct Debit logo is mandatory on every DDI.

The instruction half of the standard form carries these fields: the service user's name and address, the six-digit Service User Number, the name or names of the account holder or holders, the branch sort code, the bank or building society account number, the name and full postal address of the payer's bank or building society, a reference, a signature and a date. The operative sentence is fixed:

Please pay [service user] Direct Debits from the account detailed in this Instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this Instruction may remain with [service user] and, if so, details will be passed electronically to my Bank/Building Society.

The Guarantee sits on a detachable portion, headed with the instruction that it "should be detached and retained by the payer".

Several operational rules follow from the form and are routinely broken by first-time service users. The customer must not add amount or date restrictions to the mandate; if they have, the mandate is not valid as submitted. The mandate may only be lodged with the paying PSP by the service user,

never posted in by the customer. Tear-off slips carrying extra marketing or account information must be detached before the form is sent onward. And when a mandate is submitted electronically the amount field must be zero — every Direct Debit Instruction is for an unlimited amount, and a non-zero value produces an AUDDIS rejection with reason code 7, “DDI amount not zero”.

■ Three ways to create a mandate

Non-AUDDIS (paper). The historic method. The payer signs a paper DDI, the service user sends the physical form to the payer’s bank branch or to that bank’s central lodgement address, and the bank keys it in. Slow, error-prone, and shrinking. Pay.UK has announced that from 1 July 2027 no new non-AUDDIS Service User Numbers will be created by PSPs. Existing non-AUDDIS SUNs may continue to lodge new paper instructions and to collect against existing ones, but any SUN created after that date will be electronic only. Note also that a service user cannot mix the two on one number: separate SUNs are required for AUDDIS and non-AUDDIS instructions.

AUDDIS (Automated Direct Debit Instruction Service). The service user keeps the signed original and transmits the mandate details electronically through Bacs to the paying PSP, which validates them and, if it accepts, sets up the instruction on its database. AUDDIS is mandatory for all new service users that submit directly to Bacs. Three transaction codes carry the traffic:

Code	Meaning
0N	New Direct Debit Instruction lodged
0S	Conversion of an existing manually lodged DDI to AUDDIS
0C	Cancellation of an existing DDI by the service user

Rejections come back through the Bacs messaging service and appear on the AUDDIS Bank Rejected Direct Debit Instruction Report. Bacs states that notification of any rejected instruction will reach the service user within five working days of the instruction being sent to the paying PSP. Some failures — a failed modulus check, for instance — are caught by Bacs before the file ever reaches the paying bank and are returned immediately.

Paperless Direct Debit (PDD). Sign-up by telephone, website, telephone keypad, face to face or interactive television, with no signature at any point. PDD is available only to service users who already have live AUDDIS status and who satisfy additional criteria, and it carries a distinct set of obligations:

- The service user, not the bank, is responsible for verifying the customer’s identity, address and account details. Bacs publishes a *Direct Debit verification measures* document listing acceptable methods.
- Modulus checking at the point of sale is strongly recommended. It does not verify that the account belongs to the payer.
- Every telephone script, internet screen, confirmation letter and item of supporting material must be approved by the sponsoring PSP before use, and again after any change.
- Two members of staff must successfully complete the mandatory Paperless Direct Debit training course.
- One type of application per SUN is recommended, so that different sign-up channels can be monitored and controlled separately.
- Voice recording of telephone sign-ups is recommended.
- Service users may be subject to annual review of their PDD processes, and failure to comply with the scheme rules may result in removal of sponsorship and exclusion from the Direct Debit Scheme altogether.

PDD carries its own confirmation rule, and it is one of the most frequently misquoted requirements in the scheme. The service user must send confirmation to the customer within three working days of the verbal or internet instruction. Alternatively — and only if the first collection is made within a month of the Direct Debit being set up — the confirmation may incorporate the advance notice, in which case it must be received by the customer no later than the service user's advance notice period before the first collection.

Two further timing rules apply to all mandates regardless of channel. All Direct Debit Instructions should be submitted to the paying PSP within ten working days of the customer signing, even if the first collection is months away. And if a paying PSP asks to see a copy of a mandate, the service user must provide it within seven working days of the request; failure to do so exposes the service user to liability for every Direct Debit collected under that instruction. For a PDD service user there is no signed copy to produce, which is precisely why a disputed authority nearly always ends in a refund.

Mandates also expire from disuse. Bacs' published material is not internally consistent here, and practitioners should check the current rulebook rather than rely on either figure quoted second-hand: the Bacs glossary describes a default dormancy period of thirteen months, noting that this was temporarily increased to twenty-four months to limit the impact of pandemic payment holidays, while the Bacs FAQ states that paying PSPs hold DDI details for a minimum of twenty-four months from receipt or last payment. An extended dormancy period may be applied for through the sponsoring PSP at the point of joining AUDDIS. Whichever number applies, the operational consequence is the same: a service user collecting annually, or seasonally, must know its dormancy setting or its mandates will evaporate between collections.

■ Advance notice

Advance notice is the single most important operational obligation in the scheme, because it is the obligation that most indemnity claims allege was breached.

The default period, in the absence of any other agreement, is a minimum of ten working days plus postal time. Working days for scheme purposes means English bank working days, excluding Saturdays, Sundays and public holidays, regardless of where in the United Kingdom either party sits. The period is recorded against the SUN on the Direct Debit Originator database, and it can be shortened by agreement with the payer — many subscription businesses run on three or five working days, stated in their terms — but the shortened period must be agreed, documented and reflected on the DDO, not simply adopted.

Notice is required in two circumstances. Before the first collection under a new mandate. And before any change to the amount, date or frequency of an existing one. Note what is not on that list: notice is not required before an unchanged, already-notified collection, which is why an annual notice covering a fixed schedule is sufficient for a fixed-amount arrangement.

The notice must reach the payer — the account holder, not a spouse, not an account administrator, not a purchasing department that happens to have signed the contract. It must state the amount and the date. It must identify the service user in the name the payer will actually see on the bank statement, which for organisations trading under a different name from their registered name means using the trading name that appears in field 9 of the Bacs record. And the notice document itself, like every other piece of Direct Debit literature, requires written approval from the sponsoring PSP before use.

Where the Guarantee says "or as otherwise agreed", it is also worth noting the sentence that follows in the published wording: if the payer *requests* the organisation to collect a payment, confirmation of the amount and date is given at the time of the request rather than in advance. That is the clause that

makes one-off, payer-initiated collections workable, and it is the clause that subscription businesses most often try to stretch beyond its meaning.

Failures of advance notice surface in three different message streams, which is a good indication of how seriously the scheme takes it: a collection returned unpaid with ARUDD reason code 4, “advance notice disputed”; a mandate amended or cancelled with ADDACS reason code D, “advance notice disputed”; and indemnity claims under DDICA reason codes 1 and 2. Bacs publishes best-practice material specifically aimed at reducing indemnity claim volumes, and clear, timely, correctly addressed advance notice is the first item on it.

■ The Guarantee, clause by clause

The authoritative text, as published by Bacs and reproduced on every compliant mandate, is five bullets. Taken one at a time:

“This Guarantee is offered by all banks and building societies that accept instructions to pay Direct Debits.” The promisor is the payer’s own bank, not the biller and not Pay.UK. A payer never has to deal with the collecting organisation to obtain a refund, and a bank cannot decline on the basis that it disagrees with the payer’s account of events.

“If there are any changes to the amount, date or frequency of your Direct Debit [the organisation] will notify you (normally 10 working days) in advance of your account being debited or as otherwise agreed. If you request [the organisation] to collect a payment, confirmation of the amount and date will be given to you at the time of the request.” The notice obligation, restated to the payer as a right. The parenthesis is why “ten working days” is universally quoted as the rule despite being a default rather than a fixed requirement.

“If an error is made in the payment of your Direct Debit, by [the organisation] or your bank or building society, you are entitled to a full and immediate refund of the amount paid from your bank or building society.” The operative promise. Note the breadth: it covers errors by either party, it is full rather than partial, and it is immediate. There is no stated time limit, no minimum value, and no condition that the payer exhaust the biller’s complaints process first.

“If you receive a refund you are not entitled to, you must pay it back when [the organisation] asks you to.” The clawback. Practically, this is the clause that makes the whole design tolerable to service users: a refund obtained wrongly is recoverable as a debt, and a service user with good records can pursue it.

“You can cancel a Direct Debit at any time by simply contacting your bank or building society. Written confirmation may be required. Please also notify [the organisation].” Cancellation is unilateral, immediate and requires no justification.

■ How that compares with the statutory position

It is worth being precise about the relationship between the Guarantee and the Payment Services Regulations 2017 (SI 2017/752), because a good deal of published commentary treats them as the same thing.

Under regulation 67, a payment transaction is authorised if the payer has consented to its execution; a Direct Debit Instruction is that consent, given in advance and for an unspecified amount.

Regulation 79 then gives the payer a right to a refund from its own PSP for an authorised transaction initiated by or through the payee, but only where two conditions are both met: the authorisation did not specify the exact amount, and the amount exceeded what the payer could reasonably have expected given previous spending, the framework contract and the circumstances. Regulation 80 requires the

payer to make that request within eight weeks of the debit, and requires the PSP either to refund or to justify its refusal within ten business days.

Regulation 79(3) does provide an *unconditional* refund right, but it applies to euro direct debits within the scope of Regulation (EU) 260/2012 — the SEPA regulation — and not to sterling Bacs Direct Debits. This is the point most often got wrong: for a sterling Direct Debit, the unconditional right comes from the scheme Guarantee, not from statute.

Regulation 74 sets the long-stop: a payment service user loses the right to redress for an unauthorised or incorrectly executed transaction unless it notifies its PSP without undue delay and in any event no later than thirteen months after the debit date. Where a transaction genuinely was unauthorised — a mandate the payer never gave — regulation 76 requires the PSP to refund as soon as practicable and in any event no later than the end of the business day following the day it becomes aware, subject to a suspected-fraud exception.

Set side by side:

	Direct Debit Guarantee	PSRs 2017
Source	Bacs scheme rules (Pay.UK)	SI 2017/752
Trigger	Any error in the payment of a Direct Debit	Unauthorised transaction, or authorised but unexpectedly large
Conditions	None stated	Reg 79(2) reasonableness test for authorised transactions
Speed	Full and immediate refund	Reg 76: end of next business day (unauthorised); reg 80: 10 business days to refund or justify
Deadline to claim	None stated	Reg 80: 8 weeks; reg 74: 13 months long-stop
Who refunds	Payer's own bank or building society	Payer's own PSP

The practical consequence is that a UK Direct Debit payer almost never needs to invoke the statute. The Guarantee is faster, broader and unconditional, and banks apply it as a matter of course. Where a bank refuses, the payer's route is the bank's own complaints procedure and then the Financial Ombudsman Service.

■ Cancellation, and what it generates

A payer cancels by contacting the bank or building society — increasingly by switching a toggle in a banking app — and the scheme's own consumer guidance adds two pieces of practical advice: written confirmation may be required, and the cancellation should be made at least a full working day before a collection is due, because a collection already inside the three-day cycle may be too far along to stop. If money is nonetheless taken after cancellation, that is squarely within the Guarantee.

Cancellations and amendments reach the service user through **ADDACS**, the Automated Direct Debit Amendment and Cancellation Service. The reason codes a service user will actually see include:

Code	Meaning
0	Instruction cancelled — refer to payer
1	Instruction cancelled by payer

Code	Meaning
2	Payer deceased
3	Account transferred to a new bank or building society (new details supplied; new DDI required)
B	Account closed
C	Account transferred within the same bank or building society (no new DDI required)
D	Advance notice disputed
E	Instruction amended
R	Instruction reinstated

ADDACS advices must be actioned immediately, or within three working days of receipt. That deadline is a scheme rule, not a courtesy, and repeated failure to observe it is a compliance matter with the sponsoring PSP. Reports are available on the Payment Services Website for a limited retention period — most Bacs reports for 37 calendar days, ADDACS and AWACS for 60 calendar days — after which they are gone from the service user's view.

Cancellation from the service user's side uses transaction code 0C, and the scheme also provides a cleaner mechanism for a planned ending: the collection transaction codes distinguish the first collection under a mandate (01), a subsequent collection (17), a re-presentation of a previously unpaid item (18), and the final collection (19). Submitting a collection as code 19 tells the paying PSP to delete the instruction from its records, and may produce a "final collection" annotation on the payer's statement. Using it correctly avoids the common failure mode of a mandate left live for years after the last payment.

One structural point about cancellation deserves emphasis, because it explains a recurring category of indemnity claim. A payer who cancels with their bank and never tells the service user leaves that service user submitting collections against an instruction that no longer exists. Those collections will be returned unpaid, and if the payer is annoyed enough they may also generate a claim. The mirror-image case — the payer who cancels with the service user but not the bank, and the service user who fails to submit a 0C — leaves a live authority nobody intends to use.

■ Indemnity claims

An indemnity claim is defined by Bacs as "a claim made by the paying PSP in respect of an incorrect Direct Debit being applied to an account". It is the mechanism by which the Guarantee's cost is passed from the paying bank to the service user, under the standing indemnity the service user signed when it was sponsored.

Claims arrive as **DDICA** advices in the service user's Bacs reports, carrying one of eight reason codes:

Code	Reason
1	Amount and/or date of Direct Debit differs from advance notice
2	No advance notice received by payer, or the amount is disputed
3	DDI cancelled by paying bank
4	Payer cancelled DDI direct with the service user
5	No instruction held; payer disputes having given authority

Code	Reason
6	Signature is fraudulent or not in accordance with authorised signatory
7	Claim raised at the service user's request after the Direct Debit was applied to the payer's account
8	Service user name disputed; payer does not recognise the service user collecting the Direct Debit

Codes 1 and 2 are advance notice failures. Codes 3 and 4 are cancellation failures. Codes 5 and 6 are authority failures, and code 6 is the reason a Paperless Direct Debit service user must have documented verification: with no signature to produce, the claim is close to unanswerable. Code 8 is the quiet killer of subscription businesses trading under names their customers have never heard of.

Code 7 is different in kind. It is the route by which a service user raises a claim *on behalf of* a payer who has come to it directly asking for money back — a refund the service user has decided to make. Rather than sending the customer to their bank, the service user initiates the reversal itself through the Payment Services Website, and may enter details for up to twelve Direct Debits for a single payer in one request.

The challenge process. Since 4 March 2024 challenges are raised centrally through the Payment Services Website rather than by email to each individual paying bank; the old bank-by-bank process was closed on 31 May 2024. The timetable, as published in the Bacs DDIC Challenge Process service user guide and accurate at time of writing, works like this. The day the claim is raised and the report made available counts as working day one. The service user must raise any challenge before the end of working day nine; a claim neither accepted nor challenged by that point is automatically accepted, and the money leaves. Up to five pieces of supporting evidence may be attached — typically a copy of the Direct Debit Instruction, a copy of the advance notice, or a call recording. The paying bank then reviews and either accepts or rejects; a rejected challenge may be relodged with better evidence, still within the day-nine window. If the paying bank has neither accepted nor rejected a challenge by the end of working day eleven, the claim is automatically cancelled in the service user's favour.

Older vendor documentation frequently quotes a fourteen-working-day window. That figure belongs to the pre-2024 arrangements; the operative deadline today is the end of working day nine. The authoritative text is the *Service User's Guide and Rules to the Direct Debit Scheme* — version 5.9 at the time of writing — available to registered users on the Bacs site, and it is what a sponsoring bank will hold a service user to.

Four supporting reports are available through the Payment Services Website: a monthly challenge and response report summarising claims raised, challenged, accepted and rejected; an ad hoc challenge and response audit report produced within fifteen minutes of request; an in-progress challenge and response report; and an in-progress report for reason code 7 requests.

Refund requests are not indemnity claims. The scheme distinguishes the two. A refund request is a claim by a paying PSP for money back where the error was the paying PSP's own — for example, it failed to action a cancellation and paid the collection anyway — and therefore falls outside the valid indemnity claim criteria. Settlement of a refund request is at the discretion of the service user. A service user that treats every incoming request as an obligation is paying for other people's mistakes.

■ Who bears the loss, in order

Putting the sequence in one place, because it is the answer to the question everyone actually has:

1. The payer bears no loss. The Guarantee makes them whole immediately, from their own bank, without conditions.
2. The paying PSP bears the loss temporarily. It funds the refund out of its own money for the days it takes an indemnity claim to run.
3. The service user bears the loss finally, in the ordinary case. The claim debits its bank account under the standing indemnity, whether or not it agrees, unless it challenges successfully inside the window.
4. The sponsoring PSP bears the loss if the service user cannot. That is the whole reason sponsorship exists, and the whole reason it involves financial assessment, account limits, document approval, mandatory training and the power to withdraw. A sponsor that lets an insolvent or careless service user run up claims is buying those claims.
5. The payer may still owe the underlying debt. A successful claim reverses the payment and leaves the contract intact. If the service is delivered and the contract is valid, the service user retains its ordinary legal remedies for the debt, including debt collection and the county court. The bank's decision to refund is an administrative act under scheme rules, not a judgment on the contract.

That ordering is why the Guarantee can be as generous as it is. The cost of generosity is loaded onto the party best placed to prevent the error, and that party has posted a legally binding undertaking, is vetted annually, and can be removed from the scheme. Fraudulent set-up is a good illustration: Bacs' own guidance states that where a payment is made against a fraudulent signature, the service user is liable, but the paying PSP must refund the customer immediately and then reclaim through an indemnity claim. The customer's position is never in question. Only the allocation between the two institutions is.

■ The scale of it, as of writing

Direct Debit is not a niche instrument, and the Guarantee therefore stands behind a very large fraction of British household finance. Pay.UK's published Bacs statistics for 2025 record 5,029,499 thousand Direct Debits — 5.03 billion — from a Bacs total of 6.86 billion payments worth £6.05 trillion. Over 113 billion Direct Debits have been collected since the mechanism began, out of more than 183.6 billion Bacs transactions since 1968.

The 2025 breakdown gives a sense of what the mandate actually carries:

Sector	Direct Debits, 2025 (thousands)
Utilities and household bills	1,840,984
Insurance premiums	930,489
Subscriptions	562,919
Business to business collections	136,344

Within those, the individual lines are instructive: 648 million mobile telephone collections, 398 million general insurance, 304 million domestic fuel, 232 million council tax, 188 million water, 157 million television licences and 149 million road tax. These are not discretionary purchases. They are the fixed costs of running a household, and they are collected on an authority the payer can revoke in ten seconds and reverse without giving a reason.

■ A note for implementers

If you are building a collections product, the following is the short list that separates systems that survive their first year from systems that do not.

Store the advance notice, not just the fact that one was sent. A challenge under DDICA code 1 or 2 is won or lost on producing the document, its amount, its date and its delivery address. A boolean flag is worthless.

Reconcile ADDACS within three working days, automatically. Every code 1, 3, B and C changes what you may submit next month, and the three-working-day rule is a scheme obligation with a compliance consequence.

Treat day nine as a hard deadline in code, not in a runbook. Claims arrive in batch reports, weekends and English bank holidays do not count, and an unattended inbox in the first fortnight of January will cost real money.

Use transaction code 19 for final collections and submit 0C when a customer leaves. Live mandates against departed customers are how code 4 claims are born.

If you are running Paperless Direct Debit, treat verification as the product, not the paperwork. You have no signature. Under code 5 or code 6 you will be asked to prove authority within seven working days, and if you cannot, you pay. Voice recordings, timestamped screens, address matching and audit trails are your entire defence.

And put your trading name — the one the customer will recognise — in field 9. Code 8 claims are almost always self-inflicted.

The mandate is a small document. It authorises an indefinite series of payments of unspecified amounts, and it is backed by a promise that the payer can call in at any time, for any reason, and get their money back the same day. That combination sounds unworkable, and it would be, were it not for the fact that the entire cost of being wrong sits with the one party who can choose not to be. That is the design. It has held for over fifty years, across five billion collections a year, and it holds precisely because nobody collecting money by Direct Debit can afford to be casual about it.

36.98 Common wrong ideas

Wrong: the mandate is an instruction you give your bank. **Right:** it is an authority given by the payer to the service user, and Bacs states plainly that the DDI is not evidence of any contract between the service user and the payer's payment service provider.

Wrong: the Direct Debit Guarantee is law. **Right:** it is a scheme rule in Pay.UK's Bacs rulebook, enforced through contracts between participants, and the statutory rights underneath it in the PSRs 2017 are in several respects narrower.

Wrong: the unconditional refund right for a sterling Direct Debit comes from the Payment Services Regulations. **Right:** regulation 79(3)'s unconditional right applies to euro direct debits under Regulation (EU) 260/2012, and for sterling Bacs it is the scheme Guarantee that is unconditional.

Wrong: "immediate" describes the whole transaction. **Right:** only the payer's refund is immediate, while the claim, the challenge window and the paying bank's response all run on English bank working days, so the service user's money may not move for another fortnight.

Wrong: modulus checking confirms the customer owns the account. **Right:** it confirms only that a sort code and account number could exist together and says nothing whatever about ownership, which is why a Paperless service user must verify identity itself.

Wrong: the challenge window is fourteen working days. **Right:** that figure belongs to the arrangements closed in 2024, and the operative deadline is the end of working day nine, after which an unchallenged claim is automatically accepted and the money leaves.

Wrong: a refund under the Guarantee cancels what you owe. **Right:** the refund undoes the payment and not the contract, and the service user keeps its ordinary legal remedies for the debt.

Wrong: the Guarantee covers you if the goods are bad or the company collapses. **Right:** it covers the mechanics — amount, date and set-up — and not the merits, and Section 75 and card chargebacks, which do reach the merits, do not apply to Direct Debit at all.

Wrong: every claim a paying bank raises must be paid. **Right:** a refund request, where the error was the paying PSP's own, falls outside the valid indemnity claim criteria and is settled at the service user's discretion.

Wrong: an unused mandate stays live indefinitely. **Right:** paying PSPs remove dormant instructions after a set period, so a service user collecting annually or seasonally must know its dormancy setting or find the authority has quietly evaporated.

36.99 Chapter summary in 20 lines

1. A Direct Debit is the one payment you set up once and then stop doing, and Britain ran 5.03 billion of them in 2025 out of 6.86 billion Bacs payments worth £6.05 trillion.
2. People hand over that authority because of the Direct Debit Guarantee, five bullet points promising a full and immediate refund from the payer's own bank if an error is made.
3. The refund is unconditional in practice: no minimum amount, no stated deadline, and no requirement to argue with the biller first.
4. The bank does not stay out of pocket, because it raises an indemnity claim and debits the collecting organisation's own account.
5. That is the entire economics of the scheme — the party collecting the money carries the risk of being wrong.
6. Legally the mandate is an authority given by the payer to the service user rather than a contract with the payer's bank, and the Guarantee is a scheme rule rather than a statute.
7. Five roles matter: the payer, the service user with its six-digit Service User Number, the paying PSP that owes the Guarantee, the sponsoring PSP that vouches for the service user, and Pay.UK as rule-maker.
8. The service user's own signature appears on the indemnity, a standing, legally binding undertaking to reimburse any paying bank in the country.
9. The mandate itself is a controlled document with fixed wording and fixed field order, requiring written approval from the sponsor before use and again after any change.
10. Every Direct Debit Instruction is for an unlimited amount, which is why an electronic mandate carrying a non-zero amount is rejected under AUDDIS reason code 7.
11. Mandates are created on paper, through AUDDIS, or with no signature at all through Paperless Direct Debit, and no new non-AUDDIS Service User Numbers will be created after 1 July 2027.
12. Paperless sign-up moves verification from the bank to the service user, and brings script approval, mandatory training, recommended voice recording and possible annual review with it.

13. Advance notice, by default ten working days plus postal time and recorded against the SUN, is the obligation that most indemnity claims allege was breached.
14. Notice is required before a first collection and before any change of amount, date or frequency, must reach the payer personally, and must name the service user as the statement will show it.
15. Cancellations and amendments arrive through ADDACS and must be actioned within three working days, and a payer who cancels with the bank but not the biller leaves collections that will fail.
16. Indemnity claims arrive as DDICA advices under eight reason codes covering notice failures, cancellation failures, authority failures, service-user-requested refunds and disputed trading names.
17. The day a claim is reported counts as working day one, a challenge must be raised by the end of working day nine, and a challenge the paying bank has neither accepted nor rejected by the end of day eleven is cancelled in the service user's favour.
18. Set against the Payment Services Regulations 2017, the Guarantee is faster, broader and unconditional, so a UK payer almost never needs to invoke the statute.
19. The loss runs in a fixed order: never the payer, temporarily the paying bank, finally the service user, and the sponsoring PSP if the service user cannot pay — while the payer may still owe the underlying debt.
20. The design holds because the whole cost of being wrong sits with the one party who can choose not to be, which is why nobody collecting by Direct Debit can afford to be casual about it.

Sources: Pay.UK and Bacs published scheme material (bacs.co.uk pages on Direct Debit, AUDDIS, Paperless Direct Debit, indemnity claims, Direct Debit Instruction templates and logo, the glossary, the FAQs, the guide and rules index and the notice on non-AUDDIS service user numbers after 1 July 2027; the DDI 1 template; The Little Bacs Book), directdebit.co.uk (the Direct Debit Guarantee, how to claim, cancelling a Direct Debit, your rights and safeguards), Pay.UK's Bacs Payment System pages and Bacs annual processing statistics 2025, the Payment Services Regulations 2017 (SI 2017/752, regulations 67, 74, 76, 79 and 80) as published on legislation.gov.uk, and the published description of the Direct Debit Indemnity Claim challenge process via the Payment Services Website. Bacs message and indemnity reason codes cross-checked against independent implementation documentation; the authoritative text is The Service User's Guide and Rules to the Direct Debit Scheme, v5.9 at the time of writing, available to registered users. All timings, limits and figures accurate at time of writing.

Faster Payments

37.0 What this chapter gives you

1. You will be able to explain why a payment that reaches the plumber's account in one second does not move a penny between the two banks until Monday.
2. You will be able to say what deferred net settlement means, and why the Faster Payment System sits in exactly the same category as Bacs.
3. You will be able to work out why a payment made on Good Friday is spendable that afternoon and is not settled between banks until the following Tuesday.
4. You will be able to explain why the £1 million scheme cap is almost never the limit a customer actually meets, and name what binds instead.
5. You will be able to describe the three participation models and say which one requires a settlement account at the Bank of England.
6. You will be able to explain what prefunding and the Net Sender Cap are for, and why the July 2026 flexible model lowers the cost of entry for non-banks.
7. You will be able to distinguish the four transaction types and say why standing orders break the assumption that everything runs 24/7.
8. You will be able to state precisely what "irrevocable" means here, and which remedies still exist after a payment has gone.
9. You will be able to explain why open banking is a way of instructing a rail rather than a rail in its own right.
10. You will be able to size the netting efficiency — roughly six to nine pounds of customer payments discharged for every pound of central bank money that moves.

At 11:47 on a Sunday night in July, a man in Leeds opens his banking app, taps in a sort code, an account number and £340, and pays a plumber who fixed his boiler that afternoon. The plumber's phone buzzes before the man has locked his screen. The money is there. She can spend it.

Nothing about that sentence would have been possible in Britain in 2007. Before May 2008, a payment instructed by telephone or internet banking went into the Bacs cycle and arrived three working days later. A Sunday-night instruction would have reached the plumber on Wednesday. There was no Sunday, because Bacs does not run at the weekend, so in practice it would have been Thursday. The plumber would have gone to bed on Sunday with an invoice and nothing else, and would have chased the man on Tuesday to ask whether he had paid.

The Faster Payment System changed that, and it changed it more thoroughly than most people realise. It did not merely speed up an existing thing. It created a new expectation — that a payment between two ordinary British bank accounts is a real-time event, available at any hour, on any day, with no batch, no cut-off and no working-day arithmetic — and that expectation has since spread un-

derneath a great deal of British commercial life. Payroll advances, gig-economy earnings, marketplace payouts, insurance settlements, refunds, the entire account-to-account payments industry and every open banking payment initiation in the country all sit on top of it.

And almost none of it is true in the way it appears. The plumber's £340 did not travel from one bank to another on Sunday night. Nothing travelled. What travelled was a message, and the plumber's bank paid her out of its own money on the strength of it. The actual money between the two institutions moved during Monday, in one of three settlement cycles, netted against every other payment those institutions had exchanged. The seconds you can see and the settlement you cannot are two different systems bolted together, and the gap between them is where the risk, the liquidity cost and most of the interesting engineering live.

This chapter explains both halves: the fast one you experience, and the slow one underneath it that makes the fast one affordable.

The plain version

Imagine two villages either side of a wide river. There is one bridge, and it is a long way round.

Each village has a post office that looks after everybody's money. If Mrs Ali in the north village wants to send £180 to Mr Grant in the south village, the old way was simple and slow: the north postmaster counts £180 out of his drawer, puts it in a locked box, and waits for the van. The van crosses the bridge once a day. Mr Grant gets his money the following afternoon, or the one after that if the van has already gone.

Now somebody installs a telephone line between the two post offices.

The north postmaster picks up the phone and says: "Mrs Ali is sending £180 to Mr Grant. Pay him out of your drawer now. I'll square it with you when the van comes."

The south postmaster puts down the phone, opens his own drawer, and hands Mr Grant £180 of his own cash. Mr Grant has his money in under a minute. Nothing has crossed the bridge. There is no van in this story at all yet.

That is the whole trick, and it is worth sitting with for a moment because everything else follows from it. The speed does not come from moving money quickly. It comes from *not moving money at all* and having somebody at the far end pay out on a promise. The message is fast because messages are cheap and light. The money is slow because money is heavy, and the clever part of the design is realising you do not have to move the money in order to make the payment happen.

■ What the van still does

The van has not been abolished. It runs three times during each working day, and it carries the difference.

Suppose that during Monday morning the north post office phones the south fourteen times, for payments totalling £8,420, and the south phones the north nine times, for payments totalling £7,955. At lunchtime the two postmasters compare notes. North owes south £8,420; south owes north £7,955. Subtract one from the other and north owes south £465.

The van crosses the bridge carrying £465. Twenty-three payments worth £16,375 have been settled by moving four hundred and sixty-five pounds.

This is why the arrangement is cheap enough to give away free. If the van had to carry every payment separately it would need to cross twenty-three times with sixteen thousand pounds aboard, and somebody would have to pay for that. Instead it crosses three times a day with the leftovers.

■ What the south postmaster is worried about

Here is the part a bright twelve-year-old spots straight away, usually before an adult does.

The south postmaster handed Mr Grant £180 of his own cash on the strength of a phone call. What if, between the phone call and the van, the north post office burns down? Or the north postmaster turns out to have been lying, or has run out of money, or has simply disappeared? The south postmaster has given away real cash in exchange for a sentence spoken down a telephone line.

That worry is the single most important fact about this system, and it is invisible to everybody using it.

The way it is dealt with is not trust. It is a locked box. Before the north post office is allowed to use the telephone at all, it must take a quantity of its own cash to the county courthouse — a building neither village controls, which will not lose the money and will not go bust — and lock it in a box there. The box is not the north post office's spending money. It sits untouched. Its only job is to be there if the north post office fails, so the south postmaster can be paid from it and nobody who took money out of a drawer that morning is left short.

The size of the box is set by a rule: the north post office may never, between one van and the next, run up a debt to the other village larger than what is in the box. Once it has phoned through enough payments to reach that figure, the telephone stops working for outgoing calls until the van has been and the slate has been cleared.

So the picture is: a telephone that is always on, a van that runs three times a working day, and a locked box at the courthouse that makes the telephone safe.

■ The rules on the telephone

There are three more rules, and each of them has a reason.

The first is that there is a ceiling. The south postmaster will accept a phone call for anything up to one million pounds. Above that, the money goes by armoured van, immediately, with two guards and a signature — that is a different service, and it costs money, and it is what people buying houses use. The telephone is for ordinary amounts, and one million pounds is where “ordinary” is judged to stop.

The second is that each post office may set a *lower* ceiling of its own for its own customers, and they all do. The village rule says a million. Your own postmaster may say that you personally may phone through twenty-five thousand pounds and no more, because he has known you for two years and that is as far as his nerve goes. Both rules apply, and the lower one is the one that stops you.

The third is that the telephone is on all the time, and the van is not. The telephone works at three in the morning on Christmas Day. The van runs on working days only. That is fine, because of the locked box — but it means that over a long bank holiday weekend the north post office can be phoning payments across for four days before a single van crosses the bridge.

■ The worked example

Take the man in Leeds and the plumber.

He instructs £340 at 11:47pm on Sunday 12 July 2026. His bank checks he has the money, checks the payment against its fraud rules, and sends a message to the central system. The central system

routes it to the plumber's bank, which decides in under a second whether it can accept the payment, and answers. The whole round trip — sending bank to central system to receiving bank and back — is designed to complete inside fifteen seconds, and normally takes a small fraction of that. The plumber's bank credits her account immediately, out of its own funds. She can spend the £340 at 11:48pm on Sunday night.

The two banks square up on Monday. Three times during Monday's working hours, every bank in the system compares what it owes with what it is owed and one number per bank moves at the Bank of England. The man's £340 is a rounding error inside that number. Neither he nor the plumber will ever know which of Monday's three cycles carried it, and neither of them needs to.

Now scale it up. In 2025 the Faster Payment System carried 5.55 billion payments worth £4.84 trillion. That is roughly fifteen million payments a day, averaging about £870 each. Yet the money that actually had to move between the banks to settle all of it averaged around £2.3 billion a day in the first half of 2026. Somewhere between six and nine pounds of customer payments are being discharged for every pound of real money that moves. The rest cancels out — north owing south, south owing north, and the subtraction doing the work.

You should now be able to say the following at dinner and be right. When you send money by bank transfer in Britain, nothing is transferred. A message goes across in about a second, the receiving bank pays its customer out of its own pocket, and the banks square up with each other three times on each working day by moving only the difference. It feels instant because it is instant for you. Underneath, it is a promise with a locked box behind it.

Where the plain version stops being true

The village is a good picture of the mechanism and a misleading picture of the risk, the availability and the guarantees. Four corrections.

The first correction is that “instant” describes the customer experience, and only the customer experience. No money moves between institutions in seconds. Not some of it, not most of it — none. The Faster Payment System is a deferred net settlement system, which is precisely the same category as Bacs, and the only reason it does not feel like Bacs is that the receiving institution is contractually obliged to credit its customer immediately rather than waiting for settlement. The interbank leg happens three times each business day in the Bank of England's real-time gross settlement system, and the Bank's own description of the lag is worth quoting because it is more candid than anything the industry says in public: the gap between a customer payment and its settlement “could be several hours for Faster Payments during the week”, but over an Easter weekend “the gap between settlements is currently from between Thursday evening and Tuesday morning”. Read that again. A payment made on Good Friday is spendable on Good Friday and is not settled between banks until the following Tuesday. For four days, one institution has paid out real money and holds nothing but an obligation. That is not a defect and it is not a secret; it is the design, and the prefunded cash at the Bank of England exists precisely because of it. But anybody who tells you Faster Payments “settles in real time” has confused two layers of the system, and if they are designing a treasury function on that belief they will get their liquidity model wrong.

The second correction is that “24/7” is true of the service and false of the settlement infrastructure. The Faster Payment System is available twenty-four hours a day, every day of the year. The Bank of England's RTGS system is not. As of the time of writing, RTGS operates from 06:00 to 18:00, and the Bank has announced an early-morning extension to 01:30 from September 2027; a move towards near round-the-clock settlement, including a possible additional weekend settlement day, is at the

consultation stage, with the consultation paper published on 18 May 2026. So the always-on property that everybody markets belongs to the messaging layer. Beneath it sits a settlement layer that keeps banking hours, does not work weekends, and is only now being asked whether it might. The mismatch between the two is not a temporary lag in modernisation — it is the reason the prefunding regime exists, and any change to settlement hours changes the economics of prefunding directly.

The third correction is that the £1 million cap is almost never the limit you actually hit. One million pounds is the central scheme ceiling per transaction, and it has been since 10 February 2022, when it was raised from £250,000. Every participating institution sets its own lower limits, and those limits vary enormously by institution, by channel and by whether the account is personal or business. On Pay.UK's own published comparison, as last updated in July 2026, HSBC's online personal limit was £25,000 against a business limit of £1,000,000; Lloyds was £25,000 personal and £99,999 business; Barclays £50,000 personal; NatWest £100,000 personal; Metro Bank £21,000 personal. Starling, Monzo and several others publish the full £1,000,000 for both, subject to their own additional checks. So "Faster Payments can carry a million pounds" is a statement about the scheme, not about your account, and in most retail cases the binding constraint is two orders of magnitude below the headline. The cap is also per transaction rather than per day, and institutions layer daily aggregate limits, new-payee cooling-off periods and value-triggered verification on top of it. None of that is in the rulebook figure.

The fourth correction is that "irrevocable" and "guaranteed" are being used to mean four different things, and none of them is quite what the customer thinks. It is true that once a Faster Payment has been sent, the payer cannot cancel it — Pay.UK states this plainly, and it follows from the real-time design. But that is not the same as saying nothing can be done. Where a payment has gone to the wrong account by mistake, the industry operates a Credit Payment Recovery process. Where it has gone to the wrong account because the payer was defrauded, a mandatory reimbursement regime has applied since 7 October 2024. And "instant" itself carries an exception: the scheme's own published expectation is that funds are usually available almost immediately but can take up to two hours, and a very small proportion of UK accounts — Pay.UK puts it at under 0.1%, including some savings accounts — do not accept Faster Payments at all. A receiving institution that cannot post immediately is expected to tell the customer when the funds will arrive, not to fail the payment. So the honest formulation is: unstoppable by the sender, usually instant, occasionally two hours, sometimes recoverable afterwards, and increasingly reimbursable if it was a scam.

The technical version

■ What the system is, and who runs what

The Faster Payment System is one of the retail interbank payment systems operated by Pay.UK, the recognised operator and standards body for UK interbank retail payments. Pay.UK owns the scheme: the rules, the participation criteria, the standards and the liability arrangements. It does not build or run the machines. The central infrastructure is supplied and operated under contract by Vocalink, a Mastercard company, which also provides the central infrastructure for Bacs and for the Image Clearing System. On 17 December 2025 Pay.UK announced multi-year extensions to all three of those infrastructure contracts, providing continuity into the early 2030s.

Settlement is provided by the Bank of England. The Bank is not the operator of Faster Payments and has no role in the scheme rules; it is the settlement agent, holding the accounts in which the interbank obligations are finally discharged in central bank money.

The system launched on 27 May 2008, following a testing period in which very small-value transactions were processed. Its origins lie in the Office of Fair Trading's Payment Systems Task Force,

established in March 2004, which by May 2005 had secured an industry agreement to reduce clearing times for telephone, internet and standing order payments. Vocalink was awarded the infrastructure contract in October 2005 and has operated it ever since. Ten institutions went live at launch: Abbey (now Santander UK), Barclays, HSBC, Lloyds TSB, The Royal Bank of Scotland, Citi, Bank of Scotland, Clydesdale Bank, The Co-operative Bank and Nationwide Building Society.

The programme intended to replace the underlying infrastructure — the New Payments Architecture — has been through a substantial reset. Following the cancellation of the NPA procurement and the government's National Payments Vision, the work was renamed Interbank Infrastructure Renewal. The Payment Systems Regulator revoked Specific Direction 3 in May 2025 and Specific Direction 2 in August 2025, and closed its NPA project page on 14 October 2025. The stated long-term destination remains a single central infrastructure clearing and settling Bacs, cheques and Faster Payments, rather than three separate ones. The Vocalink contract extensions announced in December 2025 should be read against that: the existing infrastructure is now expected to carry the traffic into the early 2030s.

Scale, as reported by Pay.UK: 5.55 billion transactions worth £4.84 trillion during 2025. For comparison, the system handled 967.6 million transactions worth £771.4 billion in 2013.

■ The four transaction types

Pay.UK defines four categories of Faster Payment. They differ in initiation channel, availability and timing, not in the underlying clearing mechanism.

Type	Availability	Central limit	Timing characteristics
Single Immediate Payment	24 hours, 7 days	£1,000,000	Usually near-instant; may take up to two hours; central infrastructure returns accept or reject to the sending PSP within 15 seconds
Forward-dated payment	Can be set up 24 hours, 7 days	£1,000,000	One-off, executed on a pre-arranged future date
Standing order	Monday to Friday, excluding bank holidays	£1,000,000	Fixed amount, recurring date; rolls to next working day if the date falls on a weekend or bank holiday; at least 90% sent by 6am
Direct Corporate Access	24 hours	£1,000,000	Bulk file submission by business customers, same-day transmission; not offered by all participants

Several points in that table repay attention.

The fifteen-second figure is a property of the central infrastructure, not of the end-to-end customer experience. It is the window within which the sending PSP receives confirmation that the receiving PSP has accepted or rejected the payment, and where the answer is a rejection, the receiving PSP supplies a reason. This is the mechanism that makes the customer-facing “sent” and “received” states

meaningful. The two-hour outer bound covers the cases where the receiving institution cannot post immediately — an operational problem at its end, or an account type that does not permit immediate posting — and in those cases the scheme expects the customer to be given an indication of when the funds will arrive rather than a bare failure.

Standing orders are the outlier and are frequently misunderstood by engineers who assume the 24/7 property extends to everything. It does not. Standing orders execute on working days only, roll forward when the nominal date is not a working day, and cluster heavily in the early morning: Pay.UK's published expectation is that at least 90% are sent by 6am. If you are building a system whose behaviour depends on when a standing order will land, that early-morning clustering and the working-day roll are the two facts that matter.

Direct Corporate Access is the least discussed of the four and the most relevant to anyone with a bulk payment file. It allows a business customer to submit files of payment instructions directly into the Faster Payment System, using a secure-IP solution comparable to Bacstel-IP and, notably, the same standard file format as Bacs transactions. That last detail is the practical bridge for organisations that already produce Standard 18 files for Bacs: the file format does not have to change to move the same payment run onto a same-day rail. Not every participant offers Direct Corporate Access, and those that do not typically offer an alternative bulk mechanism of their own. A number of software providers are accredited to supply Direct Corporate Access services.

■ The value cap and its history

The central scheme limit is £1,000,000 per transaction. It has stood at that level since 10 February 2022, when Pay.UK raised it from £250,000; at the time of the increase there were 38 directly connected payment service providers and around 400 PSPs connected indirectly through sponsors. Earlier in the system's life the ceiling was substantially lower, having been raised to £100,000 for all payment types on 6 September 2010.

Institution-set limits are what bind in practice. Pay.UK maintains a published comparison of the limits set by individual banks and building societies; the figures below are drawn from that page as last updated on 24 July 2026, and are accurate at time of writing. They will change.

Institution	Online personal	Online business
HSBC	£25,000	£1,000,000
Barclays	£50,000	£1,000,000 (high value)
Lloyds Bank	£25,000	£99,999
NatWest	£100,000	£250,000 to £1,000,000
Metro Bank	£21,000	£250,000
Starling Bank	£1,000,000	£1,000,000
Monzo	£1,000,000	£1,000,000

Monzo's entry carries a rider worth generalising: payments above £10,000 are subject to further checks. Value-triggered friction of that kind is now routine across the industry and is not visible in any published limit table. A £1,000,000 headline limit and a £10,000 verification threshold coexist comfortably, and the second is what a customer will actually encounter.

■ Access: three participation models

Pay.UK defines three ways to participate, and the distinction between them is the most commercially consequential thing in this chapter for anybody building a payments business.

A **Directly Connected Settling Participant** holds its own technical connection to the central infrastructure and its own settlement account at the Bank of England, and performs its own settlement. It sets its own transaction limits. This is the full-fat model, appropriate to institutions with high volumes and eligibility for a Bank of England settlement account.

A **Directly Connected Non-Settling Participant** holds its own technical connection to the central infrastructure but does not settle at the Bank of England. A sponsoring participant performs the Bank of England settlement on its behalf. This model exists precisely for organisations that want the operational control and latency of a direct connection but are not eligible for, or do not want, a settlement account.

An **Indirect PSP** has neither its own connection nor its own settlement. It sends and receives Faster Payments through a sponsoring participant, which also performs the Bank of England settlement. The sponsor allocates the indirect participant a sort code from its own range and registers it with the payment systems and with the industry reference data. The indirect participant still processes its own payment messages, handles its own exception flows, and manages its own settlement and safeguarding accounts with the sponsor — those latter accounts being commercial bank accounts, and not to be confused with a Bank of England settlement account.

The eligibility criteria for the settlement account are set by the Bank of England, not by Pay.UK. To hold a settlement account, a firm must either hold a reserves account or be a financial market infrastructure, or an e-money or payment institution authorised in the UK by the Financial Conduct Authority, and must meet the settlement participation criteria of the relevant payment system.

That last clause is the one that changed the shape of the industry. In July 2017 the Bank of England extended direct access to RTGS settlement accounts to non-bank payment service providers. The consequences arrived quickly: Ebury became the first directly connected non-settling participant in February 2018, and TransferWise — now Wise — became the first non-bank settling participant in April 2018, followed by Equals Money in February 2019. Pay.UK's own access programme, launched in 2014, has seen the number of direct Faster Payment participants more than triple.

The current shape of that is visible in the published participant lists. As last updated in April 2026, the Faster Payment System participant list named 47 organisations, against 32 for Bacs. Among the 47 are Adyen, Banking Circle, ClearBank, Corpay, Modulr, PPS, Revolut, Soldo, Square, Stripe UK, SumUp, The Bank of London and Wise — a set that would have been unimaginable as direct participants in a UK clearing system fifteen years ago. Faster Payments is now, by participant count, materially more open than Bacs.

Pay.UK also operates an accreditation regime for technical aggregators, solution providers and consultants, with a published fee structure for Faster Payment System technical accreditation. For most new entrants the practical route to a direct connection runs through one of those accredited providers rather than through building the gateway in-house.

■ Settlement: deferred net, prefunded

Settlement of Faster Payments takes place three times each business day in the Bank of England's RTGS system. The Bank's classification is explicit: this is one of four settlement models it supports, and Faster Payments sits in the prefunded deferred net settlement category, alongside Bacs and the Image Clearing System.

The mechanics are multilateral netting. Between settlement cycles, participants exchange customer payments through the central infrastructure, entirely outside RTGS — the Bank notes that RTGS does not need to be open for this exchange to happen, which is what makes the 24/7 service possible at all. Obligations accrue. At each settlement cycle, the operator calculates each settling participant's multilateral net position against all others, and a single net amount per participant is settled across accounts in RTGS.

Because obligations accrue between cycles, settlement risk arises between cycles. The mitigation is prefunding, introduced for both Bacs and Faster Payments in September 2015. Each settling participant holds cash in a dedicated prefunding account in RTGS. That balance caps the participant's maximum net obligation to the system. If a settlement participant defaults, the segregated cash is used to complete settlement. The effect, in the Bank's own words, is that prefunding “eliminates settlement risk in systems that use it by capping the maximum net obligations of settlement participants” and “eliminates credit risk between settlement participants”. Before 2015, the exposure was managed by a mutualised loss-sharing arrangement among participants; prefunding removed that mutualisation entirely.

The cap itself is the Net Sender Cap. It limits a participant's indebtedness to other participants within each settlement cycle, and it is backed by the prefunded balance. Historically Pay.UK set Net Sender Caps prescriptively, using a Minimum Net Sender Cap together with a Peak Contingency Value that participants had to hold at all times. The observed problem was that participants held more liquidity than they needed, with the high-usage cases occurring only infrequently — dead capital sitting in a segregated account at the Bank of England against a peak that rarely came.

On 1 July 2026, Pay.UK went live with a flexible model. Participants may now determine their own Net Sender Cap values according to their own business requirements, submitting plans quarterly for Pay.UK review. Adoption is optional; participants may remain on the existing prescriptive framework. Real-time alerts and the existing settlement protections are retained, and Pay.UK states that no technical change is required to adopt the new model. The intended effects are lower liquidity cost for existing participants and a lower barrier to entry for new ones, particularly non-banks, for whom the prefunding requirement is a direct and highly visible cost of participation.

The numbers give a sense of the netting efficiency. The Bank of England's published average daily RTGS settlement values for Faster Payments were £2,008 million in Q3 2025, £2,104 million in Q4 2025, £2,287 million in Q1 2026 and £2,298 million in Q2 2026. Set that against Pay.UK's reported £4.84 trillion of Faster Payments in 2025 — roughly £13.3 billion for every day of the calendar year, and closer to £19 billion for every business day, since payments accrue over weekends and settle on working days. The ratio of customer payments to central bank money moved therefore sits somewhere between six and nine to one depending on how the days are counted. Either way, the great majority of the value cancels before anything settles. It is also worth noting that Bacs, with far fewer participants and a slower cycle, settles a larger daily net value than Faster Payments: £6,205 million average daily in Q2 2026. Faster Payments dominates by volume and by public visibility, not by settled value.

RTGS operating hours are the constraint on all of this. At the time of writing they run 06:00 to 18:00. The Bank has announced that CHAPS will open at 01:30 from September 2027, giving a 16.5-hour

settlement window, and on 18 May 2026 published a consultation paper on further extension towards near 24x7 settlement, including the possibility of an additional weekend settlement day, most likely Sunday. The Bank identifies liquidity and risk management as one of the three principal drivers, noting explicitly that longer operating hours “could enable more frequent net settlement cycles, reduce settlement risk, and allow liquidity to be used more efficiently, lowering prefunding and liquidity cost”.

The Bank’s own worked illustration of the current exposure is the sharpest statement of the seam in this chapter. Between settlement cycles on a normal weekday, the lag “could be several hours”. Over an Easter weekend, the gap between settlements runs “from between Thursday evening and Tuesday morning”. In 2026 that means the last settlement before Good Friday on 3 April, and the next on Tuesday 7 April — four full days during which Faster Payments continue to flow at full volume, receiving institutions continue to credit customers, and the obligations pile up behind prefunded balances. The Bank makes a related observation in the same paper that deserves to be read by anyone modelling stress: outflows from banks can already take place at the weekend through Faster Payments, because it operates 24x7, at a time when most other sterling flows are closed.

■ Messages and standards

The Faster Payment System’s native message set is derived from ISO 8583 — the same lineage as the card messages in Volume III, which is why an engineer who has worked on card switching finds the FPS message structure oddly familiar. Pay.UK maintains an FPS ISO 20022 Standards Library, described by Pay.UK as “a reference source providing information on the recommended conversion of ISO 8583 messages to the ISO 20022 messaging standard”. The library sits behind registration on Pay.UK’s standards portal rather than on the open website.

Beyond the payment messages themselves, participants exchange status reporting — in Faster Payments terminology, the fate of a payment — and a set of exception messages covering recalls, returns and related error flows. A direct or indirect participant must be able to generate, receive and process all of them; supporting only the happy path is not participation.

■ Confirmation of Payee

Confirmation of Payee is an overlay service, not part of the Faster Payment System itself, but no honest account of a UK account-to-account payment in 2026 omits it. Launched by Pay.UK in 2020, it checks the payee’s account name against the account details entered before the payment is initiated, and returns a result to the payer. Its purpose is to reduce misdirected payments and to attack the specific fraud pattern in which a customer is persuaded to send money to an account in a different name from the one they believe they are paying.

Adoption is now broad: Pay.UK reports over 300 organisations participating and more than two million checks completed every day. Participation follows the same shape as the payment systems themselves — direct participants that are account-holding PSPs, and indirect participants reaching the service through accredited PSP aggregators. The Payment Systems Regulator mandated an expansion of coverage in 2024.

■ Liability: who pays when it goes wrong

Two regimes matter.

For payments sent to the wrong account in error, the industry operates the Credit Payment Recovery process. It is a best-efforts recovery mechanism, not a right of reversal; the receiving customer’s consent, or a legal process, is generally required to get the money back.

For authorised push payment fraud — where the customer was deceived into authorising the payment themselves — a mandatory reimbursement requirement has applied to Faster Payments since 7 October 2024. The maximum level of mandatory reimbursement is £85,000 per claim, confirmed by the Payment Systems Regulator in policy statement PS24/7; the Bank of England, as operator of CHAPS, set the same £85,000 maximum for CHAPS to keep the two consistent. The PSR's assessment was that £85,000 covers 99.8% of Faster Payments APP scams by volume and 90% by value. Firms may apply an optional excess of up to £100, but may not apply it to vulnerable consumers. Claims must be reported within 13 months of the payment. Reimbursement is due within five business days of the claim, with the clock stoppable where the firm needs more information, and a decision required within 35 business days. The cost is split 50:50 between the sending and the receiving payment service provider — a design choice that deliberately puts the receiving institution, which onboarded the mule account, on the hook for half.

Underneath both regimes sits the statutory backstop in the UK's payment services regulations, derived from PSD2: a payment instructed by mobile, internet or telephone banking must arrive by the end of the following business day at the latest. Faster Payments beats that by roughly a day and a half. The regulation is the floor, not the standard.

■ What Faster Payments is not

It is not the rail for large value. Above £1 million, and for anything requiring same-day irrevocable settlement in central bank money on a per-payment basis, the answer is CHAPS, which settles gross across RTGS with an early-afternoon cut-off for the same-day guarantee and carries the large majority of sterling value on a tiny fraction of the volume.

It is not real-time settlement, and describing it as such in a prospectus, an architecture document or a liquidity model is a material error rather than a simplification.

It is not the same thing as open banking. Open banking is an API layer that allows a third party to initiate a payment from a customer's account with the customer's consent; the payment it initiates is, overwhelmingly, a Faster Payment. Open banking has no clearing, no settlement, no participant list and no prefunding. It is a way of instructing a rail, not a rail.

And it is not, any longer, a banks-only club. Nearly half the named direct participants are non-banks, they hold settlement accounts at the Bank of England on the same terms as the clearers, and the July 2026 change to the Net Sender Cap regime was explicitly designed to make that easier still.

■ What this means if you are building on it

Three practical consequences follow from the technical picture, and they are the ones most often got wrong.

The first is that prefunding is a real and continuous cost of direct settling participation, sized to your peak net sending position between cycles rather than to your average. It is cash, it is segregated, and until the flexible Net Sender Cap model it was sized conservatively by rule rather than by your own traffic. If your business is heavily one-directional — a payout business that sends far more than it receives — your net sender position, and therefore your prefunded balance, is close to your gross volume, and the economics of direct settling participation look very different from those of an institution with balanced two-way flow.

The second is that there is no cut-off to design around, and that is an unfamiliar discipline for teams coming from Bacs. There is no submission window, no processing day, no entry day. There is a queue, a fifteen-second response expectation, a two-hour outer bound, and a settlement cycle you do not

control and cannot see. Reconciliation must therefore be built against your own message-level fate reporting rather than against a settlement file, because the settlement file tells you about a net figure that bears no reconstructable relationship to individual payments.

The third is that the weekend is the risk. Volume does not stop on Friday evening; settlement does. Whatever exposure your model tolerates on a Tuesday afternoon, multiply it by the length of the longest bank holiday weekend in the calendar and check that the prefunded balance and the Net Sender Cap still work. That is the calculation the Bank of England is effectively inviting the industry to revisit in its May 2026 consultation, and it is the single most instructive number in the whole system: the measure of how far a payment that arrives in one second can outrun the money behind it.

37.98 Common wrong ideas

Wrong: Faster Payments settles in real time. **Right:** It is a deferred net settlement system; the interbank leg settles three times each business day at the Bank of England, and only the credit to the receiving customer is immediate.

Wrong: The whole system runs twenty-four hours a day. **Right:** The messaging service does; RTGS runs 06:00 to 18:00 on working days, with an extension to 01:30 announced for September 2027 and near round-the-clock settlement still at consultation.

Wrong: You can send £1 million by Faster Payment from any account. **Right:** £1 million is the central scheme ceiling per transaction; institution-set limits bind first, with HSBC's online personal limit at £25,000 and Metro Bank's at £21,000 as published in July 2026.

Wrong: A Faster Payment is always instant. **Right:** Funds are usually available almost immediately, but the scheme's own published expectation allows up to two hours, and under 0.1% of UK accounts do not accept Faster Payments at all.

Wrong: Once sent, nothing whatever can be done. **Right:** The payer cannot cancel, but Credit Payment Recovery exists for payments sent to the wrong account in error, and a mandatory reimbursement regime capped at £85,000 per claim has applied to APP fraud since 7 October 2024.

Wrong: Faster Payments settles more money than Bacs. **Right:** Bacs settles a larger daily net value — £6,205 million average daily in Q2 2026 against about £2,298 million for Faster Payments — because Faster Payments dominates by volume and visibility, not by settled value.

Wrong: Open banking is a rail competing with Faster Payments. **Right:** Open banking has no clearing, no settlement, no participant list and no prefunding; it is an API layer, and the payment it initiates is overwhelmingly a Faster Payment.

Wrong: Standing orders inherit the 24/7 property of the scheme. **Right:** They execute on working days only, roll forward when the nominal date is a weekend or bank holiday, and cluster early, with at least 90% sent by 6am.

Wrong: Faster Payments is a banks-only club. **Right:** The participant list named 47 organisations as last updated in April 2026, including Adyen, ClearBank, Modulr, Revolut, Square, Stripe UK, SumUp and Wise, making it materially more open than Bacs.

Wrong: You can reconcile individual payments against the settlement file. **Right:** The settlement file reports one net figure per participant per cycle, with no reconstructable relationship to individual payments; reconciliation must be built against message-level fate reporting.

37.99 Chapter summary in 20 lines

1. Before May 2008 a British payment instructed by telephone or internet banking went into the Bacs cycle and arrived three working days later.
2. The Faster Payment System launched on 27 May 2008 and created the expectation that an account-to-account payment is a real-time event at any hour on any day.
3. Nothing actually travels between the banks when you send one: a message goes across, and the receiving bank credits its customer out of its own funds.
4. The speed comes from not moving money at all and having somebody at the far end pay out on a promise.
5. The banks square up three times each business day, netting everything they have exchanged and moving one figure per participant at the Bank of England.
6. That netting is what makes the service cheap enough to give away, with somewhere between six and nine pounds of customer payments discharged for every pound of central bank money moved.
7. The exposure created between the message and the settlement is capped by prefunded cash held in a segregated account in RTGS.
8. Prefunding, introduced in September 2015, replaced a mutualised loss-sharing arrangement and caps each participant's maximum net obligation through the Net Sender Cap.
9. On 1 July 2026 Pay.UK went live with a flexible model letting participants set their own Net Sender Cap values, lowering liquidity cost and the barrier to entry.
10. "Instant" describes the customer experience only, because Faster Payments is a deferred net settlement system in precisely the same category as Bacs.
11. The gap between a customer payment and its settlement can be several hours on a weekday and, over an Easter weekend, runs from Thursday evening to Tuesday morning.
12. "24/7" belongs to the messaging layer, while the settlement layer beneath it keeps banking hours and does not work weekends.
13. The £1 million per-transaction ceiling, in force since 10 February 2022, is rarely the binding limit, because institution-set limits are often two orders of magnitude lower.
14. Four transaction types share one clearing mechanism: Single Immediate Payments, forward-dated payments, standing orders and Direct Corporate Access.
15. Standing orders are the exception to everything always-on, running on working days only, rolling forward, and clustering before 6am.
16. Participation comes in three shapes — directly connected settling, directly connected non-settling, and indirect through a sponsor — and only the first holds its own Bank of England settlement account.
17. Since the Bank opened settlement accounts to non-bank payment service providers in July 2017, the direct participant list has grown to 47 organisations, nearly half of them non-banks.
18. The sender cannot cancel a payment, but Credit Payment Recovery covers misdirection and mandatory reimbursement, split 50:50 between sending and receiving firms, covers APP fraud.
19. Faster Payments is not the rail for large value, is not real-time settlement, and is not the same thing as open banking.
20. If you build on it, prefunding is a continuous cost sized to your peak net sending position, there is no cut-off to design around, and the weekend is where a payment that arrives in one second most outruns the money behind it.

Sources: Pay.UK published scheme material (wearepay.uk Faster Payment System pages covering how Faster Payments work, transaction types, transaction limits, participation, participant lists and the FPS ISO 20022 Standards Library; Confirmation of Payee pages; the 1 July 2026 flexible Net Sender Cap announcement and

explainer; the 17 December 2025 Vocalink contract extension announcement; and the February 2022 £1 million limit press release), the Bank of England (payment and settlement pages, payment and settlement statistics, and the consultation paper “Extending RTGS and CHAPS settlement hours — next steps towards near 24x7 settlement”, published 18 May 2026, together with the 2016 RTGS PFMI self-assessment for the September 2015 introduction of prefunding), and the Payment Systems Regulator (PS24/7 on the maximum level of reimbursement, APP fraud reimbursement protections, and the Interbank Infrastructure Renewal project page). Launch history and early participant milestones cross-checked against public record. All limits, timings, figures and participant counts accurate at time of writing.

CHAPS and RTGS

38.0 What this chapter gives you

1. You will be able to explain why a conveyancer waits for CHAPS rather than sending a Faster Payment, even though Faster Payments is quicker and free.
2. You will be able to state the trade at the heart of gross settlement: it removes the risk that a counterparty fails while owing you, and charges for it in cash kept parked and idle.
3. You will be able to describe how a direct participant manufactures intraday liquidity by pledging collateral, and why the real constraint is encumbrance rather than owned cash.
4. You will be able to explain what the Liquidity Saving Mechanism does, and why marking a payment urgent or non-urgent is a real decision with consequences in minutes and in liquidity.
5. You will be able to name the three points of finality in the CHAPS rulebook and say which one makes a payment unrecallable.
6. You will be able to explain why designation under the Settlement Finality Regulations 1999 is what lets the seller's solicitor pick up the phone and release the keys.
7. You will be able to work backwards from the 2pm conveyancing convention through your own bank's cut-off to the scheme's 17:40 deadline, and say why the earliest one governs.
8. You will be able to explain tiering, and why "settled in central bank money" describes one leg of the journey only.
9. You will be able to recount what happened on 20 October 2014 and say why the contingency service was available and was not invoked.
10. You will be able to decide, for a given payment, whether CHAPS is worth £25 to £30 or whether you are buying a certainty you do not need.

At about eleven on a Tuesday morning, in an office above a dry cleaner's in a market town, a licensed conveyancer refreshes a screen. She has been refreshing it for twenty minutes. On the other end of a telephone line is a client sitting in a hired van outside a house she does not yet own, with a cat in a carrier on the passenger seat and a removal crew being paid by the hour.

Nothing about this scene is unusual, and nothing about it is optional. Until a specific number appears on the conveyancer's client account screen, the house belongs to somebody else, the keys stay in the estate agent's safe, and the van stays parked. When the number appears, everything happens at once: a phone call, a release of keys, a chain of four other households doing the same thing in sequence within the next ninety minutes.

That number appears because of CHAPS, and CHAPS works because of a machine at the Bank of England called RTGS. Together they are the most important piece of financial infrastructure in the United Kingdom that almost nobody has heard of. In 2024 CHAPS carried 52.7 million payments

worth £87.5 trillion, which is roughly 0.5% of all UK payment volumes and 91% of all sterling payment value. The average CHAPS payment was £1.7 million. Most people use it once or twice in a lifetime, pay between £25 and £30 for the privilege, and assume it is simply an expensive bank transfer that goes a bit faster.

It is not faster. Faster Payments is faster, and free, and works at midnight on a Sunday. What CHAPS gives you is something else entirely, and this chapter is about what that something is: finality, in central bank money, that cannot be undone, and what the country pays to have it.

All timings, limits, fees and figures in this chapter are accurate at the time of writing, August 2026. Payment infrastructure is a moving target, and the Bank of England has published plans to change several of these numbers before 2030.

The plain version

Imagine a school where every class has a tub of counters in the head teacher's office. Class 4 has a tub, Class 5 has a tub, Class 6 has a tub. Nobody carries counters around the school. If Class 5 owes Class 6 forty counters, somebody walks to the office, the head teacher takes forty counters out of Class 5's tub and puts them in Class 6's tub, and that is that. The counters never leave the office. What moves is which tub they are sitting in.

Now, the classes owe each other things all day long. Hundreds of times a day. There are two ways to run this.

The first way is the whiteboard. Every time one class owes another, you write it on a whiteboard in the corridor. Class 5 owes Class 6 forty. Class 6 owes Class 4 twenty. Class 4 owes Class 5 fifteen. At four o'clock the head teacher looks at the whiteboard, works out who is up and who is down overall, and moves counters between the tubs once, at the end of the day. This is enormously efficient. Hundreds of debts get settled by moving a handful of counters. Nobody has to keep many counters in their tub, because most of what they owe cancels against most of what they are owed.

There is one problem with the whiteboard, and it is a serious one. At half past three, Class 5 might collapse. The teacher walks out, the class dissolves, there is nobody to pay. All those entries on the whiteboard saying "Class 5 owes" are now just chalk. Class 6 spent the afternoon handing out biscuits on the strength of a promise that Class 5 would pay at four, and at four Class 5 does not exist. Worse, Class 6 was counting on that money to pay Class 4, so now Class 4 is short too, and the whole whiteboard has to be rubbed out and recalculated as though Class 5 had never traded at all. One collapse at 3.30 rewrites everybody's day.

The second way is to walk to the office every single time. Class 5 owes Class 6 forty counters? Somebody walks to the office, forty counters move from one tub to the other, immediately, and the matter is closed forever. There is no whiteboard. There is nothing to rub out. If Class 5 collapses at 3.31, every payment it made before 3.30 is simply done, and nobody else is affected at all.

That second way is real-time gross settlement. "Real-time" means it happens when you ask, not at four o'clock. "Gross" means each payment moves on its own, at its full size, rather than being cancelled off against other payments first. RTGS is the head teacher's office. CHAPS is the queue of people walking to it.

The second way is safer. It is also more expensive, and the expense is a specific and slightly surprising one: you need a lot more counters.

Here is why. Suppose Class 6 is going to send 900 counters' worth of payments today and receive 850 back. Under the whiteboard system it only ever needs to find the 50 counter difference at four o'clock. Under the walk-to-the-office system, it depends entirely on the order. If all 900 of Class 6's payments have to go out in the morning and the 850 coming in do not arrive until the afternoon, then Class 6 needs 900 counters sitting in its tub at nine in the morning, doing nothing, just waiting. Those 900 counters could have been doing something useful. Having them sit in a tub all morning is the price of never having a whiteboard.

That is the whole trade, and it is worth saying in one sentence, because it is the sentence this chapter exists to explain. Gross settlement removes the risk that somebody fails while owing you, and it charges you for that in cash you have to keep parked and idle.

■ The house

Now the conveyancer above the dry cleaner's.

Her client is buying a house for £340,000. At exchange of contracts, three weeks ago, the client paid a 10% deposit of £34,000. On completion day the remaining £306,000 has to reach the seller's solicitor. Of that, £246,000 has come from a mortgage lender, and £60,000 is the client's own savings plus the deposit already held.

At 10.52, the conveyancer instructs her bank to send £306,000 by CHAPS to the seller's solicitor's client account. Her firm is charged about £25 by its bank for that single instruction. Somewhere inside the machinery, the Bank of England charges her firm's bank 48.7 pence for it, which is the 2026 CHAPS tariff per payment sent. The rest is the bank's margin, its own liquidity cost, and its compliance department.

Within seconds, a message reaches the Bank of England. The buyer's bank has an account there. The seller's solicitor's bank has an account there. £306,000 comes off one and goes onto the other. Not a promise, not an entry on a whiteboard: the actual balance changes, at the Bank of England, immediately and permanently.

And here is the part that matters to the woman in the van with the cat. From the instant that balance changes, the money cannot come back. There is no cooling-off period. The buyer cannot ring up and cancel. The bank cannot reverse it because it made a mistake. A court can order someone to pay it back afterwards, but that is a new payment going the other way, not an undoing of the first one. The seller's solicitor, seeing the credit on her screen, can lift the phone to the estate agent and say "release the keys", and she can do that with total confidence, because the money on her screen is not a claim on the buyer's bank. It is settled.

That confidence is the product. Everything else, including the fee, is the cost of manufacturing it.

Now multiply it. That house is in a chain. The seller of the £340,000 house is buying a £455,000 house two streets away, and cannot send that money until the £306,000 arrives. The seller of the £455,000 house is buying a flat for £238,000. Five households, five vans, five sets of keys, all keyed off a sequence of payments that must each be final before the next one can start. If any link in that chain were reversible, the whole chain would have to wait until everyone was certain, which in practice means days. The reason a chain can complete in a single morning is that each payment in it is dead certain the moment it lands.

	The whiteboard way	The walk-to-the-office way
What moves	One net figure at the end of the day	Every payment, one at a time, in full
Cash you must keep idle	Very little	A great deal
If a member collapses mid-afternoon	Everyone's day is rewritten	Nothing already paid is affected
Cost per payment	Fractions of a penny	Tens of pence to the bank, tens of pounds to you
UK examples	Bacs, Faster Payments	CHAPS

Where the plain version stops being true

■ The counters are not scarce in the way the analogy implies

The school story says a bank needs a big pile of money sitting idle to make payments early in the day. That is the right shape but the wrong mechanism, and the difference is the single most important thing a practitioner knows that a layperson does not.

A CHAPS direct participant does not fund its day only out of its own accumulated cash. It can, within limits, manufacture settlement money during the day by pledging assets. Eligible participants in the Bank's Sterling Monetary Framework can generate intraday liquidity against collateral: they pledge gilts and other very high quality assets into a pool at the Bank, the Bank values that collateral, applies a haircut, and credits the participant with settlement funds it can spend across the day. At the end of the day that liquidity is extinguished and the collateral is released. Provided the collateral is there and the day balances by the close, the participant has borrowed several billion pounds from the central bank and given it back before dinner.

So the real constraint is not "how much cash does the bank own" but "how much eligible collateral can it encumber, and what is it giving up by encumbering it". That is a subtler and more interesting cost, and it changed the industry's behaviour after the post-crisis liquidity rules made intraday positions expensive to carry. The tub of counters is really a tub plus an overdraft facility secured on the school's silverware.

■ It is not one payment at a time, economically speaking

The plain version says each payment walks to the office on its own. Legally that is exactly right, and it stays right. Economically it stopped being the whole story on 15 April 2013, when the Bank introduced a Liquidity Saving Mechanism into CHAPS.

Participants can mark a payment as urgent or non-urgent. Urgent payments settle one at a time, immediately, exactly as described. Non-urgent payments go into a central scheduler and wait a little. Every couple of minutes the system pauses the urgent stream, runs a matching algorithm for around twenty seconds, and looks for groups of payments across different banks that broadly offset each other. Those it finds, it settles all at the same instant. Each of them is still a separate gross settlement in law, with its own finality; but because they land simultaneously, the outflows are funded by the inflows and nobody has to pre-fund the full amount.

The measured effect was large. Before the mechanism, the combined intraday liquidity requirement across CHAPS averaged around £21.2 billion. In the months after, it averaged around £16.9 billion, and the Bank's own analysis attributed roughly £4 billion of that, about a 20% saving, to the mechanism itself. So the honest summary is that CHAPS is a gross system that has quietly borrowed one of netting's best tricks, while keeping gross settlement's legal guarantees. The plain version's clean

opposition between “the whiteboard” and “the walk” is a teaching device, not a description of the modern machine.

■ **Almost nobody is actually in the room**

The school has three classes and three tubs. CHAPS has, at the time of writing, over 35 direct participants, and several thousand financial institutions that reach it through one of those participants. More than 70 organisations hold accounts allowing them to settle directly in one or more payment systems at the Bank, but the CHAPS membership itself is a small club.

This is called tiering, and it means that the sentence “CHAPS is settled in central bank money and is irrevocable” describes one leg of the journey only. If your bank is not a direct participant, your payment reaches CHAPS through a correspondent, and the leg between your bank and its correspondent is an ordinary commercial arrangement with ordinary commercial credit risk. The Bank of England has worried publicly about this for years: a 2013 Bank Quarterly Bulletin article on tiering counted around 4,500 indirect participants and identified the credit, liquidity and operational risks that concentration creates, and the Bank subsequently pressed the largest indirect participants to join directly. Those particular figures are more than a decade old and the population has changed, but the structural point has not. Central bank money is the settlement asset between direct participants. Between everybody else it is commercial bank money wearing the same coat.

And your own leg, as a customer, is not central bank money at all. What lands in the seller’s solicitor’s client account is a claim on her bank. What settled at the Bank of England is a transfer between two banks. They happen close together in time, they are causally linked, and they are not the same event.

■ **“Same day” and “irrevocable” both need footnotes**

The plain version implies you can send a CHAPS payment whenever you like during the working day and it will land. Two things complicate that.

First, cut-offs come in layers. The CHAPS system opens at 6am each working day and closes at 6pm for bank-to-bank payments, with customer payments having to be submitted by 5.40pm. That is the scheme deadline. It is not your deadline. Every bank sets its own, earlier, customer cut-off, and those are frequently in the early afternoon, sometimes considerably earlier for instructions given in a branch or by a business channel that needs a callback. On top of that sits a market convention: under the Law Society’s Standard Conditions of Sale, if the completion money is received after 2pm, completion is treated as taking place the next working day for the purposes of the compensation and apportionment provisions. So the practical deadline in a house purchase is not 5.40pm and not the bank’s cut-off but 2pm, and every conveyancer in the country works backwards from it. The chapter brief for a house sale is early afternoon; the scheme’s own limit is nearly four hours later.

Second, “irrevocable” is a statement about the settlement, not a statement about your money being unreachable. A settled CHAPS payment cannot be pulled back; the receiving bank can be asked to return funds and may agree, but it is not obliged, and there is nothing resembling a card chargeback. That is the point of the rail. But since 7 October 2024 the Bank of England has applied reimbursement rules to in-scope retail CHAPS payments broadly aligned with those for Faster Payments, with a maximum reimbursement of £85,000 per claim for authorised push payment fraud. So a victim of a conveyancing fraud is not necessarily without remedy. The payment is still final. The remedy is a different payment, running the other way, from a different party, under a different rulebook.

CHAPS is also, second by second, not a system without limits in the way people assume. Faster Payments caps a single transaction at £1 million. CHAPS scheme rules set no maximum value at all,

which is why a £4 billion money-market payment and a £306,000 house purchase travel on the same rail. What limits your payment is your own bank's policy, not the scheme.

The technical version

■ RTGS is a ledger, not a payment system

The first correction a practitioner makes is terminological. RTGS is not a payment scheme. It is the Bank of England's accounting infrastructure: the ledger on which sterling settlement accounts are held and on which balances move. The Bank describes it as an accounting system that underpins the settlement of sterling payments. CHAPS is a payment system that settles on it. So do others.

That distinction matters because the ledger is shared. On an average working day, over £800 billion settles across RTGS. That total is not CHAPS alone. It includes CREST, the securities settlement system, whose deliveries against payment settle in central bank money on the same ledger; it includes the net settlement obligations arising from Bacs, Faster Payments, LINK, cheque clearing and card schemes, each of which computes its own multilateral net positions and then discharges them by a small number of movements across RTGS accounts. The retail rails you use every day are, from the Bank's point of view, a handful of large entries per day. CHAPS is the system that pushes each payment individually onto the ledger.

The core ledger itself was replaced under the Bank's RTGS Renewal Programme. The renewed core settlement engine, RT2, completed its final transition stage and went live on 28 April 2025, and the Bank's revised tariff framework took effect on the same date.

Settlement account balances at the Bank are central bank money: a liability of the central bank rather than of a commercial bank, with no credit risk in any conventional sense. The Bank characterises it as the ultimate secure and liquid sterling asset. When a CHAPS payment settles, the sending participant's balance falls and the receiving participant's balance rises, in the Bank's books, in real time.

■ What "real-time gross settlement" removes, precisely

The risk RTGS removes is settlement risk in its principal form: the risk that you have given value and your counterparty fails before giving value back. In a deferred net settlement system, an obligation exists from the moment a payment is exchanged until the moment the net position settles. Between those two moments there is an exposure. In the retail systems that exposure is managed by prefunding, caps and loss-sharing. In a high-value system moving hundreds of billions a day, an intraday exposure of that size between a small number of institutions would be a financial stability problem in its own right, and there is no plausible amount of collateral that would make it comfortable.

Real-time gross settlement removes it by construction. Because settlement occurs in real time, in central bank money, participant by participant, no intraday exposures accumulate between account holders at all. There is nothing to unwind, so there is no unwind risk: the scenario in which one participant's failure forces recalculation of everybody else's net positions simply cannot arise.

What it does not remove is liquidity risk, and this is where the design cost sits. Because payments settle one by one, they settle only if funds are available at that instant. A participant that has sent more than it has received needs a balance to cover the difference. Aggregate across the system and you get a queueing problem with a free-rider incentive: everybody would prefer to receive early and send late, and if everybody acts on that preference the system gridlocks. The architecture responds to that incentive in three ways: liquidity supply, liquidity saving, and behavioural rules.

■ Liquidity supply

Direct participants fund their CHAPS activity from their settlement account balances. Sterling Monetary Framework participants, and certain other account holders case by case, can generate intraday liquidity secured against the very highest quality collateral with prudent haircuts, using a collateral pooling arrangement at the Bank. In the Bank's published service description, intraday liquidity provision to support CHAPS averaged around £18 billion a day; that figure dates from the mid-2010s edition of that document and should be read as indicative of scale rather than current. CREST separately generated large volumes of liquidity through auto-collateralising repo.

The point of design principle is that the central bank is willing to supply intraday liquidity generously, against good collateral, because the alternative is a gridlocked high-value system. The constraint on a participant is therefore its collateral, its opportunity cost, and the regulatory treatment of intraday positions, rather than a hard cash balance.

■ Liquidity saving

The Liquidity Saving Mechanism, live in CHAPS since 15 April 2013 and retained in the renewed service, splits the flow into urgent and non-urgent streams via a central scheduler. Urgent payments settle individually and immediately. Non-urgent payments queue for a matching cycle, which as documented runs for around twenty seconds with roughly two-minute gaps between cycles, during which the algorithm identifies groups of broadly offsetting payments from different participants and settles them simultaneously.

The legal characterisation is unchanged: each payment settles gross, individually, with its own finality. The economic characterisation is that the participants economise, because the simultaneity means outflows are funded by matched inflows. The Bank measured combined intraday liquidity requirements falling from around £21.2 billion before implementation to around £16.9 billion afterwards, and attributed around £4 billion, roughly 20%, to the mechanism.

A practitioner reading this should note the operational consequence. Marking a payment non-urgent is cheaper in liquidity terms and slower in wall-clock terms, by a matter of minutes. For a completion payment at 1.50pm, urgent is the correct choice and the extra liquidity cost is the bank's problem. For an internal treasury sweep at 9.30am, it is not.

■ Behavioural rules

The Bank sets throughput targets for CHAPS direct participants, varying by category, and monitors adherence to them. The CHAPS Reference Manual distinguishes Category A direct participants as those whose average daily CHAPS volumes, sent and received, are greater than or equal to 40,000 payments, or whose CHAPS values are greater than or equal to 3.5% of total payment values, measured over a rolling twelve-month period. The specific percentage-by-time obligations sit in participant-facing operational documentation rather than the published reference manual, so no figures are quoted here.

Alongside throughput, the scheme sets availability obligations across the day. Participants must be open to receive by 8am and must send by 10am, which prevents the pathology of a participant that logs on at noon and drains everybody else's liquidity for two hours.

■ The CHAPS operating day

The published timetable, accurate at the time of writing, is as follows. All times are UK local time on a working day, meaning Monday to Friday excluding bank and public holidays in England and Wales.

Time	Event
06:00	CHAPS opens for settlement
08:00	Direct participants must be open to receive
10:00	Direct participants must have sent
17:40	Deadline for submission of customer payments
18:00	CHAPS closes for bank-to-bank payments
20:00	Latest contingency extension of the settlement day

RTGS itself does not close when CHAPS does. The Bank's 2026 consultation describes CHAPS as operating twelve hours a day, 06:00 to 18:00, Monday to Friday, with RTGS closing between 19:00 and 00:15 after end-of-day processes.

The customer and bank-to-bank distinction is a message-type distinction. Since the June 2023 migration, CHAPS uses ISO 20022. A customer credit transfer is a pacs.008; a financial institution transfer is a pacs.009. The 17:40 deadline is the pacs.008 deadline, and the twenty minutes to 18:00 exist so that participants can square their own positions between themselves after the customer flow has stopped.

■ Finality, in three named points

The most precise thing in the CHAPS rulebook, and the thing most often garbled in secondary sources, is that finality is not one event. The scheme defines three.

The Point of Entry is when the payment message has entered the Swift network and is acknowledged within that network by Swift. The Point of Irrevocability is after the point at which the settlement status in RTGS changes to "settling". The Point of Final Settlement is the point at which the settlement status in RTGS changes to "settled". Under contingency operation in MIRS Active Mode these points are modified, with finality occurring when the settlement status in MIRS changes to "settled".

Underpinning this is statute. CHAPS is a designated system under the Financial Markets and Insolvency (Settlement Finality) Regulations 1999. Designation is what makes finality survive an insolvency: a transfer order that has entered a designated system cannot be unwound by the ordinary retrospective effect of insolvency law, and the system's default rules take precedence. Without designation, the whole architecture would rest on contract and would be vulnerable to a liquidator arguing that payments made on the morning of a failure should be clawed back. With designation, the settled payment is settled, and the conveyancer can release the keys.

This is also why "recall" is a misnomer in CHAPS. A participant may transmit a request to the receiving participant asking for funds to be returned, and the receiving participant may co-operate, subject to its own obligations to its customer. That is a request for a new payment in the opposite direction. Nothing reverses.

■ Who participates, and what it costs

There are over 35 direct participants and several thousand financial institutions that make CHAPS payments through one of them. The Bank has progressively widened access: non-bank payment service providers became eligible for settlement accounts from 2017, and an omnibus account model was introduced in 2021 for certain arrangements. The Bank took over direct delivery of the CHAPS service itself in November 2017, having previously overseen a separate scheme company.

The economics of participation, at the time of writing, are published. Direct participants pay a per-item tariff of 42.7 pence per CHAPS payment sent for 2025 and 48.7 pence for 2026, plus an annual participation fee calculated from each participant's percentage share of total CHAPS values over a

defined historical window, applied to half of total CHAPS values. Joining is a capital project in its own right: the published onboarding fees include £30,000 to open a settlement account and £90,000 to join CHAPS, before a participant has spent anything on its own systems, testing, liquidity or staffing. Certain manual contingency operations carry fixed charges, such as £550 for a manual defund, and party-to-party transfers attract a £550 fixed administration fee if used more than twice a year.

Those numbers explain the shape of the market. The tariff is trivial per payment. The fixed costs of participation, technical conformance and round-the-clock operational cover are not, which is why the direct membership is a few dozen institutions and the indirect population is several thousand.

Against that, the end customer typically pays between £25 and £30 per CHAPS payment, and indirect participants have historically paid their correspondents in the region of £2 to £3 per payment with a maximum of about £30, according to figures published by the Bank. The gap between 48.7 pence and £25 is not primarily infrastructure cost. It is the sending bank's liquidity, its manual checks on high-value payments, its fraud controls and its acceptance of operational risk on an irreversible instruction.

■ What CHAPS is actually used for

Three families of use dominate.

Wholesale financial activity comes first by value. Financial institutions and the largest businesses use CHAPS to settle money market transactions and the sterling legs of foreign exchange trades. These are the payments that make the average transaction value £1.7 million and produce the extraordinary ratio of 91% of sterling value on 0.5% of volume.

Property is the use the public sees. CHAPS is the standard instrument for completion in England, Wales, Scotland and Northern Ireland conveyancing, because completion requires cleared, final, same-day funds and there is no other sterling rail that provides all three above £1 million, and none at all that provides finality of this legal quality. The Standard Conditions of Sale contemplate payment on completion by direct credit to the seller's solicitor's bank, which in practice means CHAPS.

Corporate treasury is the third. Time-critical supplier payments, tax payments, margin calls, intra-group funding and anything where a missed cut-off has a contractual or regulatory consequence.

The ISO 20022 enhanced data mandate has made the property use case explicit in the data. From 1 May 2025 the Bank mandated purpose codes and Legal Entity Identifiers for defined categories of CHAPS payment: payments between financial institutions, meaning all pacs.009 messages and pacs.008 messages where both ultimate parties fall within defined categories such as PRA-authorized deposit-takers, broker-dealers, Bank-supervised financial market infrastructures or RTGS account holders; and property payments, which must carry one of a defined set of codes.

Code	Meaning
PCOM	Property completion payment
PDEP	Property deposit
PLDS	Property loan disbursement
HLRP	Property loan repayment
HLST	Property loan settlement
PLRF	Property loan refinancing

The mandate applies to payments originated through channels controlled by the direct participant. The wider ISO 20022 programme continues: structured addresses were encouraged from November

2025 and unstructured addresses are to be rejected from November 2026, with a mandate for structured remittance data expected from November 2027.

■ The size of the thing

Measure	Figure	Period
CHAPS payments	52.7 million, a record, up 3.1%	2024
CHAPS value	£87.5 trillion	2024
Average daily value	£344.4 billion	2024
Average daily volume	207,609 payments	2024
Average payment value	£1.7 million	2024
Share of sterling payment value	91%	2024
Share of UK payment volume	around 0.5%	2024
Average daily value	£365.7 billion	2025, to Q3
Average daily volume	209,401 payments	2025, to Q3
Total daily settlement across RTGS	over £800 billion	2024/25

For comparison against the rails covered elsewhere in this volume: Bacs runs a three working-day cycle of submission, processing and entry, and settles on a deferred net basis; Faster Payments settles in seconds, runs 24/7, and caps a single transaction at £1 million; CHAPS is same-day, gross, final, and uncapped by the scheme.

■ When it stops: 20 October 2014

The best way to understand what RTGS is for is to read the report of the day it did not work, and the Bank commissioned and published one.

RTGS opened at 06:00 on Monday 20 October 2014. At 06:02 a system error was detected. The proximate cause was a set of configuration changes made over the weekend of 18 and 19 October to transfer a CHAPS membership from one bank to another. Those changes exposed three latent defects that had been sitting in the system unnoticed: one design defect introduced in April 2014, and two functional defects dating from the implementation of the contingency service. Nothing was hacked. Nothing was overloaded. A routine membership change met three sleeping bugs at the same moment.

The cause was identified between 07:00 and 07:32. A remediation plan was finalised at 10:30. A controlled restart began between 14:15 and 14:30. Settlement resumed at 15:15. The Bank extended the settlement day from its then-scheduled 16:20 close to 20:00. By 18:00 around 99% of payments had settled by both value and volume, and by 20:00 all of them had. The day's total was 142,759 CHAPS payments worth £289.3 billion, against a forecast of about 145,000 payments and £276 billion. In other words, the day completed. Nothing was lost. Nine hours of the country's high-value payment capability were simply unavailable, and then they were available again, and the backlog cleared in five hours.

The human cost landed on the property market, because property completions are the CHAPS use case with a hard deadline and vans outside. A Law Society survey with 157 responses recorded that 18% of transactions completed within the usual times, 27% were delayed by up to three hours, 24% by more than three hours, and 30% did not complete until the next day or later. Against a typical Monday of around 2,340 residential sales, the review estimated approximately 700 transactions may not have completed until the next day.

The most instructive part of the report is not the failure but the contingency decision. The Bank had, and has, a contingency service: MIRS, the Market Infrastructure Resiliency Service, operated by Swift

on alternative technology at geographically remote sites, designed to let CHAPS keep settling without falling back to a deferred net model. It was not invoked. The review sets out why. Switching to MIRS was believed to be a one-way move in practice, because failing back to RTGS intraweek was a manual process carrying an increased risk of error if attempted overnight. MIRS lacked full functionality, notably the Liquidity Saving Mechanism, and offered reduced resilience relative to RTGS. And at the 14:20 decision meeting the operational and technical team advised that the fix-forward plan was well advanced and RTGS should be running by 15:00 to 15:15, which it was.

That is a rational decision and it is also the anatomy of a near miss. The contingency existed and was not used because using it was expensive and irreversible, which is exactly the condition under which contingency arrangements decay. The independent review's recommendations, all of which the Bank accepted, addressed precisely that: test intraweek fallback from MIRS so that invoking it is not a one-way door; educate participants on MIRS scenarios so the decision is not made cold; build a loss-of-integrity contingency manual with reconciliation scripts and explicit MIRS decision criteria. Alongside those came governance and engineering changes, including reconstituting the RTGS board under the Deputy Governor for Markets and Banking, which first met in that form on 12 March 2015, deferring further functional changes pending business assurance, and building a proper testing regime with independent testing, regression testing and separated test and pre-production environments.

■ What "RTGS goes down" actually means

There is no second CHAPS. This is the uncomfortable structural fact. The £344 billion a day that moves through CHAPS has no alternative rail: Faster Payments cannot take a £4 billion money market payment, both because of the £1 million cap and because deferred net settlement of that value would create precisely the exposure CHAPS exists to eliminate. So resilience in CHAPS cannot be delivered by diversity of rails. It has to be delivered inside the one rail.

It is delivered in four layers. First, the primary RTGS infrastructure, now RT2, with its own redundancy. Second, MIRS, which is a genuinely different technology stack run by a different organisation in different locations, so that a defect in RTGS does not propagate. Third, the ability to extend the settlement day, to 20:00 as a contingency, which converts an availability failure into a lateness problem as long as the fault is fixed within the day. Fourth, manual and bilateral fallbacks such as party-to-party transfers and contingency transfers, which are slow, expensive, individually authorised and intended for a handful of critical payments rather than the day's flow.

The Bank now frames its operational performance in terms of impact tolerances for important business services rather than a single uptime number. In its 2024/25 annual report it documented five incidents between July 2024 and March 2025 and stated that impact tolerances were met in every case, with all payments settling by end of day. That framing is deliberate and worth understanding: for a system like this, the question is not "was there an outage" but "did the day complete". A ninety-minute interruption at 07:00 is an inconvenience. The same interruption at 16:30 is a housing chain that does not complete and a money market position that has to be rolled.

■ What changes next

Two changes are confirmed or consulted on at the time of writing, and both are worth tracking by anyone building on these rails.

The first is hours. On 24 February 2026 the Bank published a policy statement confirming an early morning extension: CHAPS settlement will begin at 01:30 rather than 06:00, targeted for September 2027 subject to confirmation with participants. Sending during the extended window will be optional for direct participants, but all participants will be able to receive from 01:30, and each institution de-

cides for itself when to credit received funds to customer accounts. The Bank will run an alert-and-respond support model in the early hours, with 24/7 cover retained for critical incidents.

On 18 May 2026 the Bank went further, consulting until 10 August 2026 on the path beyond that: weekend, specifically Sunday, and bank holiday settlement from 01:30 to 18:00 not before 2029; 22-hours-a-day operation across six days not before 2031; and a long-run target of either 22x7 or 23.5x7, with short daily closures for housekeeping. The stated drivers are innovation, including synchronisation and settlement of tokenised assets and stablecoin arrangements, better alignment with overseas markets for cross-border payments, and liquidity efficiency from more frequent settlement.

The second is data. The ISO 20022 enhanced data programme is not finished. Purpose codes and LEIs became mandatory for defined payment types from 1 May 2025; structured addresses become compulsory in effect from November 2026 when unstructured ones start being rejected; a structured remittance mandate is expected from November 2027. For anyone integrating with CHAPS through a direct participant, these are hard dates with rejection consequences, not aspirations.

■ A practitioner's summary

If you are choosing a rail, the decision tree is short. If the value exceeds £1 million, or the payment must be irrevocable on receipt, or a third party will act irreversibly on the credit the moment it lands, use CHAPS. If none of those is true, you are paying £25 to £30 for a certainty you do not need.

If you are operating CHAPS payments, five things matter more than the rest. Know your own bank's customer cut-off, which is earlier than the scheme's 17:40, and know that the conveyancing convention of 2pm is earlier still. Get the beneficiary details right the first time, because the rail has no reverse gear and your remedy is capped at £85,000 per claim for in-scope retail payments under the reimbursement rules that applied from 7 October 2024. Populate the purpose code correctly for property and financial institution payments, because since 1 May 2025 it is mandatory, not decorative. Understand that urgent and non-urgent are real settings with real consequences in minutes and in your bank's liquidity cost. And build your operational process on the assumption that one day in some year, RTGS will be unavailable at 11am, the day will be extended, and everything will settle by evening, because that is what the record actually shows happens.

The woman in the van gets her keys because a ledger at the Bank of England changed by £306,000 and cannot change back. Everything in this chapter, all the collateral, the matching cycles, the throughput targets, the contingency stack, the three named points of finality, exists to make that one sentence true 209,401 times a day.

38.98 Common wrong ideas

Wrong: CHAPS is the fastest UK payment rail, which is what the fee buys. **Right:** Faster Payments is faster, free and works at midnight on a Sunday; what CHAPS sells is finality in central bank money that cannot be undone.

Wrong: A bank needs a large pile of its own cash sitting idle to make CHAPS payments in the morning. **Right:** Sterling Monetary Framework participants generate intraday liquidity against pledged high-quality collateral with haircuts, so the constraint is eligible collateral and its opportunity cost, not owned cash.

Wrong: Every CHAPS payment settles strictly one at a time. **Right:** Legally each one does, but since 15 April 2013 the Liquidity Saving Mechanism settles groups of broadly offsetting non-urgent payments simultaneously, taking combined intraday requirements from around £21.2 billion to around £16.9 billion.

Wrong: The money in the seller's solicitor's client account is central bank money. **Right:** It is a claim on her bank; what settled at the Bank of England was a transfer between two banks, a causally linked but different event.

Wrong: The CHAPS deadline is 17:40. **Right:** That is the scheme deadline for customer payments; every bank sets an earlier customer cut-off, and the conveyancing convention drawn from the Standard Conditions of Sale puts the practical deadline at 2pm.

Wrong: Because a settled CHAPS payment is irrevocable, a fraud victim has no remedy. **Right:** The payment stays final, but since 7 October 2024 the Bank has applied reimbursement rules to in-scope retail CHAPS payments with a maximum of £85,000 per claim; the remedy is a different payment under a different rulebook.

Wrong: CHAPS has a maximum transaction value in the way Faster Payments does. **Right:** The scheme rules set no maximum at all, which is why a £4 billion money-market payment and a £306,000 house purchase travel on the same rail; what limits you is your own bank's policy.

Wrong: RTGS is a payment system. **Right:** RTGS is the Bank of England's accounting ledger, over £800 billion a day, carrying CREST and the net obligations of Bacs, Faster Payments, LINK, cheques and cards as well as CHAPS.

Wrong: A CHAPS payment can be recalled if you spot the error quickly. **Right:** A participant may request that the receiving participant return the funds, and it may refuse; a return is a new payment in the opposite direction, and nothing reverses.

Wrong: Resilience comes from being able to route round CHAPS on another rail. **Right:** There is no second CHAPS, so resilience is delivered inside the one rail through RT2, MIRS, extension of the settlement day to 20:00, and slow manual fallbacks for a handful of critical payments.

38.99 Chapter summary in 20 lines

1. CHAPS carried 52.7 million payments worth £87.5 trillion in 2024, roughly 0.5% of UK payment volume and 91% of sterling payment value, averaging £1.7 million a payment.
2. It is not fast, since Faster Payments is faster and free, and what it sells instead is finality in central bank money that cannot be undone.
3. Deferred net settlement is efficient because obligations cancel, but a failure before the net settles forces everybody's positions to be recalculated.
4. Real-time gross settlement removes that by settling each payment individually, in full, in central bank money, at the moment it is made.
5. The cost of gross settlement is liquidity, because a participant that sends before it receives needs a balance sitting ready to cover the difference.
6. Direct participants do not fund that from their own cash alone; they generate intraday liquidity against pledged collateral, extinguished by the end of the day.
7. Since 15 April 2013 the Liquidity Saving Mechanism has settled groups of offsetting non-urgent payments simultaneously, saving around £4 billion of intraday liquidity, about 20%.
8. Each of those payments still settles gross and individually in law, so CHAPS has borrowed netting's best trick without giving up gross settlement's guarantees.
9. RTGS is not a payment scheme but the Bank of England's ledger, on which over £800 billion settles daily, including CREST and the net positions of every retail rail.
10. The renewed core settlement engine, RT2, went live on 28 April 2025 alongside a revised tariff framework.

11. Finality is not one event: the rulebook names the Point of Entry, the Point of Irrevocability and the Point of Final Settlement.
12. Designation under the Financial Markets and Insolvency (Settlement Finality) Regulations 1999 is what makes that finality survive an insolvency, and therefore what lets the keys be released.
13. Tiering means the guarantee describes one leg only, with over 35 direct participants sitting above several thousand institutions reaching CHAPS through a correspondent on ordinary commercial credit terms.
14. Cut-offs come in layers — 17:40 at the scheme, earlier at each bank, and 2pm by conveyancing convention — and the earliest one governs your day.
15. Direct participants pay 48.7 pence per payment sent in 2026 plus a participation fee, £30,000 to open a settlement account and £90,000 to join, while the end customer pays £25 to £30.
16. That gap is not infrastructure cost but liquidity, manual high-value checks, fraud controls and the acceptance of operational risk on an irreversible instruction.
17. Wholesale financial activity dominates by value, property completion is the use the public sees, and corporate treasury is the third family.
18. Since 1 May 2025 purpose codes and Legal Entity Identifiers have been mandatory for defined categories, with PCOM, PDEP, PLDS, HLRP, HLST and PLRF naming the property cases.
19. The outage of 20 October 2014 shows the shape of failure: a routine membership change met three latent defects, nine hours were lost, MIRS was not invoked because using it was effectively a one-way door, and the day still completed by 20:00.
20. Everything in the system — collateral, matching cycles, throughput targets, the contingency stack, the three named points of finality — exists so that a ledger entry at the Bank of England changes and cannot change back.

Sources: Bank of England published material, principally the CHAPS service pages (operating hours, direct participant numbers, uses, and reimbursement position), “A brief introduction to the Real-Time Gross Settlement system and CHAPS”, the payment and settlement statistics pages, the CHAPS Reference Manual version 28 April 2025 (points of entry, irrevocability and final settlement, Category A definition, contingency modes), the RTGS and CHAPS Annual Report 2024/25 (RT2 go-live, ISO 20022 milestones, incidents and impact tolerances, daily RTGS settlement value), the 2023 self-assessment against the Principles for Financial Market Infrastructures (Settlement Finality Regulations designation, intraday liquidity and collateral, throughput targets, MIRS), the RTGS service description (daily timetable and liquidity provision, mid-2010s edition), the RTGS and CHAPS fees pages (2025 and 2026 tariffs, participation and onboarding fees), the Quarterly Bulletin articles “How has the Liquidity Saving Mechanism reduced banks’ intraday liquidity costs in CHAPS?” (2014) and “Tiering in CHAPS” (2013), the guidance on mandating ISO 20022 enhanced data in CHAPS (purpose codes and LEIs), the policy statement “Extending RTGS and CHAPS settlement hours – early morning extension” (24 February 2026) and the consultation “Extending RTGS and CHAPS settlement hours – next steps towards near 24x7 settlement” (18 May 2026), together with the independent review of the RTGS outage on 20 October 2014 and the Bank’s published response (March 2015). Conveyancing completion practice from the Law Society Standard Conditions of Sale, fifth edition. All figures, fees and timings accurate at time of writing, August 2026.

Addressing: Sort Codes, IBANs and Confirmation of Payee

39.0 What this chapter gives you

1. You will be able to explain why a payment to a valid account in somebody else's name arrives without a single system in the chain failing.
2. You will be able to take a sort code and account number apart and say exactly what each half addresses, and what neither of them proves.
3. You will be able to run a MOD11 and a DBLAL modulus check by hand, and say why summing the products instead of the digits of the products gives the wrong answer.
4. You will be able to state what a passing modulus check is worth, given that a MOD10 rule admits about one in ten arbitrary numbers and a MOD11 rule about one in eleven.
5. You will be able to generate and validate the two check digits of a British IBAN, and explain why a correctly formed IBAN always leaves a remainder of exactly 1 modulo 97.
6. You will be able to explain why a BIC can never address an account, and why its presence in a directory does not prove network reachability.
7. You will be able to read the Confirmation of Payee reason codes and say which ones signal the shape of invoice redirection fraud.
8. You will be able to explain why allocation of a sorting code is not reachability, and why the EISCD needs refreshing weekly rather than inferring institutions from digit ranges.
9. You will be able to say why a collection account moves the real address into an eighteen-character reference field with no check digits and no validation.
10. You will be able to explain why over 99 per cent of Faster Payments initiators now check names and APP losses still rose 19 per cent in 2025.

A payroll clerk at a scaffolding firm in Wolverhampton opens an email from a supplier she has paid twenty-two times. The invoice is for £61,400. The reference is the one she expects, the letterhead is right, the VAT number is right, and the message arrives inside the same thread as the last four invoices. Exactly one thing has changed. At the bottom of the page the bank details are different, with a short apologetic note: new banking provider, please update your records before you run Friday's payment.

She updates the records. The payment goes out on Friday. It arrives, in seconds, at an account that exists, at a bank that is regulated, in the name of somebody who is not the supplier. Nothing failed. Every system in the chain did exactly what it was built to do, which was to take an address and deliver to it.

This chapter is about that address: what it is made of, what the numbers inside it prove, what they cannot prove, and what was bolted on top of the whole arrangement in 2020 to try to catch precisely the

Friday afternoon described above. All limits, timings, figures and version numbers here are accurate at the time of writing, August 2026, and payments infrastructure changes.

The plain version

Think about how a letter gets to a house.

You write two different kinds of thing on an envelope. You write a name, and you write an address. Only one of them is used. The postal system reads the postcode and the house number, sorts the letter through a series of increasingly local buildings, and pushes it through a letterbox. Nobody at any stage compares the name on the envelope with the name of the person who lives there. If you put the right name and the wrong house number, the letter goes to the wrong house. If you put the wrong name and the right house number, the letter goes to the right house.

Bank payments were built the same way, and for the same reason: it is far faster to sort by number than to check by name.

■ The two numbers

In Britain the address has two parts.

The first is the **sort code**: six digits, usually printed in three pairs, like 20-00-00. Think of it as the postcode. It says which organisation, and historically which branch of that organisation, the money is going to.

The second is the **account number**: eight digits, like 55779911. Think of it as the house number on that street. It says which specific customer, inside that organisation, gets the money.

Together they are fourteen digits, and those fourteen digits are the entire address. There is no third element. When you pay a plumber £2,400 from your phone, you type a name into a box, but the name is a label for your own records. The money is moved by the fourteen digits and nothing else.

■ The trick hidden inside the numbers

Now, fourteen digits is a lot of digits to type, and people mistype. If one digit in a postcode is wrong, the letter usually bounces, because most wrong postcodes do not exist. But almost every combination of fourteen digits looks like a plausible bank address. So the banks did something clever: they built a hidden test into the account number itself.

Here is the idea in miniature. Suppose account numbers were only four digits long, and suppose the rule was this: multiply the first digit by 4, the second by 3, the third by 2 and the fourth by 1, add up the answers, and the total must divide exactly by 11.

Take the number 1325.

Digit	1	3	2	5
Multiplier	4	3	2	1
Answer	4	9	4	5

Four plus nine plus four plus five is 22. Twenty-two divided by eleven is exactly two, with nothing left over. So 1325 passes.

Now mistype it. Swap two digits, so you type 1352: 4 plus 9 plus 10 plus 2 is 25, which does not divide by 11. Rejected. Change one digit, so you type 1825: 4 plus 24 plus 4 plus 5 is 37, which does not divide by 11. Rejected.

That is the whole idea, and the real system is the same idea made bigger. The real test uses all fourteen digits, sort code and account number together, and each bank has its own set of multipliers. It is called a **modulus check**, and it runs before the payment screen will even let you press Next. It catches the overwhelming majority of typing mistakes, which is exactly what it was designed for and exactly the limit of what it does.

■ Writing the address for abroad

Every country invented its own address format, and they do not fit together. A German account is eighteen digits arranged differently; a French one has letters in it. So Europe agreed on a standard envelope you can put any country's domestic address inside, called an **IBAN**, which stands for International Bank Account Number.

A British IBAN looks like this:

GB29 NWBK 6016 1331 9268 19

Take it apart and there are no surprises in it:

Piece	Meaning
GB	The country: Great Britain
29	Two check digits, explained in a moment
NWBK	A short code for the bank, in this case NatWest
601613	The sort code you already know
31926819	The account number you already know

So a British IBAN is a British bank address with a country label and a checksum stapled to the front. There is nothing new inside it.

The two check digits are the interesting part, because they protect the whole string rather than just the account number. The recipe is odd but it works. First, move the four characters at the front to the back. Then turn every letter into a number, where A is 10, B is 11, and so on up to Z as 35. You now have one very long number. Divide it by 97. If the IBAN was typed correctly, the remainder is exactly 1. Every time. For any country.

Working it the other way produces the check digits in the first place. Start with the account details and two zeros where the check digits will go, convert to a number the same way, divide by 97, and you get a remainder of 69. Subtract that from 98 and you get 29. That is where the 29 in GB29 comes from. It is not a reference number, it is not chosen by the bank, and it is not a secret. It is arithmetic.

There is one more code you will see on international paperwork, the **BIC** or SWIFT code, which looks like NWBKGB2L. It names the bank, not the account: four letters for the institution, two for the country, two for the town. It is the equivalent of writing "Barclays, Leeds" on an envelope. Useful, but on its own it will not get money to a person.

■ The name on the envelope, finally

Return to the scaffolding firm in Wolverhampton. Every system did its job. The digits were valid, the sort code existed, the account existed, and the money arrived. What nobody did was look at the name.

In 2020, British banks started doing exactly that. The service is called **Confirmation of Payee**. When you type in a new payee, your bank sends the sort code, the account number and the name you typed to the bank at the other end, and asks a single question: is this the name on that account? The other

bank answers in about a second, and you get one of four answers on your screen before you have sent anything.

Yes. The name matches. Carry on.

Close. It is nearly right, and here is the actual name on the account: you typed “Dan Okafor” and the account says “Daniel Okafor”. Check it and decide.

No. The name is wrong. You are shown no name at all, only a warning, because telling you the real name would let anyone with a keyboard harvest account holders’ names.

Cannot check. The other bank did not answer in time, or the account holder has opted out, or that kind of account is not covered.

The scale of the thing it is fighting is not small. In 2025, people and businesses in the United Kingdom lost £576.4 million to payments they authorised themselves after being deceived, which was 19 per cent more than the year before, and about a third of all payment fraud losses. Banks reimbursed £354.3 million of it, roughly 61 per cent. That is the number Confirmation of Payee exists to bring down: not money stolen from accounts, but money pushed out of them by the account holder, into an address that was perfectly valid and completely wrong.

Where the plain version stops being true

The envelope is a good picture of routing and a poor picture of nearly everything else. Four corrections.

The first correction is that the address is not a place, is not permanent, and is very often not unique to one person. A house stands still. A sort code and account number pair is a row in a directory, and directories are edited constantly. Sort codes move between institutions when a bank is bought or a brand retired. Crucially, a sort code does not reliably tell you which bank holds the money: under indirect access, a sponsoring bank allocates sort codes from its own range to smaller firms that settle through it, so the institution named against the sort code and the institution where your payee actually banks can be two different companies. Meanwhile, the Current Account Switch Service moves customers between banks and redirects payments arriving at the old address to the new one, silently, for years. And a very large number of real-world accounts are **collection accounts**: one sort code and account number belonging to a platform, a broker or a savings marketplace, into which tens of thousands of customers pay, with the individual identified only by the payment reference. In that arrangement the fourteen digits are the building and the reference field is the flat number, and the reference field has no check digits, no validation and no protection whatsoever.

The second correction is that every check described above tests the number against itself, and none of them tests the number against reality. This is the single most expensive misunderstanding in the field. A modulus check answers one question: could this pattern of digits be a legitimate combination for this sort code? It does not say the account exists, is open, has ever been issued, or belongs to anyone. Worse, passing is weak evidence by construction. A modulus 10 check accepts roughly one in ten random numbers by chance; a modulus 11 check roughly one in eleven. Failing is strong evidence; passing is barely evidence at all. The IBAN check digits are weaker still in this respect: they verify that a string was transcribed correctly, and say nothing about whether the account exists or whether the bank code and sort code inside it even agree. And for some sort codes there is no arithmetic at all. The published rules are explicit that when a sort code falls outside every range in the weight table, no modulus check can be performed and the details should be presumed valid unless other evidence suggests otherwise. Presumed valid. That is the honest position of the whole domestic validation layer, and it sits a long way from the confidence implied by a green tick.

The third correction is that Confirmation of Payee is advisory, partial and deliberately fuzzy, and is none of the three things people assume. It is not a block: the payer sees the result and can carry on regardless, and in a competently run scam the payer has already been told to expect a mismatch and given a reason for it. It is not comprehensive: the responding bank can decline to answer, the account holder can opt out, certain account types are not supported, and payments where either party is outside the United Kingdom are outside the requirement entirely. And it is not exact. Name matching is intentionally tolerant, because real account names contain initials, married names, trading names, punctuation and the ordinary noise of human typing, and a literal comparison would reject correct payments all day. The “close match” answer exists because rejecting every payment to “J Smith” when the account says “John Smith” would have made the service unusable. Somebody at each responding institution tuned that tolerance, and the tuning is not published. Most fundamentally, the service checks whether a name matches an account. It cannot check whether the person behind that account is who they claim to be, or whether the goods will ever arrive. Against purchase scams, where the payee’s real name really is on the account, it does almost nothing.

The fourth correction is that there is no single worldwide address format, the IBAN is not it, and Britain does not use IBANs domestically anyway. IBAN covers Europe and a long tail of adopters elsewhere; the United States, Canada, India, China, Japan and Australia do not use it. A domestic Faster Payment, a domestic CHAPS payment and a Bacs Direct Debit are every one of them addressed with a bare sort code and account number, and a British consumer can bank for forty years without seeing an IBAN. Nor is the BIC an address in the same sense: it identifies an institution, or a branch or department within one, and never an account. An IBAN without a BIC is usually a complete destination. A BIC without an account identifier is a building with no room number.

The technical version

■ The domestic identifier pair

The addressing primitive for all three principal sterling interbank rails — Bacs, Faster Payments and CHAPS — is a six-digit sorting code paired with an account number. Nothing in the domestic schemes routes on name, and nothing routes on IBAN.

Sorting codes are allocated by Pay.UK on behalf of the industry. Direct participants receive codes through the onboarding process for the payment system concerned; indirect participants receive codes sub-allocated by their sponsoring direct participant, and settle through that sponsor’s agency arrangement. Pay.UK also allocates sorting codes to firms that issue GB IBANs for international payments even where those firms do not participate in a UK payment system at all. The operational rules are overseen by the Clearing Codes Operational Group, whose membership includes direct participants, Pay.UK and Vocalink.

The governing rule for practitioners is that allocation is not reachability. A sorting code must be separately registered with each payment system before it can be used in that system. A code registered for Bacs Direct Credit is not thereby usable for Direct Debit; a code registered for Faster Payments is not thereby usable for CHAPS. Systems that validate a sorting code’s existence but not its scheme registration produce a characteristic class of rejection that looks like a data error and is in fact a reachability error.

■ The Extended Industry Sort Code Directory

The public extract of the central sorting code database is the **Extended Industry Sort Code Directory (EISCD)**, successor to the Industry Sort Code Directory, published under Pay.UK's governance and distributed by Vocalink. It is the authority, and no heuristic about digit ranges should ever be used in its place.

For each sorting code the EISCD carries, among other fields, the branch title, the postal address and telephone number, the associated BIC, and clearing participation records showing which schemes that code is registered for and which settlement member it clears through. Subscription is on a quarterly fee basis with unlimited downloads; the guidance is that monthly refresh is the minimum and weekly is advisable, and many originators take it through their Bacs software rather than directly. For a Bacs originator or approved supplier this file is not optional infrastructure. It is the difference between rejecting a payment at capture and discovering the problem in an ARUCS or ARUDD report three working days later.

Historically the leading digits carry institutional meaning: blocks were assigned to the clearing banks when the six-digit format was introduced, and the legacy is visible in the ranges the large groups still occupy. It is a legacy and not a rule. New entrants have been allocated ranges with no historical logic, groups have merged and codes have transferred. Any code that infers the institution from the first two digits will eventually be wrong, silently, in production.

■ Account number normalisation

The account number is nominally eight digits, and institutions have never been uniform about it. Before any validation is attempted, the account number must be normalised to eight digits according to the published rules. Shorter numbers are zero-prefixed. Several institutions issue nine- or ten-digit numbers and specify which eight digits are significant, and in one case specify that a digit of the account number displaces a digit of the sorting code before checking. Building societies add a further layer: many hold balances against a **roll number** rather than an account number, with payments routed to a single society account and allocated internally from the reference field. Roll numbers have institution-specific formats, cannot be modulus checked, and are the reason the Confirmation of Payee schema carries a secondary identification field at all.

■ Modulus checking

Modulus checking is the domestic self-consistency test on the fourteen-digit pair. The specification is published by Vocalink as *Validating Account Numbers: UK Modulus Checking*; at the time of writing the current release is version 8.70, dated 10 February 2026. Two reference files accompany it: `valacdos.txt`, the weighting table, and `scsubtab.txt`, the sorting code substitution table used by certain exceptions.

Positions are named by letter throughout the specification. The six sorting code digits are `u v w x y z`; the eight account number digits are `a b c d e f g h`. Every rule is expressed against those fourteen positions.

Rows in `valacdos.txt` are fixed-length, space-separated, and consist of a sorting code range start, a range end, a modulus check type, fourteen weights for `u` through `h`, and an optional exception number. A row therefore looks like this in shape:

```
010004 016715 MOD11 0 0 0 0 0 0 8 7 6 5 4 3 2 1
```

A sorting code may appear in two rows. That means two checks are specified, and both are ordinarily required to pass, except where a numbered exception states that either check passing is sufficient.

Three algorithms are defined.

Type	Method
MOD10	Multiply each of the fourteen digits by its weight, sum the products, divide by 10. Remainder zero passes.
MOD11	Multiply each of the fourteen digits by its weight, sum the products, divide by 11. Remainder zero passes.
DBLAL	Multiply each digit by its weight, then sum the <i>individual digits</i> of every product rather than the products themselves, and divide the total by 10. Remainder zero passes.

The digit-summing step in DBLAL is where most reimplementations go wrong, and it is worth doing once by hand.

Worked MOD11. Sorting code 000000, account number 58177632, weights 0 0 0 0 0 0 7 5 8 3 4 6 2 1.

Position	u	v	w	x	y	z	a	b	c	d	e	f	g	h
Digit	0	0	0	0	0	0	5	8	1	7	7	6	3	2
Weight	0	0	0	0	0	0	7	5	8	3	4	6	2	1
Product	0	0	0	0	0	0	35	40	8	21	28	36	6	2

The sum is 176. 176 divided by 11 is 16 with no remainder. The check passes.

Worked DBLAL. Sorting code 499273, account number 12345678, weights 2 1 2 1 2 1 2 1 2 1 2 1 2 1 2 1.

Position	u	v	w	x	y	z	a	b	c	d	e	f	g	h
Digit	4	9	9	2	7	3	1	2	3	4	5	6	7	8
Weight	2	1	2	1	2	1	2	1	2	1	2	1	2	1
Product	8	9	18	2	14	3	2	2	6	4	10	6	14	8
Digits summed	8	9	9	2	5	3	2	2	6	4	1	6	5	8

The sum of the bottom row is 70. 70 divided by 10 is 7 with no remainder. The check passes. Note that summing the *products* instead would give 106, and 106 fails. The two routes disagree, and only one of them is the specification.

Fourteen numbered exceptions modify the standard procedure. They exist because individual institutions carried their own pre-existing check schemes into the industry standard and were not going to renumber their customers. Summarised, and always to be implemented from the current specification rather than from any summary:

Exception	Effect
1	Add 27 to the total before dividing.

Exception	Effect
2 and 9	Alternative weightings selected by the values of a and g, for certain euro accounts.
3	Skip the double alternate check where c is 6 or 9.
4	The remainder must equal the two-digit check value formed by g and h.
5	Substitute a sorting code from scsubtab.txt and apply specific check-digit logic.
6	Foreign currency accounts: treat as valid where a is 4 to 8 and g equals h.
7	Zeroise the weights for positions u to b where g is 9.
8	Substitute sorting code 090126 for checking purposes only.
10 and 11	Either check passing deems the account valid.
12 and 13	Either check passing deems the account valid.
14	If the first MOD11 check fails and h is 0, 1 or 9, drop h, prefix a zero, and recheck.

Exception 14 is worth working through because it shows how much of the specification is history rather than mathematics. Sorting code 180002, weights 0 0 0 0 0 8 7 6 5 4 3 2 1.

Account 98093517 gives products 72, 56, 0, 45, 12, 15, 2 and 7 across positions a to h, summing to 209. 209 divided by 11 is 19 exactly. It passes on the first check and exception 14 is never invoked.

Account 00000190 sums to 21. 21 divided by 11 leaves a remainder of 10, so the first check fails. Because h is 0, exception 14 applies: drop the trailing digit and prefix a zero to make 00000019. That sums to 11, which divides by 11 exactly. The account passes.

Two properties of modulus checking must be stated plainly to anyone building on it. First, the specification itself directs that where a sorting code is not found in the weighting table, no check can be performed and the details should be presumed valid. Second, a passing check carries very little information. A MOD10 rule admits about one in ten arbitrary numbers and a MOD11 rule about one in eleven. Modulus checking is a typo filter of very high value at the point of capture and a fraud control of no value whatsoever.

■ The IBAN

The IBAN is specified by **ISO 13616**, with Swift acting as the registration authority and publishing the IBAN Registry that lists each participating country's format.

The United Kingdom's entry is:

Attribute	Value
IBAN structure	GB2!n4!a6!n8!n
IBAN length	22
BBAN structure	4!a6!n8!n
BBAN length	18
Bank identifier	Positions 1 to 4 of the BBAN, 4!a
Branch identifier	Positions 5 to 10 of the BBAN, 6!n
Registry example	GB29NWBK60161331926819

Read in the registry's notation, 2!n means exactly two numeric characters, 4!a exactly four alphabetic, and ! denotes a fixed rather than maximum length. So a GB IBAN is: the literal country code GB, two check digits, a four-letter bank identifier which is conventionally the first four characters of the institution's BIC, the six-digit sorting code as the branch identifier, and the eight-digit account number. The sorting code and account number are carried unaltered. A GB IBAN contains no information a domestic payment instruction does not already have.

Other formats vary in length and content, which is why length-based validation alone is insufficient:

Country	Structure	Length	Registry example
United Kingdom	GB2!n4!a6!n8!n	22	GB29NWBK60161331926819
Germany	DE2!n8!n10!n	22	DE89370400440532013000
France	FR2!n5!n5!n11!c2!n	27	FR1420041010050500013M02606
Netherlands	NL2!n4!a10!n	18	NL91ABNA0417164300

The check digits are computed under **ISO/IEC 7064, MOD 97-10**. The generation procedure is:

1. Assemble the BBAN, then append the country code and the two literal characters 00.
2. Replace each letter with two digits: A is 10, B is 11, through to Z as 35.
3. Interpret the result as a single integer and take it modulo 97.
4. Subtract the remainder from 98. If the answer is a single digit, pad it to two digits with a leading zero.

Worked through for the registry example. The BBAN is NWBK60161331926819. Appending GB00 gives NWBK60161331926819GB00. Converting the letters — N is 23, W is 32, B is 11, K is 20, G is 16, B is 11 — produces:

2332112060161331926819161100

That integer modulo 97 is 69. 98 minus 69 is 29, giving the check digits 29 and the complete IBAN GB29NWBK60161331926819.

Validation is the same arithmetic run the other way. Move the first four characters to the end, convert letters to digits, and take the result modulo 97. A correctly formed IBAN always yields exactly 1. For the example, rearranging gives NWBK60161331926819GB29, which converts to 2332112060161331926819161129, and that modulo 97 is 1.

That integer is twenty-eight digits long and will overflow a 64-bit type, so implementations that do not use arbitrary-precision arithmetic must compute the modulus piecewise. The standard technique takes the first nine digits, reduces modulo 97, then repeatedly prefixes the running remainder to the next seven digits and reduces again. For the string above:

Step	Value reduced	Remainder
1	233211206	23
2	230161331	22
3	229268191	58
4	5861100	69

Which is the 69 the generation step needs. Nine and seven are chosen because the largest intermediate value stays inside the range of a 32-bit signed integer.

MOD 97-10 detects all single-character substitutions and all transpositions of adjacent characters, which is the entire design goal. It detects nothing about the account behind the string.

■ The BIC

The BIC is specified by **ISO 9362**, with Swift as registration authority. Since the 2022 revision the expansion is Business Identifier Code rather than Bank Identifier Code, reflecting that non-bank entities hold them.

The structure is four characters of business party prefix, two characters of ISO 3166-1 country code, and two characters of business party suffix, giving the eight-character BIC8. An optional three-character branch identifier extends this to the eleven-character BIC11, used to identify a specific location, department, service or unit of the same business party. By convention XXX in the branch position denotes the head office, and a BIC8 and the corresponding BIC11 ending XXX address the same party.

Swift distinguishes **connected** BICs, which have access to the Swift network, from **non-connected** BICs, registered for reference purposes only and with no right or authorisation to exchange messages over it. A BIC present in a directory is therefore not proof of reachability, any more than an EISCD sorting code is proof of scheme registration. The second character of the business party suffix carries a convention of its own, and Swift has published enough about it to settle the point. Swift's information paper on the ISO 9362 implementation states that addresses ending in a zero in position eight are used only for the Test and Training service, that they are not published in the directory of BICs, and that this convention will not change. A one in the same position historically meant the BIC was not connected, and Swift's published BIC Policy still illustrates a non-connected BIC as an eight-character BIC with a one in the eighth position. The connectivity reading, however, is no longer safe to compute. Since November 2018 an organisation connecting or disconnecting no longer needs a new BIC, and Swift's own guidance is that users should rely on the network connectivity status attribute in the SWIFTRef directories rather than on the intelligence once carried in position eight.

The essential point is categorical: the BIC identifies an institution or a unit within one, never an account. In cross-border messaging it names the parties and the route; the account is named by the IBAN or by a domestic identifier carried alongside.

■ Confirmation of Payee

Confirmation of Payee is an account name-checking overlay operated under Pay.UK's standards, launched in 2020. It is not a payment system and it moves no money. Architecturally it is a peer-to-peer API service: there is no central switch, only a directory identifying the organisations in the ecosystem, and requests pass directly between the requesting PSP and the responding PSP. Account-holding PSPs may build directly or use a technical service provider; smaller PSPs may reach the service through a CoP aggregator, of which Pay.UK lists nine approved.

The service is payment-system agnostic: Pay.UK describes it as applicable to domestic UK payments across CHAPS, Faster Payments and Bacs, covering both credit and debit flows, and a companion capability, Payer Name Verification, extends name checking to the setting up of Bacs Direct Debits.

A request carries a small, fixed set of fields:

Field	Content
SchemeName	SortCodeAccountNumber
Identification	Fourteen numerics: six-digit sorting code followed by eight-digit account number

Field	Content
Name	The payee name as entered, up to 140 characters
AccountType	Personal or Business
SecondaryIdentification	Optional secondary reference data, up to 140 characters

The matching is performed by the responding PSP against its own records; the requester never receives the account name on a negative result. The response carries a boolean match indicator and, where the match is not a straight yes, a reason code. The codes in use are:

Code	Meaning
ANNM	The responder performed the match and confirms the name does not match
MBAM	Close match; the responder returns the actual account name
BANM	A personal account was expected; a business account was found, with a name match
PANM	A business account was expected; a personal account was found, with a name match
BAMM	A personal account was expected; a business account was found, with a close name match
PAMM	A business account was expected; a personal account was found, with a close name match
AC01	The account does not exist in the responder's books
IVCR	The secondary reference data did not locate the account
ACNS	The account is not supported for CoP
OPTO	The payee has opted out of CoP
CASS	The account has been switched through the Current Account Switch Service
SCNS	The sorting code does not belong to the responding participant
SCNF	Sorting code not found at the routing layer

The account type distinction and the four codes covering it are frequently overlooked in implementation and are among the most operationally useful signals the service produces. A payment addressed to a business but landing on a personal account is a common shape for invoice redirection fraud.

■ Phase 1, Phase 2 and the regulatory timetable

Phase 1 went live in 2020 and covered the six largest banking groups, which between them accounted for around 90 per cent of Faster Payments and CHAPS transactions. It required full Open Banking membership to participate and, structurally, only supported PSPs holding their own unique sorting codes and identifying customers by sorting code and account number alone. That left two gaps: firms without their own sorting code, and accounts identified by something else, such as building society roll numbers or credit card numbers.

Phase 2 addressed both. A dedicated CoP-only role profile was created in the Open Banking Directory so that a PSP could enrol for name checking without taking on the full Open Banking estate, and

the standard was extended to carry **secondary reference data (SRD)** so that roll numbers and card numbers could identify the target account. Pay.UK's guidance gives the two canonical SRD cases as a sixteen-digit card number when paying a credit card bill and a roll number for some building society accounts.

The regulatory backing sits with the Payment Systems Regulator. Specific Direction 10 mandated Phase 1 for the six largest groups. **Specific Direction 17**, issued in October 2022, expanded the mandate in two tranches: Group 1 PSPs, named in the direction, by 31 October 2023, and Group 2 PSPs by 31 October 2024. Group 2 is defined by having a unique sorting code, or being a building society, and sending funds for customers over CHAPS or the Faster Payment Service into accounts consumers and businesses can pay into and be paid from. Directed firms must implement both the sending and the responding sides unless they hold only account types where responding alone is appropriate.

The exclusions matter as much as the scope. The direction does not apply where either party is outside the United Kingdom, to loan repayments including mortgages, credit cards and buy-now-pay-later, to a PSP moving its own money, or to bulk and batch payments submitted by indirect PSPs.

On coverage, the PSR's compliance report of 18 November 2025 recorded that over 99 per cent of organisations initiating Faster Payments transactions in the United Kingdom offer CoP checks, that over 320 organisations offer the service, and that over two million checks are completed each day.

■ Where the address sits on each rail

Addressing is the same fourteen digits everywhere, but the container differs.

Bacs carries payment instructions in the Standard 18 format: fixed-position records of 100 characters, with an optional extension for multi-processing files. The destination sorting code occupies positions 1 to 6 and the destination account number positions 7 to 14, so the fourteen-digit address is literally the first fourteen characters of the record. Position 15 is the destination account type, positions 16 to 17 the transaction code, positions 18 to 23 and 24 to 31 the originating sorting code and account number, positions 32 to 35 a four-character free format field, positions 36 to 46 the amount in pence, positions 47 to 64 the originator's name in eighteen characters, and positions 65 to 82 the eighteen-character payment reference. Multi-processing files carry a processing date at positions 101 to 106 in the format bYYDDD.

The eighteen-character reference is the pinch point. In a collection-account architecture the reference is the real address, and eighteen characters with no validation is all the protection it has. Pay.UK also publishes an ISO 20022 to Standard 18 translation guide for originators moving between the two.

CHAPS has carried ISO 20022 messages since 19 June 2023, with customer credit transfers as pacs .008. The Bank of England has been progressively mandating enhanced data within those messages. From 1 May 2025, Legal Entity Identifiers became mandatory for payments between financial institutions, and Purpose Codes became mandatory both for payments between financial institutions and for property transactions, using six specified codes: HLRP, HLST, PLDS, PDEP, PCOM and PLRF. Hybrid addresses — town name, country and up to two unstructured lines — were enabled from November 2025, and fully unstructured addresses are to be rejected by the CHAPS validation library from November 2026. The Bank has consulted on extending mandatory enhanced data to all initiation channels and Purpose Codes to all CHAPS payments from November 2027. The direction of travel is unambiguous: the address is being made structured and machine-readable, party by party and field by field.

Faster Payments carries the same sorting code and account number in scheme messages, with Pay.UK publishing an ISO 20022 standards library as part of the migration away from the legacy message

format. The relevant scheme constraint for addressing purposes is the £1 million maximum per transaction; individual PSPs set their own lower limits by channel, account type and payment type, and those limits, not the scheme limit, are what a customer meets in practice.

Open Banking payment initiation is an API layer over these rails, chiefly Faster Payments. It changes who constructs the instruction, not how the destination is addressed.

■ APP fraud and the reimbursement regime

Confirmation of Payee was introduced because addressing by number alone left one specific hole, and the hole has a name: authorised push payment fraud, in which the customer is deceived into instructing a payment to an account the fraudster controls. Every technical control described in this chapter is satisfied in an APP scam. The modulus check passes. The IBAN, if there is one, validates. The sorting code is registered. The account exists.

The regulatory response went beyond name checking. Since **7 October 2024**, the PSR's reimbursement requirement has applied to in-scope APP scam payments made over Faster Payments and CHAPS. The consolidated policy sets a maximum of £85,000 per claim, permits a claim excess of up to £100 which a PSP may set lower or not apply at all, and requires reimbursement within five business days, with defined circumstances in which the clock may be stopped and an absolute backstop of the thirty-fifth business day following the report. The cost is split evenly: receiving PSPs must pay sending PSPs 50 per cent of the reimbursement made to the consumer. A consumer standard of caution exception allows a claim to be refused for gross negligence against four defined requirements, including responding to interventions, reporting within thirteen months, sharing information and reporting to the police. Neither the standard of caution exception nor the excess may be applied to vulnerable consumers.

The scale, from UK Finance's 2026 Annual Fraud Report covering 2025: total payment fraud losses of £1.28 billion across more than four million cases, of which APP fraud accounted for £576.4 million, up 19 per cent year on year and 32 per cent of all losses. Investment fraud within that rose 40 per cent to £221.5 million. Around two-thirds of APP fraud originated online and 17 per cent through telecommunications. Firms reimbursed £354.3 million, some 61 per cent of APP losses.

Those two facts sit together and explain the shape of the whole control stack. Coverage of name checking is above 99 per cent of Faster Payments initiators, and APP losses still rose 19 per cent. Confirmation of Payee closed the misdirection problem. It did not close the deception problem, because it was never capable of doing so.

■ The summary a practitioner should carry

A UK domestic payment is addressed by six digits of sorting code and eight digits of account number, and by nothing else. The sorting code is allocated by Pay.UK, sub-allocated by sponsors for indirect participants, and must be separately registered with each payment system it is to be used in; the EISCD is the authority for both ownership and scheme registration, and it should be refreshed weekly. Account numbers must be normalised to eight digits before validation.

Modulus checking, specified by Vocalink and driven by `valacdos.txt` and `scsubtab.txt`, runs one or two of MOD10, MOD11 or DBLAL against the fourteen positions u to h, subject to fourteen numbered exceptions. It is a high-value typo filter, it admits roughly one in ten or one in eleven arbitrary numbers by chance, and where a sorting code is absent from the table the specification directs that the details be presumed valid.

The IBAN under ISO 13616 wraps the same domestic pair in a country code and two MOD 97-10 check digits computed under ISO/IEC 7064; validation is the rearranged string modulo 97 equalling 1, and the check protects transcription and nothing else. The BIC under ISO 9362 names an institution or a unit within one, never an account, and its presence in a directory does not imply network connectivity.

Confirmation of Payee is a peer-to-peer API overlay, mandated in tranches under PSR Specific Direction 17 to 31 October 2023 and 31 October 2024, covering Faster Payments and CHAPS and excluding international payments, loan repayments, a PSP's own funds and indirect bulk submissions. It returns a match indicator and a reason code, performs matching at the responding PSP, and never reveals a name on a mismatch. It catches the wrong address. It cannot catch the wrong intention, which is why the reimbursement regime of 7 October 2024 exists alongside it and why, in a system where over 99 per cent of Faster Payments initiators check names, £576.4 million still went to the wrong people in 2025.

The scaffolding firm's payroll clerk had every number right. That was always the problem.

39.98 Common wrong ideas

Wrong: The payee name is part of the address and is checked before delivery. **Right:** A UK domestic payment is routed by six digits of sorting code and eight of account number and by nothing else; the name is a label for the payer's own records.

Wrong: A passing modulus check means the account exists. **Right:** It tests the fourteen digits against themselves only, and a MOD10 rule admits about one in ten arbitrary numbers while a MOD11 rule admits about one in eleven; failing is strong evidence, passing is barely evidence at all.

Wrong: Every sorting code can be modulus checked. **Right:** Where a sorting code falls outside every range in the weighting table no check can be performed, and the published specification directs that the details be presumed valid.

Wrong: A sorting code identifies the institution holding the money, and once allocated it can be paid on any rail. **Right:** Sponsors sub-allocate codes from their own range to indirect participants, so the named institution and the payee's actual bank can differ, and a code must be separately registered with each payment system before it is usable there.

Wrong: The IBAN check digits prove the account is real. **Right:** MOD 97-10 detects all single-character substitutions and all adjacent transpositions and nothing else, saying nothing about whether the account exists or whether the bank code and sorting code inside the string even agree.

Wrong: A BIC identifies an account and proves you can send to it. **Right:** The BIC names an institution or a unit within one, never an account, and non-connected BICs are registered for reference only with no right to exchange messages over the network.

Wrong: British payments are addressed by IBAN. **Right:** Domestic Faster Payments, CHAPS and Bacs Direct Debits all carry a bare sorting code and account number, and a British consumer can bank for forty years without seeing an IBAN.

Wrong: Confirmation of Payee blocks a payment when the name does not match. **Right:** It is advisory; the payer sees the result and may carry on, and in a competently run scam the payer has already been told to expect a mismatch and given a reason for it.

Wrong: Confirmation of Payee covers every payment. **Right:** The direction excludes payments where either party is outside the United Kingdom, loan repayments, a PSP moving its own money and indirect bulk submissions, and the payee can opt out.

Wrong: Name matching is an exact comparison. **Right:** It is deliberately tolerant, which is why the close-match answer exists, because rejecting every payment to “J Smith” against an account reading “John Smith” would make the service unusable; each responder tunes that tolerance and the tuning is not published.

39.99 Chapter summary in 20 lines

1. A UK domestic payment is addressed by a six-digit sorting code and an eight-digit account number, and by nothing else.
2. The name typed into the payee box is a label for the payer’s records, because no domestic sterling rail routes on name.
3. Modulus checking is a self-consistency test built into those fourteen digits, specified by Vocalink and driven by the weighting table `valacdos.txt` and the substitution table `scsubtab.txt`.
4. Three algorithms are defined, MOD10, MOD11 and DBLAL, and DBLAL’s step of summing the digits of each product rather than the products themselves is where reimplementations go wrong.
5. Fourteen numbered exceptions modify the standard procedure, because institutions carried their own pre-existing check schemes into the industry standard rather than renumber their customers.
6. A passing check carries very little information, and where the sorting code is absent from the weighting table the specification directs that the details be presumed valid.
7. Modulus checking is therefore a typo filter of very high value at the point of capture and a fraud control of no value whatsoever.
8. Sorting codes are allocated by Pay.UK, sub-allocated by sponsors to indirect participants, and must be separately registered with each payment system, so allocation is not reachability.
9. The Extended Industry Sort Code Directory is the authority for both ownership and scheme registration, and should be refreshed weekly rather than guessed at from leading digits.
10. Account numbers must be normalised to eight digits before validation, and building society roll numbers cannot be modulus checked at all.
11. Very many real accounts are collection accounts, where the fourteen digits are the building and an unvalidated eighteen-character reference is the flat number.
12. The IBAN under ISO 13616 wraps the same domestic pair in a country code and two check digits, adding no information a domestic instruction does not already carry.
13. Those check digits are computed under ISO/IEC 7064 MOD 97-10, and validation rearranges the string, converts letters to digits and takes the remainder modulo 97, which is always exactly 1.
14. The resulting integer overflows a 64-bit type, so implementations without arbitrary-precision arithmetic must reduce it piecewise, nine digits and then seven at a time.
15. MOD 97-10 protects transcription and detects nothing about the account behind the string.
16. The BIC under ISO 9362 names an institution or a unit within one, never an account, and directory presence does not imply network connectivity.
17. Confirmation of Payee, launched in 2020, is a peer-to-peer API overlay with no central switch, matching at the responding PSP and never revealing a name on a mismatch.
18. Its reason codes carry more than a yes or no, and the four account-type codes flag a business payment landing on a personal account, a common shape for invoice redirection fraud.
19. Specific Direction 17 mandated the service in tranches to 31 October 2023 and 31 October 2024, excluding international payments, loan repayments, a PSP’s own funds and indirect bulk submissions.
20. Over 99 per cent of Faster Payments initiators now check names and APP losses still rose 19 per cent to £576.4 million in 2025, because the address was never the thing that was wrong.

Sources: Pay.UK, including the Confirmation of Payee service pages and FAQs, the Extended Industry Sort Code Directory pages, sort code information and allocation guidance, Faster Payment System transaction limits, and the ISO 20022 to Standard 18 translation guide; Vocalink, Validating Account Numbers: UK Modulus Checking, version 8.70 dated 10 February 2026; Swift, the IBAN Registry as ISO 13616 registration authority, and the BIC pages for ISO 9362; Payment Systems Regulator, Specific Direction 17 and the associated requirements pages, CP21/6 on Confirmation of Payee Phase 2, the Confirmation of Payee compliance report of 18 November 2025, and PS25/5, the consolidated policy statement on the APP scams reimbursement requirement; Bank of England, policy statement on mandating ISO 20022 enhanced data in CHAPS; UK Finance, Annual Fraud Report 2026 covering 2025; Bacs Standard 18 message implementation guidance.

SEPA

40.0 What this chapter gives you

1. You will be able to explain why there is no SEPA computer and no SEPA clearing house, and say what Europe built instead.
2. You will be able to distinguish scheme scope from legal scope, and explain why a euro transfer from Dublin to Zurich uses the same rulebook without the EU charge-equality rule.
3. You will be able to state the SCT Inst timing budget as five, seven and nine seconds, and say what the tenth second belongs to.
4. You will be able to explain why sanctions screening had to be taken out of the transaction path before instant payments could be made compulsory.
5. You will be able to say who holds a SEPA mandate, why the debtor's bank never sees the signature, and what the Unique Mandate Reference and Creditor Identifier are for.
6. You will be able to compare SEPA Direct Debit Core with UK Bacs Direct Debit and explain why the British model can promise an immediate refund with a straight face.
7. You will be able to say why the business-to-business scheme abolishes the refund right and what obligation it puts on the paying institution instead.
8. You will be able to name the five things that must all be true at once for a Cork to Lisbon payment to be domestic in character.
9. You will be able to explain why SEPA never had to migrate to ISO 20022, and what the 15 November 2026 address deadline requires of every message you send.
10. You will be able to list what SEPA did not unify: cards, currencies, the front end and national variation.

Europe did something in payments that almost nobody else has managed. It took thirty-odd countries with thirty-odd national payment systems, thirty-odd account numbering conventions, thirty-odd sets of file formats and thirty-odd sets of consumer protection rules, and it made a euro payment between any two of them behave, from the customer's side of the counter, like a payment inside one country.

It did not do this by building a single machine. There is no SEPA computer. There is no SEPA clearing house. What Europe built instead was a set of rulebooks, a legal obligation to follow them, and a shared message format, and then it let the existing infrastructure compete underneath. That distinction is the whole chapter. Understand it and the rest of SEPA falls into place. Miss it and you will spend years looking for a network that does not exist.

This chapter covers what SEPA unified and what it conspicuously did not, the four schemes that carry the traffic, the mandate model that makes European direct debits work differently from British ones,

the reachability rules that make a payment from Cork to Lisbon domestic in character, and the fact that SEPA was born speaking ISO 20022 while everyone else was still translating.

The plain version

Imagine that every country in Europe has its own postal service, and that each one uses a different shape of envelope.

In Germany the envelopes are tall and narrow. In France they are square. In Ireland the address goes on the back. Each country has its own rules about how long a letter takes, its own price for a stamp, and its own way of writing a house number. Posting a letter inside your own country is easy, because you learned the shape of your own envelope as a child. Posting a letter to the next country along is a small nightmare. You need a special form, you pay a special price, and it takes an unpredictable number of days.

Now imagine that all of those countries sat down together and agreed three things.

First, one envelope. Same size, same shape, same place for the address, everywhere.

Second, one way of writing an address, which always starts with the country. So instead of a house number that means something only to the local postman, everyone writes an address that any post office on the continent can read at a glance.

Third, and this is the one that changed people's lives, the same price. Posting a letter from Ireland to Portugal costs what posting a letter from Cork to Dublin costs. Not a bit more. The same.

That is SEPA. It stands for the Single Euro Payments Area, and it is an agreement about the shape of the envelope.

■ A payment, with real numbers

Aoife lives in Cork. She has found a second-hand guitar advertised by Nuno, who lives in Lisbon, for 480 euros.

In 1999 this would have been an adventure. She would have gone into a branch, filled in an international payment form, given the name and address of Nuno's bank, paid a fee somewhere between fifteen and thirty euros, and waited three to five days while wondering whether Nuno would receive 480 euros or 462 euros after somebody along the way took a slice.

Today she opens her banking app and types one thing: Nuno's account address. It looks like this.

PT50 0002 0123 1234 5678 9015 4

The first two letters are PT, for Portugal. That is the whole trick. Aoife's bank does not need to know anything about Portuguese banking to know where this payment is going, because the address tells it. Her own account address starts IE. Her German friend's starts DE. Everyone in the area writes their account address the same way, and it always begins with the country.

She types the amount, 480.00, and Nuno's name. Her app checks the name against the account before she confirms, and tells her it matches. She presses send.

The money arrives in Nuno's account in about four seconds. He gets a notification on his phone while Aoife is still looking at hers. It is a Sunday evening. No bank in either country is open.

The fee Aoife pays is whatever her bank charges her to send 480 euros to someone in Cork. If that is nothing, this is nothing. Nuno receives 480.00, not 479.20. Nobody in the middle takes a slice, because the rules say nobody may.

■ A regular payment, with real numbers

The second thing SEPA agreed a shape for is the opposite arrow: the payment you let somebody else pull.

Nuno subscribes to a language app run by a company in Dublin. It costs 12.99 euros a month. Instead of him remembering to pay every month, he signs a permission slip that says the Dublin company may take 12.99 euros from his Portuguese account on the fifteenth of each month.

That slip is the important object. In Europe, the company keeps it. Nuno's bank in Lisbon never sees his signature. What his bank sees, when the first collection turns up, is a claim from the Dublin company that says, in effect, we have a signed permission from this person, here is its reference number, here is our registered collector number, please pay.

Three rules protect Nuno.

The company must tell him in advance what it is going to take and when. The default is fourteen days' notice, though he and the company can agree something different.

The collection has to arrive at his bank at least one working day before the day it is due, so that his bank has time to look at it rather than having it dumped on the doorstep at the last minute.

And if he changes his mind, he can ring his bank within eight weeks of the money leaving and ask for it back, and his bank must give it back without asking him why. If he says he never signed anything at all, he has thirteen months.

That is a SEPA Direct Debit. It works identically whether the collector is in Dublin and the payer in Lisbon, or both are in the same street in Munich.

■ What SEPA is not

Two things are worth saying plainly at the dinner table, because they are the source of most confusion.

SEPA is euros only. If Aoife wanted to send Swedish kronor to Stockholm, none of this applies. The euro is the point. The area is defined by the currency, not by the European Union.

And SEPA is not a company. There is no SEPA. If you go looking for the building, there is no building. It is an agreement about the envelope, written down as a set of rulebooks, made compulsory by European law. The actual sorting offices, the machines that move the money, are separate businesses that compete with each other, and every one of them has agreed to use the same envelope.

Where the plain version stops being true

■ SEPA is a rulebook, not a rail

The envelope analogy hides the entire infrastructure layer, and the infrastructure layer is where a practitioner spends their life.

A SEPA payment does not travel over "SEPA". It travels over a clearing and settlement mechanism, of which there are many. In the euro area alone the European Central Bank counted thirty-five retail payment systems processing around 55.7 billion transactions in the first half of 2025. The pan-European ones are STEP2-T and RT1, both operated by EBA CLEARING, and TIPS, operated by the Eurosystem.

There are also large national systems, of which CORE in France is the biggest. Two of the three largest retail systems in the euro area by volume, STEP2-T and CORE, are SEPA systems. The largest, MCMS, is Mastercard's, which tells you something about what SEPA did not cover.

Underneath all of them sits central bank money. Deferred net obligations from STEP2-T settle in T2, the Eurosystem's real-time gross settlement service. Instant payments settle across dedicated accounts in TIPS. When you say "SEPA payment" you have named the rulebook and said nothing about which of these carried it, and the choice matters enormously for cut-offs, liquidity and failure modes.

■ **"Domestic in character" is a legal claim, and it has edges**

Aoife's payment to Nuno is domestic in character because European law makes it so, not because the physics changed. Three specific edges are hidden.

The charging rule is a law with a geographical limit. It applies within the European Union. But the geographical scope of the EPC schemes covers forty-one countries and territories, eleven of which are outside the European Economic Area, including the United Kingdom, Switzerland, Monaco, Andorra, San Marino and Serbia. A euro credit transfer from Dublin to Zurich uses the same rulebook and the same message, but the EU charge-equality rule and the instant payments law do not reach Switzerland. Same envelope, different legal weather.

Non-EEA legs also change the data you must send. Under the EU funds transfer rules, when either the sending or the receiving payment service provider sits in a SEPA country outside the European Economic Area, the payer's full address must be carried in the message. Within the EEA it is optional. So a payment to a British account is not, in fact, format-identical to a payment to a Belgian one.

And the promise that nobody may refuse your account because of the country it is in is a law precisely because it kept happening, and it still does. Payroll systems, insurers, utilities and subscription platforms that only accept a national IBAN are a live enforcement problem across the area, not a solved one.

■ **"Instant" is a defined term with a budget and a set of trapdoors**

Four seconds is what the customer sees. The scheme sees a timing budget with named deadlines, and the deadlines have consequences.

The rulebook does not say "fast". It says that the confirmation must have reached the originating institution by a specific second, and that if it has not, the payer's institution must put the money back. Money that "arrived instantly" may in fact have been reserved rather than debited, and may be unreserved a few seconds later.

There are also caps that the plain version never mentions. At launch in November 2017 the scheme had a hard ceiling of 15,000 euros per transaction. That rose to 100,000 euros in 2020, and the scheme-level ceiling has since been removed, but individual institutions still set their own value limits according to their own risk appetite, and European law explicitly lets each customer set a personal per-transaction or per-day limit. "No limit" at scheme level does not mean no limit in your app.

Finally, the reason instant payments could be made compulsory at all is a legal manoeuvre most people never notice. Sanctions screening was taken out of the transaction path. Institutions now screen their own customer base against targeted financial restrictive measures at least once every calendar day, and are forbidden from screening the payer and payee again during the execution of an instant transfer. Without that change, ten-second settlement across thirty countries would not have been achievable.

■ Direct debit is the least unified of the unified things

The plain version implies that a European direct debit is one thing. It is closer to two things wearing a shared uniform, sitting on top of thirty national habits.

There are two schemes, not one, and they behave very differently. The consumer scheme gives an eight-week no-questions-asked refund. The business-to-business scheme abolishes the refund entirely and instead requires the payer's bank to check the mandate before it debits, which is a completely different risk model wearing a similar-looking message.

Adoption is wildly uneven. In the first half of 2025, direct debits were around 14 per cent of non-cash payments across the euro area, but in Germany they were around 31 per cent of all non-cash payments. A German product manager and an Italian one are describing different worlds when they say "our customers pay by direct debit".

And the electronic mandate, the thing that ought to make all this frictionless, is still a minority. Of euro area direct debits in the first half of 2025, those with an electronic mandate accounted for 12 per cent by number. Eighty-eight per cent still rest on consent given in some other form, which in practice means paper, or a national scheme's own e-mandate service, or a recorded telephone call.

The technical version

■ What SEPA is, institutionally

SEPA is a European Union policy objective, delivered through two instruments that are frequently confused.

The first is scheme rules. These are written and maintained by the European Payments Council, an international non-profit association under Belgian law which is explicitly not part of the EU institutional framework. The EPC publishes rulebooks, each institution adheres to them individually and appears on a published register of scheme participants, and adherence is a contractual matter between the participant and the EPC.

The second is law. Regulation (EU) No 260/2012, universally called the SEPA Regulation, establishes technical and business requirements for credit transfers and direct debits in euro. It was amended in 2024 by Regulation (EU) 2024/886, the Instant Payments Regulation, which inserted new Articles 5a to 5d dealing with instant credit transfers, charges, verification of payee, and sanctions screening.

The geographical scope of the EPC schemes is set out in EPC409-09, which reached version 8.0 on 24 December 2025 and lists forty-one countries and territories. These comprise the twenty-seven EU Member States, three EEA countries (Iceland, Liechtenstein and Norway), eleven non-EEA countries (Albania, Andorra, Moldova, Monaco, Montenegro, North Macedonia, San Marino, Serbia, Switzerland, the United Kingdom and Vatican City State) and four non-EEA territories (Guernsey, Jersey, the Isle of Man and Saint-Pierre-et-Miquelon). All figures and dates in this chapter are accurate at time of writing.

The critical structural point is that scheme scope and legal scope are not the same set. Scheme rules bind a Serbian or British participant. EU regulations do not.

■ SEPA Credit Transfer

The SEPA Credit Transfer scheme, launched in January 2008, is the ordinary push payment. The governing document at time of writing is the 2025 SCT rulebook version 1.1 (EPC125-05), effective 5 October 2025.

Its core parameters:

Parameter	SEPA Credit Transfer
Currency	Euro at all process stages, including exception handling
Account identifier	IBAN, mandatory
Maximum execution time	One banking business day following the point in time of receipt of the instruction
Scheme-level amount ceiling	None; settlement and value limits may exist between participants
Remittance information	Up to 140 characters, structured or unstructured
Reject	Same day, at the latest the next banking business day
Return	Within three banking business days after settlement date
Recall	Ten business days for duplicate or technical error; thirteen months for fraud
Recall response	Within fifteen banking business days of receipt of the recall

Only one recall is permitted per transaction. The recall regime is worth dwelling on, because it is the structural answer to a question every treasury team asks: what happens when we send the wrong payment. The answer in SCT is that you may ask, the beneficiary institution must answer within a defined window, and it may say no. There is no unilateral clawback.

Volumes are large. The EPC puts SEPA credit transfer traffic at more than twenty-nine billion transactions a year across the area. The ECB's own euro area figures for the first half of 2025 record 16.8 billion credit transfers worth 107.3 trillion euros, which was 22 per cent of non-cash payments by number and 92 per cent by value.

■ SEPA Instant Credit Transfer

SCT Inst went live in November 2017, described by the EPC at the time as a world first for a region of this size. The governing document at time of writing is the 2025 SCT Inst rulebook version 1.1 (EPC004-16), published 5 October 2025 and stated to remain in force until 21 November 2027 at 03:30 CET.

The scheme is euro-only, operates twenty-four hours a day on every calendar day of the year subject to short and foreseeable planned maintenance, and has no scheme-level maximum amount. Value limits are applied by the originator's institution according to its own risk appetite and risk management controls.

The timing model is the part practitioners get wrong most often. It is not one number. The 2025 rulebook replaced the original 10 / 20 / 25 second structure with a 5 / 7 / 9 second structure, measured from the time stamp applied by the originator's institution at the time of receipt. Within seven seconds of that time stamp, the clearing mechanism serving the beneficiary's institution must have received a positive or negative confirmation message: this is the Timeout Deadline. That confirmation must then reach the originator's institution within a further two seconds, that is to say by the ninth second at the latest.

The tenth second belongs to the law rather than the scheme. Under Article 5a of Regulation (EU) No 260/2012 as inserted by the Instant Payments Regulation, the payee's institution must make the

amount available on the payee's account and confirm completion within ten seconds of the time of receipt of the payment order by the payer's institution; and where the payer's institution has received no such confirmation within ten seconds, it must immediately restore the payer's account to the state it would have been in had the transaction not taken place. The 2025 rulebook also moved the time stamp to millisecond precision.

Two further legal points matter. Time of receipt for an instant credit transfer is the moment the order is received by the payer's institution, regardless of hour or calendar day, notwithstanding the ordinary cut-off provisions of PSD2. And credit value date on the payee's account must be the same date as the date of crediting.

■ The Instant Payments Regulation timetable

The regulation converts SCT Inst from a commercial option into an obligation, on a staggered timetable.

Obligation	Euro area PSPs	Non-euro area PSPs
Receive instant credit transfers	9 January 2025	9 January 2027
Send instant credit transfers	9 October 2025	9 July 2027
Charge parity with non-instant transfers (Article 5b)	9 January 2025	9 January 2027
Verification of payee (Article 5c)	9 October 2025	9 July 2027
Daily sanctions screening (Article 5d)	9 January 2025	9 January 2025

Payment and electronic money institutions have their own dates, 9 April 2027 in the euro area and 9 July 2027 outside it, reflecting their later admission to settlement systems. Non-euro area institutions have a further deadline of 9 June 2028 relating to sending outside business hours in national currency.

The substantive obligations are worth stating precisely, because they are what a compliance officer will be measured against.

Article 5a(1) requires that any institution offering its customers a service of sending and receiving credit transfers must offer sending and receiving of instant credit transfers to all of its customers, and must ensure that all payment accounts reachable for credit transfers are also reachable for instant credit transfers twenty-four hours a day on any calendar day.

Article 5a(6) requires institutions, on request, to let a customer set a maximum amount per day or per transaction for instant credit transfers, modifiable at any time before an order is placed.

Article 5b(1) provides that charges for sending and receiving instant credit transfers shall not be higher than charges for other credit transfers of corresponding type, and Article 5b(2) requires that the verification service be provided free of charge.

Article 5d(1) requires institutions offering instant credit transfers to verify whether any of their customers are persons or entities subject to targeted financial restrictive measures, immediately after any new or amended measures enter into force and at least once every calendar day. Article 5d(2) then forbids them from screening payer and payee again during execution of an instant credit transfer.

There is one derogation of note. An institution in a Member State whose currency is not the euro may, with prior permission from its competent authority based on an assessment of its access to euro liquidity, decline to send instant euro transfers above a per-transaction limit from accounts denominated in

the national currency, during hours when it neither sends nor receives non-instant euro credit transfers. That limit is set by the competent authority and may not be lower than 25,000 euros. Permissions run for one year and may be extended.

Uptake is now material rather than marginal. In the first half of 2025, instant credit transfers accounted for 23 per cent of the number and 7 per cent of the value of credit transfer transactions processed by euro area retail payment systems.

■ Verification of Payee

The EPC's Verification Of Payee scheme (EPC218-23) is the scheme-level answer to Article 5c. Version 1.0 was published in October 2024; version 1.1 was issued on 16 March 2026 with an effective date of 20 September 2026.

VOP is not an ISO 20022 message flow in the way the payment schemes are. It uses application programming interface technology making use of ISO 20022 resource elements, exchanged through Routing and Verification Mechanisms of each participant's choosing, with participants registered in the EPC Directory Service.

There are four possible outcomes: match; no match; close match, in which case the name associated with the account is shown to the payer; and verification check not possible. The rulebook sets a maximum execution time of five seconds for the requesting institution to obtain the VOP response, and states a preference for one second or less.

Two things about scope catch people out. VOP applies to all credit transfers under Article 5c, not only instant ones. And it must be performed immediately after the payer provides the payee's details and before the payer is offered the possibility of authorising the transfer, regardless of the initiation channel.

■ SEPA Direct Debit Core

SDD Core launched in November 2009. The governing document at time of writing is the 2025 SDD Core rulebook version 1.1 (EPC016-06), effective 5 October 2025.

The defining architectural choice is the creditor-driven mandate flow. The debtor signs a mandate, on paper or electronically, in favour of the creditor. The creditor stores it. The debtor's institution does not receive the signed instrument. Each mandate is uniquely identified by the combination of the Unique Mandate Reference (attribute AT-M001) chosen by the creditor and the Creditor Identifier (attribute AT-E005) issued to the creditor under national arrangements. That pair, and not any document held at the paying institution, is what makes a collection traceable.

The operational parameters:

Parameter	SDD Core
Currency	Euro at inter-PSP level
Pre-notification to debtor	At least 14 calendar days before due date, unless another timeline is agreed between debtor and creditor
Presentation	Debtor's PSP must receive the collection at the latest one inter-PSP business day before due date (D-1)
Returns	Within five inter-PSP business days following settlement date

Parameter	SDD Core
Refusal	Debtor may instruct its PSP before settlement; handled as a rejection
Refund, authorised collection	Eight weeks from the debit date, on a no-questions-asked basis
Refund, unauthorised collection	Claim to be presented within thirteen months of the debit date
Mandate dormancy	If no collection is presented for 36 months, the creditor must cancel the mandate

A point of law that is easy to skip: offering SDD Core is not optional. Under the reachability provisions, an institution reachable for national direct debits in euro must be reachable for consumer direct debits under a Union-wide scheme, which in practice means SDD Core.

■ SEPA Direct Debit B2B

The business-to-business scheme launched alongside Core in November 2009 and is optional for institutions. The governing document at time of writing is the 2025 SDD B2B rulebook version 1.0 (EPC222-07), issued 28 November 2024 and effective 5 October 2025.

Three differences define it.

The debtor may not be a consumer. Participation is limited to debtors authorised under national law to opt out of the refund right.

There is no refund right for an authorised transaction. The debtor cannot ask its institution for the money back.

Because the refund right is gone, the risk control moves upstream. Before debiting the account, the debtor's institution must check that the mandate-related data received in the collection matches mandate-related data received from, or confirmed by, the debtor. This obligation is the reason B2B onboarding is slow: somebody at the paying bank has to be told about the mandate before the first collection, not after it.

Presentation is the same D-1 as Core. Returns are permitted up to three inter-PSP business days after settlement date, a shorter window than Core's five, and the mandate dormancy rule at 36 months is the same.

■ The mandate model compared with UK Direct Debit

For a British reader, and for anyone building a product that has to do both, this is the comparison that matters. The two systems reach similar outcomes by opposite routes.

The clearest way to see the difference is to ask a single question: after the mandate is signed, who has a copy of it?

In the United Kingdom, under AUDDIS, the signed Direct Debit Instruction is retained by the service user, but the instruction itself is lodged electronically with the paying institution through Bacs. Bacs guidance is explicit that the one major change AUDDIS made was that the signed DDI is retained by the service user and not by the paying institution, while the details still travel to the paying institution. The paying bank therefore holds a record of the instruction, can validate collections against it, and can bounce a collection for which no instruction exists.

In SEPA Core, no equivalent lodgement takes place. The mandate data reaches the debtor's institution with the first collection, and the debtor's institution has nothing to check it against unless it has built its own mandate store. The B2B scheme exists precisely to plug that gap for corporate payers.

	SEPA Direct Debit Core	UK Bacs Direct Debit
Who holds the signed mandate	Creditor	Service user
Does the payer's bank hold the instruction	No, receives mandate data with the first collection	Yes, lodged via AUDDIS
Mandate identifiers	Unique Mandate Reference plus Creditor Identifier	Service User Number plus payer reference
Advance notice, default	14 calendar days before due date, unless otherwise agreed	Minimum 10 working days plus postal time, in the absence of any other agreement
Collection cycle	Collection received by debtor's PSP at latest D-1	Three working day cycle: submission, processing, entry
Settlement	Deferred net at the clearing mechanism	Deferred net through Bacs
Refund, authorised	8 weeks, no questions asked	Direct Debit Guarantee: immediate refund by the paying institution, then indemnity claim against the service user
Refund, unauthorised	Claim within 13 months	Direct Debit Guarantee, with the service user liable on indemnity
Producing the mandate on request	Creditor must be able to produce it	Copy required within seven working days of request
Dormancy	Creditor must cancel after 36 months without a collection	Paying institutions hold DDI details for a minimum of 24 months from receipt or last payment
Currency and geography	Euro, 41 SEPA countries and territories	Sterling, United Kingdom

Three practical consequences follow.

The British model puts the arbiter at the payer's bank, which is why the Direct Debit Guarantee can promise an immediate refund with a straight face: the paying institution refunds first and recovers from the service user afterwards through the indemnity process. The SEPA model puts the arbiter at the creditor, and compensates by giving the payer a broad, time-boxed, no-questions-asked right that the paying institution honours from the settlement chain.

The British model is therefore harder to join and easier to police. Service users are sponsored by a payment service provider, hold a Service User Number, and are vetted. Note also that from 1 July 2027 it will no longer be possible to apply for a non-AUDDIS Service User Number, which completes the migration to electronic lodgement.

And the notice periods are not equivalent even though they look similar. Ten working days plus postal time in the United Kingdom is a longer real period than fourteen calendar days in SEPA, and the UK figure is a default that can be varied only with the sponsor's agreement, while the SEPA figure can be varied by agreement between creditor and debtor directly.

■ Reachability obligations

Reachability is the provision that turns a scheme into an area. Without it, SEPA would be a format standard that some banks supported for some corridors.

Article 3 of Regulation (EU) No 260/2012 states the principle in two directions. A payee's payment service provider that is reachable for a national credit transfer under a payment scheme must be reachable for credit transfers initiated through a scheme at Union level. A payer's payment service provider that is reachable for a national direct debit must be reachable for direct debits initiated through a scheme at Union level.

The effect is that reachability cannot be sold as a premium corridor. If an institution accepts euro credit transfers from across the street, it must accept them from across the continent.

The Instant Payments Regulation extends this to time as well as geography. Article 5a(1) requires that all payment accounts reachable for credit transfers must also be reachable for instant credit transfers twenty-four hours a day on any calendar day.

Underneath the legal obligation sits an interoperability problem, because reachability is only real if the clearing mechanisms interconnect. This is why RT1 is designed to reach all institutions adhering to the SCT Inst scheme, whether directly, through other clearing mechanisms connected to RT1, or through Eurosystem reachability measures, and why participants can use a single interface for transactions settling in RT1, in TIPS, or across the two. Reachability in law required reachability in plumbing, and the plumbing had to be built.

■ Why a euro payment from Ireland to Portugal is domestic in character

Five separate things have to be true at once for the Cork to Lisbon payment to feel like a Cork to Dublin payment, and SEPA supplies all five.

The identifier is the same. IBAN is mandatory in both cases. There is no separate international addressing scheme to learn, and the SEPA Regulation removed the requirement for customers to supply a BIC.

The message is the same. The same ISO 20022 messages, in the same versions, with the same mandatory fields, carry a domestic and a cross-border euro transfer.

The rulebook is the same. Execution time, return windows, recall procedure, remittance capacity and exception handling do not change at the border.

The reachability obligation removes the question of whether the far end can be paid at all, and Article 9 of the SEPA Regulation removes the question of whether the far end will accept the account: a payer must not be required to specify the Member State in which a payee's account is located, and a payee must not be required to specify the Member State of a payer's account, provided the account is reachable. This is the provision that makes refusing a foreign IBAN unlawful.

And the price is the same. Regulation (EC) No 924/2009, as amended by Regulation (EU) 2019/518, requires that charges for a cross-border payment in euro be the same as those for a corresponding national payment of the same value in the same currency. The Instant Payments Regulation then adds the second parity rule inside the instant product: instant may not cost more than non-instant of corresponding type.

Domestic in character is therefore a precise claim about identifiers, formats, rules, access and price. It is not a claim about routing, and it does not survive being carried outside the European Union, where the pricing and instant-payment obligations no longer apply even though the scheme rules still do.

■ ISO 20022 as SEPA's native format

Most payment systems in the world are migrating to ISO 20022. SEPA never had to, because it started there.

The first SEPA Credit Transfer rulebook in 2008 specified ISO 20022 XML, at a point when the standard was barely four years old and almost nothing else used it. Two years before the euro area migration deadline, the SEPA Regulation made it law: the Annex to Regulation (EU) No 260/2012 requires that the message format be the ISO 20022 XML standard, that the account identifier be the IBAN, that at least 140 characters of remittance information be carried, and that processing be fully automated and electronic without manual intervention.

This has three consequences that a practitioner feels every day.

There is no legacy format to translate from. Elsewhere, ISO 20022 adoption is a truncation problem: fields that existed in a fixed-width or tag-based legacy format have to be squeezed into or expanded out of the new one, and data is lost at the seams. SEPA has no such seam. The customer-to-bank message, the interbank message and the reporting message were designed as one family.

The versions are pinned. Scheme messages are not “ISO 20022” in the abstract, they are named versions specified in the implementation guidelines. At time of writing the SEPA schemes use the 2019 message versions.

Scheme	Message	Purpose
SCT and SCT Inst	pacs.008.001.08	FI to FI customer credit transfer, the payment itself
SCT and SCT Inst	pacs.002.001.10	FI to FI payment status report, positive and negative confirmation
SCT and SCT Inst	pacs.004.001.09	Payment return, and positive response to a recall
SCT and SCT Inst	camt.056.001.08	FI to FI payment cancellation request, used for recall
SCT and SCT Inst	camt.029.001.09	Resolution of investigation, negative response to a recall
SCT Inst	pacs.028.001.03	FI to FI payment status request
SDD Core and B2B	pacs.003.001.08	FI to FI customer direct debit, the collection
SDD Core and B2B	pacs.002.001.10	FI to FI payment status report, used for rejects
SDD Core and B2B	pacs.004.001.09	Payment return, covering returns and refunds
SDD Core and B2B	pacs.007.001.09	Payment reversal, creditor-initiated

And the format still moves, which is the answer to anyone who thinks a native standard means a frozen one. The current migration is addresses. Until 5 October 2025 only unstructured and structured address formats were permitted. From 5 October 2025 to 15 November 2026 all three formats are allowed, including a hybrid combining structured elements with address lines. From 15 November 2026 only structured and hybrid addresses will be allowed. The 2025 SCT Inst rulebook version 1.1 exists in large part to move that date from 22 November 2026 to 15 November 2026.

Address rules interact with the geography in a way already flagged. Where either institution is located in a SEPA country outside the European Economic Area, the payer's full address must be provided in the message under the EU funds transfer rules. Within the EEA it is optional. Screening requirements in practice only need the structured town and country elements, which is part of the argument for the structured migration.

■ The infrastructure underneath

Three systems carry most of the pan-European traffic.

STEP2-T is EBA CLEARING's pan-European automated clearing house for retail payments in euro, and one of the three largest retail payment systems in the euro area. It settles by continuous gross settlement: payments settle at any time the settlement windows are open, transaction by transaction, in immediately available central bank funds held on a TARGET Technical Account, with queuing and optimisation where liquidity is short.

RT1 is EBA CLEARING's instant service, live since 2017, providing real-time gross settlement of instant payments in central bank funds and reaching all institutions adhering to the SCT Inst scheme, directly or through connected mechanisms.

TIPS is the Eurosystem's TARGET Instant Payment Settlement service. It settles in central bank money on a 24/7/365 basis, processes transactions in under 100 milliseconds, and provides final and irrevocable settlement. It is compliant with the SCT Inst scheme. It currently handles euro, Swedish krona and Danish krone. It operates on cost recovery with no entry fee and a fixed charge of 0.2 euro cent per transaction, split equally between the sending and receiving participants.

Above them sits T2, the Eurosystem's real-time gross settlement service for euro payments in central bank money, which since March 2023 has replaced the former TARGET2 within the consolidated TARGET Services platform, and alongside it EURO1, EBA CLEARING's multilateral net large-value system whose final balances settle end of day in T2. In the first half of 2025, euro area large-value systems settled 74.0 million payments worth 235.1 trillion euros.

■ What SEPA did not unify

The honest summary of SEPA's scope is that it unified two instruments completely, one instrument partially, and left everything else alone.

It did not unify cards. The SEPA Cards Framework did not produce a pan-European card scheme, and the largest retail payment system in the euro area by volume is Mastercard's clearing system. Card payments were 57 per cent of euro area non-cash payments by number in the first half of 2025, and none of that traffic is governed by a SEPA scheme.

It did not unify currencies. Everything in this chapter is euro-denominated at the interbank level. Non-euro EU members participate as scheme countries and their institutions can hold euro accounts, but a domestic payment in Polish złoty or Swedish krona is outside SEPA entirely.

It did not unify the front end. Request-to-pay, e-invoicing, mobile initiation, wallets and account-to-account overlays sit above the schemes, not inside them. The EPC publishes a SEPA Request-To-Pay rulebook, published in November 2020, which is a messaging layer for asking, not a payment scheme. Commercial overlay services build on SCT Inst rather than replacing it.

It did not eliminate national variation. Additional Optional Services remain, national mandate services differ, and adoption patterns diverge so sharply that the same scheme supports 31 per cent of non-cash payments in one country and a rounding error in another.

And it did not, by itself, make anything instant. That took twelve more years, a second scheme, and finally a regulation with dates in it.

■ Timeline

Date	Milestone
January 2008	SEPA Credit Transfer scheme launched
November 2009	SEPA Direct Debit Core and B2B schemes launched
February 2012	Regulation (EU) No 260/2012, the SEPA Regulation, adopted
1 August 2014	End of migration for euro area countries
31 October 2016	End of migration for non-euro area countries
November 2017	SCT Inst launched, initial ceiling 15,000 euros
2020	SCT Inst ceiling raised to 100,000 euros
November 2020	SEPA Request-To-Pay rulebook published
November 2018	TIPS launched by the Eurosystem
March 2023	One-Leg Out Instant Credit Transfer rulebook published; T2 replaces TARGET2
13 March 2024	Regulation (EU) 2024/886, the Instant Payments Regulation, adopted
October 2024	First Verification Of Payee scheme rulebook published
9 January 2025	Euro area PSPs must receive instant credit transfers; charge parity; daily sanctions screening
5 October 2025	2025 rulebooks enter force for SCT, SCT Inst, SDD Core and SDD B2B
9 October 2025	Euro area PSPs must send instant credit transfers and offer verification of payee
15 November 2026	Only structured and hybrid address formats permitted
9 January 2027 onwards	Staggered obligations for non-euro area PSPs and for payment and e-money institutions

■ What to check before you build

If you are integrating, five questions decide most of your design.

Which scheme, and therefore which timing model: D+1 for SCT, seconds for SCT Inst, D-1 presentation for both direct debit schemes.

Which clearing mechanism your institution uses, because that determines your real cut-offs, your file formats at the edges, and your liquidity behaviour, none of which the rulebook tells you.

Whether any leg touches a SEPA country outside the European Economic Area, which changes your address data obligations and removes your customer's EU pricing and instant-payment protections.

Whether your direct debit population contains consumers, which decides Core versus B2B and therefore whether you are exposed to an eight-week no-questions-asked refund or to a mandate-checking obligation at the paying institution.

And whether your address data is structured. That deadline is 15 November 2026 and it applies to every message you send.

40.98 Common wrong ideas

Wrong: SEPA is a payment system that carries the money. **Right:** SEPA is a set of rulebooks made compulsory by law, and the traffic runs over competing clearing and settlement mechanisms such as STEP2-T, RT1, TIPS and national systems like CORE.

Wrong: SEPA covers the European Union. **Right:** The area is defined by the currency rather than the Union, and the EPC schemes cover forty-one countries and territories, eleven of them outside the EEA, including the United Kingdom and Switzerland.

Wrong: A euro payment anywhere in SEPA costs what a domestic one costs. **Right:** The charge-equality rule is EU law with a geographical limit, and it does not reach Switzerland or the United Kingdom even though the scheme rules do.

Wrong: A SEPA message is format-identical wherever it is sent. **Right:** Where either institution sits in a SEPA country outside the European Economic Area, the EU funds transfer rules require the payer's full address, which is optional within the EEA.

Wrong: SCT Inst has no value limit now that the scheme ceiling has gone. **Right:** Institutions still apply their own limits according to their risk appetite, and Article 5a(6) requires them to let a customer set a personal per-transaction or per-day maximum.

Wrong: Money that arrived instantly has definitely arrived. **Right:** It may have been reserved rather than debited, and where no confirmation reaches the payer's institution within ten seconds it must restore the payer's account to the state it would have been in.

Wrong: A European direct debit is one thing. **Right:** There are two schemes with opposite risk models: Core with an eight-week no-questions-asked refund, and B2B with no refund right for an authorised collection and a mandate check at the paying institution before it debits.

Wrong: The debtor's bank holds the mandate and can validate collections against it, as in the United Kingdom. **Right:** In SEPA the creditor holds it, the mandate data reaches the debtor's institution with the first collection, and that institution has nothing to check against unless it has built its own mandate store.

Wrong: European direct debits now run on electronic mandates. **Right:** In the first half of 2025 only 12 per cent of euro area direct debits by number carried an electronic mandate, with the rest resting on paper, national e-mandate services or recorded calls.

Wrong: Verification of Payee applies to instant credit transfers. **Right:** Article 5c applies it to all credit transfers, and it must be performed immediately after the payee's details are entered and before the payer is offered the possibility of authorising.

40.99 Chapter summary in 20 lines

1. SEPA made a euro payment between any two of thirty-odd countries behave, from the customer's side of the counter, like a payment inside one country.
2. It did this without building a machine, because what Europe built was a set of rulebooks, a legal obligation to follow them and a shared message format.
3. The rules are written by the European Payments Council, a Belgian non-profit explicitly outside the EU institutional framework, and adherence is contractual.
4. The law is Regulation (EU) No 260/2012, amended in 2024 by the Instant Payments Regulation, which inserted Articles 5a to 5d.

5. Scheme scope and legal scope are different sets, since the schemes cover forty-one countries and territories, eleven outside the EEA, where EU regulations do not bind.
6. SEPA Credit Transfer, launched in January 2008, executes by the following banking business day, uses IBAN, carries 140 characters of remittance information, and permits a request to recall but no unilateral clawback.
7. SEPA Instant Credit Transfer went live in November 2017 and now runs a five, seven and nine second scheme budget inside a ten-second legal obligation.
8. If no confirmation reaches the payer's institution within ten seconds, that institution must immediately restore the payer's account as though the transaction had never happened.
9. The scheme-level ceiling, once 15,000 euros and then 100,000, has been removed, but institutional limits and customer-set limits remain.
10. Instant payments could only be made compulsory because sanctions screening was moved out of the transaction path into a daily check by each institution against its own customer base.
11. Verification of Payee under Article 5c applies to all credit transfers, returns four possible outcomes, and must run before the payer is offered the chance to authorise.
12. SEPA Direct Debit Core, launched in November 2009, is creditor-driven, with the creditor holding the signed mandate and identifying it by Unique Mandate Reference plus Creditor Identifier.
13. Core gives the debtor fourteen days' default pre-notification, D-1 presentation, an eight-week no-questions-asked refund, and thirteen months for an unauthorised collection.
14. The B2B scheme abolishes the refund right for authorised collections and moves the control upstream, requiring the debtor's institution to check mandate data before it debits.
15. The British scheme reaches similar outcomes by the opposite route, because AUDDIS lodges the instruction with the paying bank, which is why the Direct Debit Guarantee can refund immediately and recover afterwards by indemnity.
16. Reachability under Article 3 is what turns a scheme into an area, since an institution reachable for a national euro credit transfer or direct debit must be reachable for the Union-wide scheme, now twenty-four hours a day.
17. Five things make Cork to Lisbon domestic in character: the same identifier, the same message, the same rulebook, the reachability and non-discrimination obligations, and the same price.
18. SEPA was born speaking ISO 20022, so there is no legacy seam to translate across, but the versions are pinned and the format still moves, with only structured and hybrid addresses permitted from 15 November 2026.
19. SEPA did not unify cards, currencies, the front end or national habit, and card payments were 57 per cent of euro area non-cash payments by number in the first half of 2025 with none of it governed by a SEPA scheme.
20. Before you build, decide which scheme and timing model, which clearing mechanism, whether any leg leaves the European Economic Area, whether your debtors are consumers, and whether your address data is structured.

Sources: European Payments Council rulebooks and guidance, namely EPC125-05 2025 SCT rulebook v1.1, EPC004-16 2025 SCT Inst rulebook v1.1, EPC016-06 2025 SDD Core rulebook v1.1, EPC222-07 2025 SDD B2B rulebook v1.0, EPC218-23 Verification Of Payee rulebook v1.1, EPC114-06 and EPC122-16 implementation guidelines, EPC409-09 List of SEPA Scheme Countries v8.0, EPC153-22 v2.1 guidance on the provision of addresses, and the EPC SEPA timeline; Regulation (EU) No 260/2012 and Regulation (EU) 2024/886 as published in the Official Journal; Regulation (EC) No 924/2009 as amended by Regulation (EU) 2019/518; European Central Bank, TARGET Instant Payment Settlement service pages and "Payments statistics: first half of 2025" (29 January 2026); EBA CLEARING STEP2-T settlement and RT1 SCT Inst service pages; Bacs, AUDDIS and Direct Debit scheme FAQs and the Service User's Guide and Rules to the Direct Debit Scheme.

SWIFT and Correspondent Banking

41.0 What this chapter gives you

1. You will be able to explain why SWIFT has never moved a penny, and say precisely what it does move instead.
2. You will be able to draw the correspondent chain for a dollar payment from Sheffield to Izmir and identify the point at which a dollar left the United States, which is never.
3. You will be able to account for eighty-five missing dollars on a USD 214,000 payment by naming the charge basis that entitled each bank to take its cut.
4. You will be able to separate the four causes of a three-day international payment, cut-off times, business calendars, value dating and compliance screening, and say which of them a faster network would fix.
5. You will be able to say which side of which balance sheet a nostro sits on and which side a vostro sits on, and why one account carries two names.
6. You will be able to map an MT 103 onto a pacs.008 and explain why translating back the other way loses data that cannot be recovered.
7. You will be able to say what changed on 22 November 2025, and what it now costs an institution that has not migrated.
8. You will be able to explain why SWIFT's "nearly 60 per cent within 30 minutes" and the Financial Stability Board's "50.6 per cent within one hour" are both honest figures.
9. You will be able to tell a customer what a UETR is, where it lives in the message, and what the gpi tracker will and will not tell them.
10. You will be able to argue that a correctly structured beneficiary address is a speed feature rather than a compliance chore, and show the mechanism by which it is one.

On a Tuesday morning a finance manager at an engineering firm in Sheffield instructs a payment of USD 214,000 to a steel mill outside Izmir. She has the mill's account number, the mill's bank, and a BIC. She types it into her bank's business portal at 09:40, approves it on a second device, and gets a reference number. The screen says the payment has been sent.

On Friday afternoon the mill emails to confirm receipt. USD 213,915 has arrived. Eighty-five dollars are missing, nobody in Sheffield authorised anyone to take them, and three working days have gone by. Her bank's helpline tells her the payment "left us on Tuesday", which is true, and that "it's with the correspondent", which is also true, and neither statement helps.

Nothing went wrong. That is the disorienting part. Every institution in the chain did precisely what it was designed to do, the money arrived, and the arithmetic is defensible line by line. An international payment is not a thing that travels. It is a sequence of separate domestic bookkeeping entries in differ-

ent countries, stitched together by messages, and this chapter is about the messages, the bookkeeping, and the gap between them.

All figures, dates and message versions here are accurate at the time of writing, August 2026. The standards in particular have just been through the largest change in their history.

The plain version

Start with the single most common misunderstanding in all of payments. Ask most people what SWIFT is and they will say it is the system that sends money abroad. It is not. SWIFT sends no money at all. It has never sent a penny, a cent or a lira, and it never will.

SWIFT sends messages.

■ Two shops and two tins

Imagine a corner shop in Sheffield and a corner shop in Istanbul. The two owners have known each other for years. Because their customers keep asking to send money to relatives in the other city, they set up an arrangement.

The Sheffield owner keeps a tin of Turkish lira at the Istanbul shop. The Istanbul owner keeps a tin of pounds at the Sheffield shop. Both tins are labelled, both are counted, and both live permanently in the shop where they sit. They do not travel.

Now a customer walks into the Sheffield shop and wants to send money to her brother in Istanbul. She hands over cash. The Sheffield owner picks up the phone and says: pay my customer's brother out of my tin. The Istanbul owner opens the Sheffield tin, counts out the lira, hands them to the brother, and writes down that the tin is now that much lighter.

Money has moved from one person to another across two thousand miles. Nothing physical crossed a border. The cash the customer handed over is still in Sheffield. The lira the brother received were already in Istanbul. All that crossed the border was a phone call.

That is correspondent banking. The tins are called **accounts**, the shops are called **banks**, and the phone call is called **SWIFT**.

■ Why SWIFT had to be invented

The phone call is the weak point. It can be misheard, it can be faked by someone who sounds convincing, and it cannot be proved afterwards. When there are not two shops but eleven thousand, in two hundred and twenty countries, all using different words for the same thing, "phone the other bank" collapses.

So in 1973 a group of banks built a shared, secure, standardised messaging network. They called it the Society for Worldwide Interbank Financial Telecommunication, which nobody says out loud, and it went live carrying real messages in 1977. It is a cooperative, owned by the banks that use it, based in Belgium.

It gives its members three things. A private network, so messages do not travel across the open internet. Cryptographic proof of who sent each message, so a fake instruction is very hard to inject. And, least glamorous but most important, a rigid vocabulary: a payment instruction has numbered fields, field 32A is always the value date, currency and amount, field 59 is always the beneficiary, and a bank in Japan and a bank in Chile fill in the same numbered boxes.

What SWIFT does not give its members is money. It has no view of anyone's balance, it holds no accounts, and it does not know whether the payment it carried actually happened. It delivers the instruction. The banks do the rest.

■ Why there is usually a chain

The two-shop story works because those two shopkeepers know each other and keep tins with each other. Real banks cannot do that with every other bank on earth. There are tens of thousands of banks. Maintaining an account with each one would mean tens of thousands of pots of idle money, tens of thousands of reconciliations, and tens of thousands of relationships to check for fraud and criminality.

So banks keep accounts with a small number of large banks, and those large banks keep accounts with each other. If you want to send dollars, the chain has to pass through a bank that has actual dollars in America, because dollars only really exist in America. If you want to send pounds, the chain has to pass through a bank that has actual pounds in Britain.

Here is the Sheffield firm's payment, drawn as a chain.

Step	Who	What they actually do
1	The firm's UK bank	Takes USD 214,000 of value from the firm. It has no dollars in America itself, but it keeps a dollar tin at a big bank in New York.
2	New York bank A	Opens the UK bank's dollar tin, takes out USD 214,000, and passes it to New York bank B. Both tins are in New York. Nothing crosses an ocean.
3	New York bank B	Puts the dollars into the Turkish bank's dollar tin, which it holds. Still in New York.
4	The Turkish bank	Sees that its New York tin is heavier, and credits the steel mill's account in Izmir with the matching amount.

Four institutions, three messages, and at no point did a dollar leave the United States. What travelled was a claim: at the end, the Turkish bank has more dollars in New York than it had before, and the UK bank has fewer.

■ Where the eighty-five dollars went, and the three days

Each of those middle banks did work. Each one has a price list. Because the payment was sent on shared-charge terms, which is the default for most business payments, the sending firm paid its own bank's fee up front and every bank downstream was entitled to take its cut out of the money itself.

Who	Deducted	Running total arriving
Sent		214,000
New York bank A	30	213,970
New York bank B	25	213,945
Turkish bank, receiving fee	30	213,915

Eighty-five dollars, all of it disclosed in somebody's tariff, none of it visible on the Sheffield screen at the moment of sending.

The days are a different story and a more interesting one. The messages themselves were delivered in seconds. What consumed Tuesday, Wednesday and Thursday was everything around the messages: a bank's daily cut-off time for dollar payments, after which your instruction is tomorrow's work; the fact that New York is five hours behind Sheffield and its dollar clearing system has its own opening hours; and, in this case, a compliance check. The steel mill's name resembled a name on a sanctions list closely enough for an automated filter to stop the payment and put it in a queue for a human being. The human being asked for the mill's full street address, the question travelled back up the chain to Sheffield, the answer travelled back down, and by the time the payment was released it had missed the next cut-off.

■ The parcel tracking number

For decades this was invisible. Once a payment left, the sending bank genuinely did not know where it was. It could ask, by sending another message, and wait.

Then in 2017 the banks agreed to attach a tracking number to every payment: a long, random, unique string generated when the payment is created and carried unchanged all the way down the chain. Every bank that touches the payment reports back to a central tracker: I received it at this time, I passed it on at this time, I took this much in fees, I used this exchange rate. The sending bank can now watch the payment move, the way you watch a parcel. That programme is called SWIFT gpi. It did not make payments faster by itself. It made them visible, which turns out to be most of what people wanted.

Where the plain version stops being true

■ The tin is not a tin

The most important simplification above is the pot of money sitting in the other shop. There is no pot. When a UK bank "holds dollars in New York", what it actually holds is a claim on the New York bank: a promise, recorded as a number on the New York bank's ledger, that the New York bank owes it that many dollars. It is a deposit, which means it is an unsecured loan to that institution. If the New York bank fails, the UK bank does not walk in and collect its dollars. It joins the queue of creditors.

This changes the shape of the whole business. It is why banks care intensely about which correspondents they use, why a correspondent relationship takes months of due diligence, and why a bank will keep balances with two or three correspondents in the same currency rather than one. It is also why the balance is not a passive store. It is a working balance, run deliberately thin, topped up during the day, drawn down by outgoing payments, and frequently supported by an intraday credit line from the correspondent. "Trapped liquidity" in nostro accounts means exactly this: money the bank cannot use for anything else, parked in other people's institutions in other people's jurisdictions, purely to be able to pay.

And because those balances sit inside another country's legal system, they are subject to that country's law. Dollar balances in New York are within reach of United States authorities, whoever owns them. That is the mechanism by which one country's sanctions policy becomes binding on banks with no other connection to it, and it is the single largest reason correspondent banking has become as cautious as it has.

■ The messages do not always follow the money

The chain drawn above implies one flow: the instruction walks down the chain alongside the value. That is one of the two ways it is done, and often not the one used.

In the **serial method** the payment instruction really does pass from bank to bank down the chain, each one debiting and crediting the accounts it holds and then forwarding a fresh instruction to the next.

In the **cover method** there are two separate flows. The originating bank sends the customer payment instruction directly to the beneficiary's bank, so the beneficiary's bank knows immediately that money is coming and for whom. Separately, it sends a bank-to-bank instruction into the correspondent chain to move the actual funds, the "cover". The customer message and the cover message travel different routes and can arrive at different times.

This matters more than it sounds. Under the cover method the intermediary banks moving the money historically saw only a bank-to-bank instruction with no detail of the underlying customer. That gap was exploited: institutions stripped out the names of sanctioned parties, and intermediaries screened a message that no longer contained anything to find. The standards were changed in 2009 to force the underlying customer details into cover messages, and the enforcement actions that followed are among the largest financial penalties ever imposed. The simple chain diagram hides all of this.

■ A day per hop is a symptom, not a mechanism

It is tempting to read the three-day journey as three banks each taking a day. Nothing in the system works like that. Delivery of the message is measured in seconds. The delay is made of four separate things, worth separating because they respond to different fixes.

The first is **cut-off times**. Every bank publishes a deadline for each currency, after which an instruction becomes next business day's work. Miss it by a minute and you have lost a day. These are bank-specific and published in each bank's own tariff and service schedules; read the ones that apply to you rather than any general claim about them.

The second is **business calendars**. A payment in dollars needs the American clearing system open, a payment in euros the European one. A chain crossing three currencies must satisfy three calendars, and every weekend and public holiday in any of them stops the clock.

The third is **value dating**, which is bookkeeping rather than movement. A payment can be credited today with a value date of tomorrow, so the recipient sees it but cannot use it.

The fourth is **compliance screening**, in practice the biggest single source of unexplained delay in the whole system. It gets its own treatment below.

None of these is fixed by making the network faster. The network was never the bottleneck.

■ SWIFT is not the only way, and not every hop takes a fee

Two further corrections. First, plenty of cross-border value moves without any SWIFT message. A bank moving money between its own branches in two countries does it on its own internal ledger. Card networks run their own messaging and settlement. The euro area's retail schemes run over their own clearing infrastructure. China and Russia have built their own interbank messaging arrangements. Money transfer operators run closed-loop networks where both ends are theirs and the "cross-border" element is an internal transfer plus local funding on each side. SWIFT is dominant in interbank correspondent messaging. It is not a synonym for international payment.

Second, the fee-per-hop is a consequence of a choice recorded in the message, not an automatic toll. Payments can be sent on terms where the sender bears every charge in the chain and the beneficiary

receives the full amount, on terms where charges are shared, or on terms where the beneficiary bears everything. Within the United Kingdom and the European Economic Area, the shared-charge basis is effectively mandated by payment services regulation for payments in the relevant currencies, which is why intra-European transfers rarely arrive short and transfers to Türkiye or Nigeria frequently do.

The technical version

■ What SWIFT actually is

The legal entity is a cooperative society under Belgian law, headquartered in La Hulpe, owned and controlled by its member institutions, governed by a board of twenty-five directors elected by shareholders. It was founded in 1973 and carried its first live traffic in 1977.

The scale, on SWIFT's own published figures for 2024: more than 11,500 connected institutions, more than 220 countries and territories, an average of more than 53 million FIN messages per day, and FIN availability of 99.99 per cent. Note carefully what that message count is and is not. FIN traffic includes securities, treasury, foreign exchange, trade finance and reporting messages as well as payments. It is not a count of payments, and each payment in a chain generates several messages.

The services are worth naming precisely, because practitioners use the names as shorthand.

Service	What it carries
SwiftNet	The underlying secure IP network and its interfaces
FIN	Store-and-forward messaging for the legacy MT message set
FINplus	Store-and-forward messaging for ISO 20022 messages, including CBPR+
InterAct	Real-time and store-and-forward exchange of structured XML messages
FileAct	Bulk file transfer, used for large payment files and reporting

Security of the endpoints, as opposed to the network, has been SWIFT's largest operational preoccupation since 2016, when attackers who had compromised a member's own systems used entirely valid credentials to issue fraudulent instructions from the Bangladesh Bank. The network worked perfectly; the customer's premises did not. The response was the Customer Security Programme and its Customer Security Controls Framework, against which members must attest annually. This is the correct frame for SWIFT security generally: the network is extremely hard to break, and every significant incident has been an endpoint compromise.

■ Governance and oversight

SWIFT is systemically important without being a financial institution, which makes it awkward to regulate. The arrangement the central banks settled on is cooperative oversight. The National Bank of Belgium is lead overseer, because SWIFT is incorporated in Belgium, supported by the G-10 central banks: Banca d'Italia, Bank of Canada, Bank of England, Bank of Japan, Banque de France, De Nederlandsche Bank, Deutsche Bundesbank, the European Central Bank, Sveriges Riksbank, the Swiss National Bank, and the Federal Reserve System, represented by the Federal Reserve Bank of New York and the Board of Governors. A wider SWIFT Oversight Forum brings in central banks from other major economies.

The oversight is conducted against a set of High Level Expectations covering risk identification and management, information security, reliability and resilience, technology planning, and communication with users. What it is not is prudential supervision, and the reason is stated plainly by SWIFT itself: it is “neither a payment nor a settlement system”, and in its own words, “we don’t actually move money.” There is no settlement risk to supervise because there is no settlement, only operational and security risk, and that is exactly what the oversight covers.

■ Nostro and vostro

The vocabulary is Italian and, once you see the trick, trivial. The same account has two names depending on whose books you are reading it in.

Nostro means “ours”. A nostro account is our account, held with you, denominated in your currency. On our balance sheet it is an asset: money owed to us.

Vostro means “yours”. A vostro account is your account, held with us, denominated in our currency. On our balance sheet it is a liability: money we owe you. Older documentation sometimes uses **loro**, “theirs”, for an account a bank services on behalf of a third party.

A UK bank’s dollar account at a New York bank is that UK bank’s nostro and the New York bank’s vostro. One account, one balance, two labels, two sides of two balance sheets. The confusion in practice comes from people using “nostro” as a generic term for correspondent accounts regardless of viewpoint, which is harmless until it appears in a reconciliation specification.

Reconciling nostro accounts is a substantial operational function in its own right. The account servicer reports movements and balances, and the account holder matches them against its own expected entries. In ISO 20022 terms that is a camt.053 end-of-day statement, a camt.052 intraday report, and camt.054 debit and credit notifications, which replaced the MT 940 and MT 950 statements, the MT 942 interim report, and the MT 900 and MT 910 confirmations respectively. Aged unmatched items are a regulatory and audit concern, because an unexplained entry in a correspondent account is exactly what a fraud or a misdirected payment looks like.

■ The shrinking network

The correspondent network has been contracting for well over a decade, and the data is unusually good because it is derived from the messages themselves. The Committee on Payments and Market Infrastructures, working with SWIFT, published an annual quantitative review built on MT 103 and MT 202 traffic, excluding MT 202 COV, covering more than two hundred jurisdictions.

The headline is stark. The number of active correspondents worldwide fell by about 22 per cent between 2011 and 2019, with a further roughly 3 per cent fall in 2019 alone, while the number of active corridors fell by about 12 per cent over the same period. The decline is universal but wildly uneven: over 2011 to 2022 the fall was about 19.4 per cent in Northern America and about 47.1 per cent in the Americas excluding Northern America, with Oceania down about 37.8 per cent. Small island developing states were hit hardest of all. Meanwhile the volume and value of payments flowing through the shrinking network kept rising. Fewer banks, carrying more. Concentration in this business means a single institution’s decision to exit a country can leave that country materially harder to pay.

The currency picture from the same data, as at December 2022, explains why the chain so often runs through New York. Of the value transferred, the US dollar accounted for 50.9 per cent, the euro for 27.3 per cent and sterling for 4.3 per cent.

The drivers of the decline are mostly regulatory in origin. A correspondent must satisfy itself about its respondent’s anti-money-laundering controls, and to some degree about the respondent’s own cus-

tomers. That due diligence is expensive, the penalties for getting it wrong are enormous, and the revenue from a small respondent in a high-risk jurisdiction is not. The rational response for a large bank is to exit, and it is called de-risking. Attempts to reduce the cost side of that calculation include the Wolfsberg Group's Correspondent Banking Due Diligence Questionnaire, which standardises what is asked, and SWIFT's KYC Registry, which standardises where the answers are kept.

■ The last mile is always domestic

An international payment always ends as a domestic settlement in the currency's home system, and this is the single most useful thing to hold in your head when diagnosing one. Dollars settle across the Federal Reserve's Fedwire Funds Service or across CHIPS. Euros settle across the Eurosystem's T2 or across EURO1. Sterling settles across CHAPS, inside the Bank of England's RTGS system, and for lower-value onward legs across Faster Payments. When a payment "arrives in the UK", what has happened is that a UK institution's sterling position at the Bank of England has increased, or a UK bank has credited a foreign bank's sterling vostro account.

For a UK practitioner the consequence is that inbound international payments land on domestic rails and behave like domestic payments. An inbound sterling payment reaching a UK clearing bank for onward delivery will go out over CHAPS if it is large or time-critical, and CHAPS is same-day, irrevocable, real-time gross settlement with an early-afternoon cut-off for the same-day guarantee. Miss that cut-off and it is tomorrow's payment. Alternatively the onward leg may go over Faster Payments, which settles in seconds and runs 24/7 with a maximum of GBP 1 million per transaction, subject to lower limits set by individual institutions. Neither rail knows or cares that the money came from abroad. They see a domestic instruction.

The domestic rails have themselves been converging on the same standard. The Fedwire Funds Service completed its migration to ISO 20022 on 14 July 2025. The Bank of England moved CHAPS to ISO 20022 and has been ratcheting up the data requirements, mandating Legal Entity Identifiers and purpose codes on certain CHAPS payments from 1 May 2025, with further enhanced data requirements consulted on for 2027. The point is that the same structured data now survives the whole journey instead of being flattened at each border.

■ MT and MX

This is where the vocabulary trips people, so take it slowly.

MT stands for Message Type. It is SWIFT's proprietary format, in service since 1977, carried over FIN. An MT message is a flat text structure of five blocks: block 1 the basic header, block 2 the application header, block 3 the user header, block 4 the text with the actual content, and block 5 the trailer. Inside block 4 the content sits in tagged fields, so field 20 is the sender's reference, field 32A is the value date, currency and amount, field 50 identifies the ordering customer, field 57 the account-with institution, field 59 the beneficiary, field 70 the remittance information, and field 71A the details of charges.

Message types are numbered by category, and the first digit tells you the business.

Category	Business
0	System messages
1	Customer payments and cheques
2	Financial institution transfers
3	Treasury markets: foreign exchange, money markets, derivatives
4	Collections and cash letters
5	Securities markets

Category	Business
6	Commodities and syndications
7	Documentary credits and guarantees
8	Travellers cheques
9	Cash management and customer status

So an MT 103 is a single customer credit transfer, an MT 202 is a general financial institution transfer, and an MT 940 is a customer statement. The messages ending in 92, 95, 96 and 99 in each category are the common group messages: cancellation requests, queries, answers and free-format.

The format's constraints were reasonable in 1977 and became crippling. Names and addresses are unstructured free text in lines of thirty-five characters, typically four of them. Remittance information in field 70 is four lines of thirty-five characters: a hundred and forty characters to describe why the payment is being made. The permitted character set is restricted, which mangles accented names. And anything too long is truncated silently. A beneficiary address running to five lines loses its fifth line, and the sanctions filter downstream screens what remains.

MX is the informal name for the ISO 20022 XML messages that replace MT. ISO 20022 is not a message format so much as a modelling methodology with a shared dictionary of business concepts, from which messages are generated. Message identifiers follow a fixed pattern: a four-letter business area, a three-digit message number, a three-digit variant and a two-digit version. So pacs.008.001.08 is business area pacs, message 008, variant 001, version 08.

The business areas you will meet in payments are pacs for payments clearing and settlement, pain for payment initiation, camt for cash management, acmt for account management, and head for the Business Application Header that wraps the message.

The broad correspondence is as follows. Treat this as a map of intent, not a field-level mapping, because the two formats do not correspond one to one and that is the entire point of the change.

MT	ISO 20022	Purpose
MT 103	pacs.008	Customer credit transfer
MT 202	pacs.009 core	Financial institution transfer
MT 202 COV	pacs.009 COV	Cover payment carrying underlying customer detail
MT 205	pacs.009	Domestic financial institution transfer
MT 103 return	pacs.004	Payment return
MT 210	camt.057	Notification to receive
MT 900 / MT 910	camt.054	Debit and credit confirmation
MT 940 / MT 950	camt.053	End-of-day statement
MT 942	camt.052	Intraday report
MT n92	camt.056	Cancellation request
MT n96	camt.029	Resolution of investigation
MT 101	pain.001	Request for transfer / credit transfer initiation

What ISO 20022 buys is data. Party names and addresses in discrete elements rather than free text, so a town is tagged as a town and a country as a country. Remittance information both structured

and unstructured, and far longer. Purpose codes drawn from a controlled list. Room for Legal Entity Identifiers alongside BICs. An extended character set. And end-to-end identifiers that survive the journey.

■ The migration, and the date that mattered

The cross-border implementation of ISO 20022 is governed by **CBPR+**, Cross-Border Payments and Reporting Plus: a set of usage guidelines defining exactly how the ISO 20022 messages are to be populated for cross-border traffic, developed with the Payments Market Practice Group and published through SWIFT's MyStandards. CBPR+ messages travel over FINplus, not FIN. The guidelines were released in tranches, with the initial set covering the core payment and reporting messages, later additions covering cheques, direct debits, cancellations, and charges notifications and requests.

The coexistence period, during which institutions could send either MT or ISO 20022 for cross-border payments and reporting, ended on **22 November 2025**. That is the date the industry had been working towards for years, and it is the reason a great deal of bank technology spending happened in 2024 and 2025.

What happens now, for an institution that has not migrated, is worth stating precisely because it is widely misreported. MT categories 1, 2 and 9 are no longer supported natively for cross-border payments and reporting. Institutions still sending MT are put through automatic conversion, with no opt-in, and institutions still needing to receive MT can use an in-flow translation service. From 1 January 2026, both the contingency processing for MT senders and the in-flow translation service are chargeable. So the MT world did not vanish on 22 November 2025; it became a paid fallback that degrades the data. Translation from ISO 20022 to MT is lossy by definition, because the structured fields have nowhere to go, and a bank running on translated MT is receiving less than it is being sent.

There is a second-order effect that shows up in the performance statistics. The Financial Stability Board's key performance indicators recorded a fall in the share of wholesale cross-border payments credited within one hour between 2023 and 2024, and attributed part of it to the migration itself: message translation takes additional processing time. Standardisation is worth it, and it is not free while it is happening.

■ The UETR and SWIFT gpi

The **UETR**, the Unique End-to-end Transaction Reference, is the identifier that makes tracking possible. It is a string of 36 characters, generated by the originator, and it must not be altered by anyone in the chain. In MT messages it lives in field 121 in block 3, the user header. In ISO 20022 it is carried in the Business Application Header and in the payment identification. It is required in the relevant payment instruction messages, including MT 103 and its variants and MT 202, MT 205 and their cover variants, for all SWIFT users originating payments, whether or not they are gpi members.

SWIFT gpi, the global payments innovation initiative, launched in 2017, is not a rail and not a message format. It is two things: a rulebook that participating banks sign up to, and a central tracker database. The rulebook commits members to same-day use of funds within the receiving bank's business day, transparency of the fees deducted, unaltered transmission of remittance information, and end-to-end tracking. The tracker is a database, hosted by SWIFT, that every bank in the chain updates as the payment passes through, recording status, timestamps, charges taken and exchange rate applied. Because the UETR is unaltered and unique, all those updates from different banks in different countries attach to the same payment.

SWIFT's own published performance figures are that nearly 60 per cent of gpi payments are credited to end beneficiaries within 30 minutes and almost 100 per cent within 24 hours, and that hundreds of the largest cash management banks send over USD 300 billion via gpi every day.

Around the tracker sits a set of services that turn visibility into action: case resolution, so investigations run as structured exchanges rather than free-format messages; stop and recall, so a payment can be chased down the chain if it was sent in error or under fraud; and pre-validation, so a beneficiary account can be checked before the payment is sent rather than bouncing days later.

■ Two sets of numbers that do not agree, and why

Here is a genuine trap. SWIFT reports that almost 60 per cent of gpi payments are credited within 30 minutes. The Financial Stability Board reports, for 2024, that 50.6 per cent of cross-border wholesale payments were credited within one hour and 92 per cent within one business day, with retail at 33.5 per cent and 69 per cent and remittances at 53 per cent and 76.6 per cent. The 2025 progress report describes only a slight improvement at global level since the indicators were first calculated in 2023, with wholesale and remittance speed improving and cost remaining stubborn.

Both sets of numbers are honest. They measure different things. The gpi figure covers gpi payments, sent by banks that joined the scheme along corridors they serve, and the clock is the bank-to-bank leg. The FSB indicators are drawn from a wider population and are closer to end-to-end. If you are asked why your bank's tracker says 30 minutes and the customer says three days, that gap is the answer, and it is mostly made of the domestic legs at each end, the cut-offs between them, and the compliance queue.

The G20 targets are the yardstick everyone is now measured against: by the end of 2027, 75 per cent of cross-border wholesale, retail and remittance payments credited or available within one hour and the remainder within one business day; a global average retail cross-border cost of no more than 1 per cent with no corridor above 3 per cent; and, restating the UN Sustainable Development Goal, a global average cost of sending a USD 200 remittance of no more than 3 per cent by 2030 with no corridor above 5 per cent. Against that last one, the World Bank's Remittance Prices Worldwide put the global average cost of sending USD 200 at 6.49 per cent in the first quarter of 2025, with digital services at 4.85 per cent and non-digital at 7.16 per cent, and the global weighted average at 4.92 per cent. The distance between 6.49 and 3 is the whole of the reform agenda.

■ Sanctions screening, and where the days actually go

Every institution in the chain screens every payment, and it screens the message, not the payment. This distinction is the source of most of the friction.

The lists screened against include, for a UK institution, the UK Sanctions List and the consolidated list of asset freeze targets maintained by the Office of Financial Sanctions Implementation; for anything touching dollars, the Specially Designated Nationals and Blocked Persons list and other lists maintained by the United States Office of Foreign Assets Control; the European Union's consolidated list; and the United Nations Security Council consolidated list. Institutions add their own lists of high-risk entities, jurisdictions, and typologies on top.

The screening engine reads names, addresses, countries, BICs, and free-text fields, and applies fuzzy matching, because criminals do not helpfully spell their names the way the list does. Fuzzy matching generates false positives at a rate that would astonish anyone outside the industry, and every one of them is a stop. The payment is held, an analyst reviews it, and if the message does not contain enough information to clear the alert, the bank sends a request for information back up the chain. In ISO 20022 that is a camt.026 unable to apply or a camt.028 additional payment information, answered by

camt.029; in the old world it was an MT 195 query and MT 196 answer. The round trip crosses time zones, and the answer often has to come from the originating customer, who is asleep.

Three consequences follow, and they are the real technical content of this section.

First, **data quality is a speed feature**. The single most effective way to make an international payment fast is to populate the beneficiary name and full structured address correctly. A truncated or free-text address is what turns a fifteen-second payment into a two-day investigation. This is the strongest practical argument for ISO 20022, and it is why the enhanced data mandates on CHAPS and elsewhere are framed around structured addresses, purpose codes and LEIs. Structured fields let the screening engine compare a town against a town rather than a hundred and forty characters of prose against everything.

Second, **the cover method has a history**. Plain MT 202 cover messages that omitted the underlying parties were used to defeat screening, deliberately, by stripping identifying data. MT 202 COV was introduced in the 2009 standards release specifically to force the underlying customer information into the cover message so that intermediaries screened the real parties. The enforcement that followed produced some of the largest financial penalties on record, including BNP Paribas's guilty plea and approximately USD 8.9 billion penalty in 2014 for sanctions violations. The lesson institutionalised across the industry is that an intermediary is responsible for what passes through it, and cannot rely on the sender.

Third, **the data requirements are about to get heavier**. At its June 2025 plenary the Financial Action Task Force agreed revisions to Recommendation 16, the wire transfer rule, published on 18 June 2025 and intended to take effect by the end of 2030. They require enhanced and verified originator information and enhanced beneficiary information on cross-border payments. Above a threshold of USD or EUR 1,000, the message must carry the name and account number of both originator and beneficiary, or a unique transaction reference; for individuals, the originator's address and the beneficiary's country and town; and for legal entities, a BIC, a Legal Entity Identifier, or another unique official identifier. Card transactions for the purchase of goods and services remain outside the full requirements. Below the threshold the data set is smaller and need not be verified absent suspicion.

Read that list of required fields against the ISO 20022 message structure and the design intent becomes obvious. The standards migration and the transparency rules are the same project seen from two directions.

■ Sanctions as an instrument, and the network as a chokepoint

Because SWIFT is a Belgian entity, it is bound by European Union law. When the EU has required it, SWIFT has disconnected named institutions from its network, most prominently a set of Russian banks from 2022. Be precise about what disconnection does. It does not freeze anyone's money; balances in correspondent accounts are frozen, if at all, by asset-freeze measures applied by the banks holding them, not by a messaging network. What disconnection removes is the standard, automated way to instruct. A disconnected bank can still be paid by any institution willing to use telex, email, a bilateral link or another network, subject to whether doing so is lawful.

That is why alternative arrangements exist and are growing: China's Cross-Border Interbank Payment System for renminbi, Russia's domestic financial messaging system, regional systems such as Buna in the Arab region, and a lengthening set of direct links between national instant payment systems. None approaches SWIFT's reach. All reduce the degree to which any single network is a chokepoint, which is precisely why some states are building them.

■ Where this is going

The standards are converging. ISO 20022 is now the format for cross-border messaging, for Fedwire, for CHAPS, for T2 and for a growing list of domestic high-value systems. The same structured payment can, in principle, travel from a corporate ledger through a domestic rail, across a correspondent chain and out through another domestic rail without being flattened.

The service levels are hardening. SWIFT gpi set rules for wholesale payments; SWIFT Go extended a service level to low-value payments; and on 25 September 2025 SWIFT announced new rules for retail cross-border payments on its network, developed with over thirty banks from seventeen countries, committing to upfront transparency on costs, full value delivery, end-to-end visibility, and instant settlement where domestic infrastructure and regulation allow.

And the underlying model is being questioned. The Bank for International Settlements has framed the correspondent banking problem explicitly in terms of the cost and failure rate of the current chain, and Project Agorá, its work with central banks and private banks on tokenised deposits and central bank money on a shared programmable platform, is an attempt to collapse the sequence of separate ledgers into one. Whether that works is not yet knowable. What is knowable is that the friction it targets is real, measured, and has not improved much in a decade.

■ The Sheffield payment, explained

The USD 214,000 never left the United States. Her bank's dollar position sat as a claim on a New York correspondent, was debited there, moved across a domestic dollar clearing system to a second New York bank, and was credited to the Turkish bank's dollar claim, also in New York. The steel mill was paid in Izmir out of the Turkish bank's own balance sheet, against the increase in what New York owed it. SWIFT carried the instruction, in seconds, and knew nothing about any of the balances involved. It could not have told her where the money was, because at every instant the money was in New York and only the entitlement to it was moving.

The eighty-five dollars were three tariffed charges taken under a shared-charge instruction, each disclosed in a document she had not read, none of them visible at the point of sending. Had the payment been sent on terms where the sender bears all charges, the mill would have received the full amount and her own bank would have billed her more. Had the corridor been intra-European, the shared basis would have been mandatory and the deductions would not have happened.

The three days were not three banks taking a day each. They were one cut-off missed by hours, one time zone, and one sanctions filter that stopped on a fuzzy name match and needed a full structured address the original message had never carried, because thirty-five characters was not enough.

Every part of that is fixable. Most of it is being fixed, and the fix is not a faster network. It is better data in the message, one identifier that survives the whole journey, and an agreement among the banks about what they owe the customer. That is what ISO 20022, the UETR and the gpi rulebook respectively are, and it is why a chapter about a messaging network turns out to be a chapter about bookkeeping, deadlines and paperwork.

41.98 Common wrong ideas

Wrong: SWIFT is the system that sends money abroad. **Right:** SWIFT sends messages, has no view of anyone's balance, holds no accounts, and does not know whether the payment it carried actually happened.

Wrong: a bank's dollars in New York sit in a pot of money that belongs to it. **Right:** they are an unsecured claim on the New York bank, recorded on that bank's ledger, and if the New York bank fails its correspondent joins the queue of creditors.

Wrong: a payment that takes three days means three banks each took a day. **Right:** message delivery is measured in seconds, and the days are made of cut-off times, business calendars, value dating and a compliance queue.

Wrong: every hop in the chain automatically takes a fee out of the money. **Right:** the charge basis is a choice recorded in the message, and within the United Kingdom and the European Economic Area the shared basis is effectively mandated for the relevant currencies.

Wrong: the payment instruction always travels alongside the funds. **Right:** under the cover method the customer message goes straight to the beneficiary's bank while a separate cover message moves the funds down the correspondent chain, and the two can arrive at different times.

Wrong: SWIFT is a synonym for international payment. **Right:** intragroup transfers, card networks, euro-area retail schemes, China's and Russia's own arrangements and money transfer operators' closed loops all move cross-border value without a SWIFT message.

Wrong: disconnecting a bank from SWIFT freezes its money. **Right:** disconnection removes the standard automated way to instruct; balances are frozen, if at all, by asset-freeze measures applied by the banks that hold them.

Wrong: MT messages ceased to exist on 22 November 2025. **Right:** the coexistence period ended, and MT became a chargeable fallback from 1 January 2026, delivered through automatic conversion and lossy in-flow translation.

Wrong: the tracker says thirty minutes, so a customer complaining of three days must be mistaken. **Right:** the gpi figure covers the bank-to-bank leg of gpi payments, while the wider indicators are closer to end to end, and the gap is the domestic legs, the cut-offs and the compliance queue.

Wrong: SWIFT's security record means the network has been broken into. **Right:** every significant incident, including the Bangladesh Bank fraud, has been an endpoint compromise using entirely valid credentials, which is why the Customer Security Programme is aimed at members' own premises.

41.99 Chapter summary in 20 lines

1. SWIFT is a Belgian messaging cooperative owned by its members, not a payment system: it has never moved money and holds no accounts.
2. An international payment is a sequence of separate domestic bookkeeping entries in different countries, stitched together by messages.
3. Correspondent banking means two banks holding accounts with each other, so that value is paid locally at each end while only the instruction crosses the border.
4. No bank can hold an account with every other bank, so payments run down a chain through a small number of large correspondents.

5. Dollars only really exist in America, so a dollar chain must pass through New York, and the dollars themselves never leave it.
6. What a bank “holds” at a correspondent is an unsecured claim on that institution, sitting inside another country’s legal system and reachable by that country’s authorities.
7. That is the mechanism by which one country’s sanctions policy binds banks with no other connection to it, and the largest single reason correspondent banking has become so cautious.
8. The eighty-five dollars missing from the Sheffield payment were three tariffed charges, each disclosed somewhere, taken lawfully under a shared-charge instruction.
9. Within the United Kingdom and the European Economic Area the shared basis is effectively mandated for the relevant currencies, which is why intra-European transfers rarely arrive short.
10. The three days were not three banks taking a day each; the messages were delivered in seconds.
11. Delay is made of cut-off times, business calendars, value dating and compliance screening, and none of them is cured by a faster network.
12. Screening reads the message rather than the payment, so truncated or unstructured data is what turns a fifteen-second payment into a two-day investigation.
13. The cover method splits the customer message from the funds message, a gap that was exploited by stripping names until MT 202 COV forced the underlying parties back into the cover.
14. MT is the 1977 flat-text format with thirty-five-character lines and silent truncation; ISO 20022 replaces it with structured, tagged and far longer data.
15. CBPR+ defines how ISO 20022 is populated for cross-border traffic over FINplus, and the coexistence period ended on 22 November 2025, leaving MT as a chargeable, lossy fallback.
16. The UETR is a thirty-six-character identifier that nobody in the chain may alter, which is what lets updates from banks in different countries attach to one payment.
17. SWIFT gpi is a rulebook plus a central tracker, not a rail: it made payments visible rather than fast, which turned out to be most of what customers wanted.
18. SWIFT’s thirty-minute figure and the Financial Stability Board’s one-hour figure are both honest, because they measure different populations over different legs.
19. The correspondent network has been shrinking for over a decade through de-risking while carrying ever more traffic, which concentrates the system and can leave a whole country harder to pay.
20. Every international payment ends as a domestic settlement in the currency’s home system, and the fix for its friction is better data in the message, one identifier that survives the journey, and an agreement about what the banks owe the customer.

Sources: Swift, on FIN traffic figures, organisation, governance and oversight, ISO 20022 and CBPR+ implementation, the UETR, gpi, and the retail cross-border rules announced 25 September 2025; National Bank of Belgium and Deutsche Bundesbank on the cooperative oversight of Swift; BIS Committee on Payments and Market Infrastructures correspondent banking data commentaries and chartpacks, and BIS Bulletin No 87; Financial Stability Board G20 targets and the 2024 and 2025 progress reports; World Bank Remittance Prices Worldwide, Issue 53, March 2025; Federal Reserve Financial Services on the Fedwire Funds Service ISO 20022 migration of 14 July 2025; Bank of England on ISO 20022 and CHAPS enhanced data; FATF, update to Recommendation 16, June 2025; ISO 20022 Registration Authority message catalogue.

UPI

42.0 What this chapter gives you

1. You will be able to explain why NPCI never holds a rupee of Meera's forty and yet determines where every one of them goes.
2. You will be able to say what happens between Ravi's speaker announcing forty rupees and his bank actually being made whole, and name the risk his bank carries in the meantime.
3. You will be able to explain why UPI carries 85.5 per cent of India's digital payment volume and only 9.5 per cent of its value, and why that is the same shape as CHAPS and Faster Payments in Britain.
4. You will be able to describe single-click two-factor authentication and say why removing the OTP round trip is what made a forty-rupee payment worth making.
5. You will be able to tell a customer which twelve-digit number to write down when a payment goes wrong, and what the Reserve Bank of India's turnaround-time circular entitles them to without asking.
6. You will be able to distinguish a failed payment from a payment that succeeded to the wrong person, and give the different remedy for each.
7. You will be able to explain deemed status, and why a payment sitting pending for hours is neither a success nor a failure.
8. You will be able to read a UPI QR code as a `upi://pay?` deep link and say which parameters make it static and which make it dynamic.
9. You will be able to say what zero merchant discount rate is, what the amendment of 6 August 2026 changed, and why nobody should model UPI economics as though free acceptance were a law of nature.
10. You will be able to design a UPI integration that survives the API rate limits, the abolition of person-to-person collect, and the fact that every number in the specification moves.

There is a country where roughly two thirds of a billion payments happen every day, where the tea seller with a plastic stool and a kettle takes forty rupees by phone without paying a fee for the privilege, and where the whole thing runs on a piece of national infrastructure owned by a not-for-profit company that no single bank controls. That country is India, and the thing is the Unified Payments Interface. It is the highest-volume retail payment system on earth, and it is worth understanding properly, because almost everything written about it is either breathless or wrong.

The plain version

Imagine a very old-fashioned telephone exchange. Every house in town has a wire running to one building in the middle. Inside that building sits an operator with a board full of sockets. When you pick up your phone you do not dial a number, you say a name: “the baker on the high street, please.” The operator knows which socket belongs to the baker, plugs your wire into his, and the two of you are connected. You never had to know which socket he was on. You never had to know that his wire runs down Mill Lane and under the canal. You said a name and the operator did the rest.

UPI is that operator, for money.

In India, every bank account has a long number and a branch code, exactly as a British account has a sort code and an account number. Nobody enjoys typing those. So UPI lets you register a short nickname that points at your account. The nickname looks like an email address: `meera@okhdfcbank`, `ravi@ybl`, `stalls9971@paytm`. It is called a virtual payment address. When somebody wants to pay you, they type the nickname, or they scan a square barcode that contains the nickname, and the operator in the middle looks it up and connects the two banks.

The operator is a company called NPCI, the National Payments Corporation of India. It sits between every bank in the country. It does not hold your money. It holds the address book and it passes the messages.

Now the worked example, with real numbers.

Meera stops at Ravi’s tea stall on her way to work. Two glasses of chai, forty rupees, which is about thirty-five pence. Ravi has no card machine. He has a laminated square of paper taped to the counter with a black-and-white pattern printed on it. Meera opens her payments app, points the camera at the square, and the app fills in Ravi’s nickname automatically. She types 40. The app shows her a name: **RAVI KUMAR**. That name did not come from the square of paper. It came back from Ravi’s bank, a quarter of a second ago, as a check that she is paying who she thinks she is paying. She taps to confirm and types a four-digit PIN.

What happens next takes about a second and a half. Her app sends a message to her bank. Her bank sends it to NPCI. NPCI asks Ravi’s bank to confirm the address is real. Meera’s bank checks her PIN, checks she has forty rupees, and takes forty rupees off her balance. NPCI then tells Ravi’s bank to put forty rupees on his balance. Ravi’s bank does it and says so. A small speaker sitting next to Ravi’s kettle, the kind every stallholder in India now owns, says out loud in Hindi: “Forty rupees received.” Meera is already walking away.

Here is the part that surprises people. No money has actually moved between the two banks. Meera’s bank is forty rupees richer. Ravi’s bank is forty rupees poorer, because it credited Ravi out of its own pocket. Both banks made a note of it. Later that day, at one of ten scheduled moments, NPCI adds up every single thing that happened between every pair of banks, works out who owes whom on balance, and tells the Reserve Bank of India to shift the net amounts between the banks’ accounts. Meera’s forty rupees never travels on its own. It is swept up into one enormous sum with millions of other payments and settled in a lump.

Think of it as a group of friends who share a taxi every week and keep a tally on the fridge rather than passing coins about each time. At the end of the week one person hands over one note and the tally goes back to zero.

The scale of that tally is what makes UPI worth a chapter. In July 2026, the most recent full month at the time of writing, UPI processed **23,658.35 million transactions**, which is 23.66 billion, worth INR

29,87,880.49 crore, which is about INR 29.9 lakh crore or roughly 29.9 trillion rupees. **741 banks** were live on the system that month. Divide the volume by thirty-one days and you get about **763 million payments a day**. Divide that by 86,400 seconds and you get roughly **8,800 payments every second**, averaged flat across the whole day and night, which means the peaks are a great deal higher than that.

Divide the value by the volume and you get the number that explains everything: the average UPI payment in July 2026 was worth **INR 1,263**, about eleven pounds. This is not a system for buying flats. It is a system for buying chai, and vegetables, and a bus ticket, and topping up a phone, hundreds of millions of times an hour.

Every one of those payments gets a receipt number. It is usually twelve digits long, and it is the single most useful thing a customer can know, because when a payment goes wrong the receipt number is the only thing that lets a bank find it. More on that shortly.

Where the plain version stops being true

The operator does not move money, and “instant” describes only what the customer sees. The telephone-exchange picture makes it sound as though NPCI holds a great pot of rupees and shifts them about. It does not. NPCI is a switch: it routes messages, applies rules, keeps score, and produces settlement files. The actual money moves between banks’ current accounts held at the Reserve Bank of India, and it moves in batches, not per payment. UPI is a **deferred net settlement** system dressed in real-time clothing. Since 3 November 2025, NPCI has run **ten settlement cycles a day for authorised transactions, between 09:00 and 21:00 at two-hour intervals**, plus two separate cycles a day for disputes, so that a chargeback can no longer delay a payments cycle. Between the moment Ravi’s phone speaks and the moment his bank is actually made whole, his bank is carrying credit risk on Meera’s bank. That risk is small per payment and enormous in aggregate, which is why NPCI caps how far into the red any member may go before it must top up, and why the number of daily cycles has crept upwards over the years. The customer experience is instant. The interbank plumbing is not, and any practitioner who confuses the two will size their liquidity wrong.

The app is not a bank, and the nickname hides less than it appears to. Meera’s app may be PhonePe or Google Pay or Paytm. None of those is a bank. They are third-party application providers, and they ride on the licence of a sponsor bank called a payment service provider bank. Money never sits with the app; the app is a user interface and a fraud-scoring engine bolted onto somebody else’s banking licence. That is a deliberate design choice and it is the single biggest reason UPI got good so fast, but it means that when an app has an outage, the correct question is which sponsor bank is down. As for the nickname: it hides the account number **from the payer**, which is the point, because a stolen account number plus branch code is a useful thing for a fraudster and a stolen ravi@ybl is not. It does not hide anything from NPCI, which resolves it, or from Ravi’s bank, which owns it. And the resolution deliberately leaks something back the other way: the payee’s registered name is returned to the payer before confirmation. That is an anti-fraud feature, and it is also a small, permanent, unavoidable disclosure of who owns a given address.

Volume is not value, and UPI is not where India’s money is. This is the correction that most commentary gets wrong. According to the Reserve Bank of India’s half-yearly Payment Systems Report covering the second half of 2025, UPI accounted for **85.5 per cent of India’s digital payment transactions by volume** but only **9.5 per cent by value**. RTGS, the high-value real-time gross settlement system, accounted for **0.1 per cent of volume and 68.6 per cent of value**. NEFT took 3.6 per cent of volume and 14.9 per cent of value. The shape is identical to the United Kingdom’s, where CHAPS carries a tiny fraction of the payments and the large majority of the sterling value. UPI is the everyday

rail. It is not the wholesale rail, it was never designed to be, and a person who says “India moves its money on UPI” is describing counts, not sums.

Free to the customer is not the same as free, and the law changed in August 2026. Since 1 January 2020, merchant discount rate on bank-account UPI and on RuPay debit has been set at zero by statute, through Section 10A of the Payment and Settlement Systems Act 2007 and Section 269SU of the Income-tax Act 1961, both inserted by the Finance (No. 2) Act 2019. Zero MDR is why Ravi accepts UPI: a card machine would have cost him a rental and a percentage, and a QR sticker costs him the price of the lamination. But zero MDR is a subsidised outcome, not a costless one. NPCI charges member banks a switching fee, banks carry the infrastructure, and the government has run an incentive scheme to compensate the ecosystem for the foregone revenue. **On 6 August 2026 the Lok Sabha passed the Taxation and Other Laws (Amendment) Bill, 2026**, which amends Section 10A of the Payment and Settlement Systems Act to give the central government the power to permit charges on UPI and other notified electronic modes. At the time of writing no charge has been notified and small-value payments are widely expected to stay free, but the statutory bar that made zero MDR permanent has been removed. Anyone modelling UPI economics past 2026 should treat the current position as a policy choice rather than a law of nature. Note also that three flows already sit outside the zero-MDR perimeter today: RuPay credit card on UPI above INR 2,000, prepaid-instrument-on-UPI, and credit line on UPI.

The technical version

■ What NPCI is, and what it is not

The National Payments Corporation of India was incorporated in 2008 as a not-for-profit company, promoted by the Reserve Bank of India and the Indian Banks’ Association, and it operates India’s retail payment systems under authorisation from the RBI under the Payment and Settlement Systems Act, 2007. Its registered office is at 1001A, The Capital, B Wing, Bandra Kurla Complex, Mumbai 400 051, and its corporate identity number, U74990MH2008NPL189067, encodes the “not-for-profit public limited” status that shapes everything about how it behaves. It runs IMPS, NACH, RuPay, NETC FASTag, the National Financial Switch, AePS, BHIM and UPI, and it has two subsidiaries of note: NPCI International Payments Limited, which exports the stack, and NPCI Bharat BillPay.

The relevant comparison for a British reader is Pay.UK rather than Visa. NPCI is the scheme operator and the switch. It is not a bank, it is not a settlement agent, and it does not take a percentage of the transaction. It levies a switching fee on member banks, and NPCI’s own circulars refer to rebate arrangements on that fee through the URCS back-office system, which is where reconciliation, adjustments and chargebacks are handled.

UPI went live as a pilot on **11 April 2016**, launched by Dr Raghuram G. Rajan, then Governor of the Reserve Bank of India, in Mumbai, with **21 banks**. As at March 2026 there were **703 banks live**, and as at July 2026, **741**.

■ The participant stack

A single UPI payment involves more legal entities than most people expect. The roles, in NPCI’s own vocabulary, are these.

The **UPI app** is the interface the customer touches. The **payer PSP** is the payment service provider bank whose licence that app rides on; if the app is a bank’s own app, the two are the same entity, and if the app is a third-party application provider such as PhonePe, Google Pay or Paytm, the PSP is a sponsor bank. Axis Bank and Yes Bank are the two largest sponsors by volume. The **remitter**

bank is the bank that actually holds the payer's account and performs the debit. The **payee PSP** is the equivalent on the receiving side, and the **beneficiary bank** holds the account that gets credited. For merchant payments there is also an **acquiring bank** and, very often, a payment aggregator sitting between the merchant and the acquirer.

The separation of the app layer from the account layer is the architectural decision that made UPI what it is. A company with no banking licence, no branch network and no balance sheet can build a payments app, compete purely on user experience, and reach every bank account in India on day one. That is why the market is so concentrated in the app layer and so fragmented in the bank layer: as at recent months PhonePe and Google Pay between them have run roughly four fifths of UPI volume, against an NPCI rule capping any single third-party application provider at 30 per cent. NPCI first announced that cap in November 2020 and has extended the compliance deadline three times; the current deadline is **31 December 2026**, accurate at the time of writing. Whether it is enforced is one of the genuinely open questions in Indian payments.

■ Addressing: the VPA, the UPI number and the mapper

The virtual payment address takes the form `handle@psp`. The left side is chosen by the customer or generated by the app; the right side identifies the PSP, so `@okhdfcbank`, `@oksbi`, `@ybl`, `@paytm`, `@axl`. The address is not a bank account. It is a pointer, resolved by NPCI's central mapper to a specific account at a specific bank. One customer may have several addresses pointing at the same account, or several accounts behind one address with one nominated as default.

Three other addressing forms exist. A **UPI Number** is a numeric alias, typically the customer's mobile number, mapped to an account through NPCI's mapper; NPCI's circular series on the Numeric UPI ID continues into FY 2026-27. A payer may also address a payment to a raw **account number plus IFSC**, which is how UPI absorbs the older IMPS use case. And a payment may be addressed by **QR code**, which is simply an encoded address plus optional parameters, covered below.

Before any money moves, the payer's app performs an address validation. The resolved **registered beneficiary name** is returned and displayed. This is mandatory, it is the reason UPI's mis-payment rate is as low as it is, and it is the reason a UPI address cannot be treated as private.

■ Authentication: device binding, the Common Library, and single-click two-factor

Indian regulation requires two-factor authentication on electronic payments. Card payments satisfy this with a card plus an OTP, which costs a round trip through a mobile network and is where a large share of card failures come from. UPI satisfies it without any round trip at all, and this is the second architectural decision that made the system work.

At registration, the app sends an outbound SMS from the handset to bind the device to the customer's mobile number as registered with the bank. That establishes the **device** as the first factor, something the customer has. The **UPI PIN**, four or six digits, is the second factor, something the customer knows. Because the device is already bound, no OTP is needed at payment time. NPCI calls the result **single-click two-factor authentication**, and it is worth pausing on how much latency and failure that removes.

The PIN is never seen by the app or by the PSP. It is captured inside NPCI's **Common Library**, a component embedded in every UPI app, which encrypts the keystrokes into a credential block under NPCI's public key. The block travels through the app, the PSP and the switch as opaque ciphertext and is decrypted and verified only at the remitter bank, inside a hardware security module. A compromised app cannot harvest PINs. A compromised PSP cannot replay them.

The PIN itself is set using a second-channel credential: historically the last six digits and expiry date of the customer's debit card, and more recently Aadhaar-based verification, which removed the requirement to hold a debit card at all and opened UPI to the very large population of Jan Dhan account holders who never had one. NPCI's circulars on user onboarding cover this route.

■ The two legs: debit and credit

A UPI push payment is two legs and a switch. Both legs are synchronous request-response messages over the UPI API, an XML message set carried over mutually authenticated TLS with digital signatures on the payload. The messages a practitioner meets are named in matched pairs: ReqPay and RespPay for the payment itself, ReqValAdd and RespValAdd for address validation, ReqAuthDetails and RespAuthDetails for the credential and authorisation exchange with the remitter bank, ReqChkTxn and RespChkTxn for status enquiry, ReqBalEnq for balance, ReqListAccount for account discovery, ReqMandate for e-mandates and ReqComplaint for disputes.

The sequence for Meera paying Ravi is as follows.

The app resolves ravi@ybl with ReqValAdd, which the switch routes to the payee PSP; the response carries the registered name, which the app displays. The customer enters the amount and the PIN; the Common Library produces the credential block. The app issues ReqPay to the payer PSP carrying a **transaction ID of up to 35 characters** generated by the initiating PSP, the amount, both addresses, a purpose code, and, for merchant payments, a merchant category code. The payer PSP signs and forwards it to the UPI switch.

The switch performs the **debit leg**: it presents the credential block and the debit instruction to the remitter bank, which validates the PIN in its HSM, applies its own velocity and value limits, checks the balance, debits the account and responds. Only on a successful debit does the switch perform the **credit leg**, presenting the credit instruction to the beneficiary bank, which credits the payee's account and responds. The switch then returns RespPay to both PSPs, and both customer apps show the outcome.

Three failure shapes matter, and they are the source of most customer distress.

If the debit fails, nothing has happened and the customer sees a decline. Common response codes are 00 for success, Z9 for insufficient balance, ZM for an incorrect UPI PIN and ZA for a transaction declined by the customer; the authoritative list lives in NPCI's API and technical specification and runs to several hundred entries across issuer, acquirer and switch categories.

If the debit succeeds and the credit fails cleanly, the switch initiates a reversal and the money returns, usually within minutes.

If the debit succeeds and the credit **times out** without a definitive answer, the transaction enters an indeterminate state. NPCI's rules define **deemed approval** and **deemed debit** handling for exactly this case, with dedicated circulars on the subject in FY 2025-26 and FY 2026-27, and the position is resolved in the next reconciliation cycle through the URCS back office. This is the state in which the customer's money has genuinely left their account and has genuinely not arrived, and it is the state the next section is about.

A **collect** or pull payment inverts the first half of the flow. The payee's PSP issues ReqPay in collect mode; the switch delivers a pending request to the payer's PSP, which notifies the payer; the payer approves with a UPI PIN inside an expiry window set by the requester; the debit and credit legs then run exactly as above. Collect is how a merchant's website asks a customer's app for money without a redirect, and it is how UPI handles subscription first-payments and IPO applications.

Collect is also, historically, UPI's worst fraud surface, because a request for money looks superficially like an offer of money to an inexperienced user. Fraudsters exploited that relentlessly. NPCI's response was progressive restriction and then abolition: the person-to-person collect limit was reduced to INR 2,000, and **from 1 October 2025 NPCI discontinued person-to-person collect requests on UPI entirely**. Merchant collect remains available to verified merchants. If you are designing a UPI integration today, plan for a world in which pull only exists in the merchant direction.

■ Settlement

Settlement is deferred and net. NPCI accumulates every authorised transaction, computes each member's multilateral net position, and produces a settlement file. Members hold current accounts with the Reserve Bank of India, and the net positions are posted against those accounts. NPCI operates net debit caps and collateral requirements so that no member can accumulate an unfunded net debit position beyond a prescribed limit.

The cycle structure changed on **3 November 2025**. Before that date, authorised transactions and disputes shared the same cycles. From that date, authorised settlement runs in **ten cycles a day between 09:00 and 21:00 at two-hourly intervals**, and disputes run in **two separate cycles**, one covering midnight to 16:00 and one covering 16:00 to midnight. RTGS posting timelines were unchanged. The purpose is to stop a large chargeback batch from delaying a payments cycle, and to make each cycle's arithmetic cleaner for members' treasury teams. This is accurate at the time of writing; settlement cycle structures change, and a practitioner should read NPCI's current circular rather than a book.

For a British reader the mental map is: UPI's customer experience resembles Faster Payments, its settlement model resembles Faster Payments' deferred net settlement rather than CHAPS' real-time gross settlement, and its addressing and app layer occupy the space that Open Banking occupies in the United Kingdom. The difference is that in India those layers are one system with one rule book, whereas in the United Kingdom the rail, the addressing and the API layer are three separate things governed by three separate regimes.

	UPI (India)	Faster Payments (UK)	Bacs (UK)
Operator	NPCI	Pay.UK	Pay.UK
Customer experience	Seconds, 24/7	Seconds, 24/7	Three working day cycle
Settlement model	Deferred net	Deferred net	Deferred net
Addressing	VPA, UPI number, QR, account plus IFSC	Sort code and account number, Paym	Sort code and account number
Per-transaction ceiling	INR 1 lakh default, higher for listed categories	GBP 1 million	Set by scheme and sponsor
Recurring	UPI AutoPay e-mandate	Standing order	Direct Debit

■ Identifiers, and what to do when money vanishes

Every UPI transaction carries at least two identifiers, and confusing them wastes hours.

The **UPI transaction ID** is generated by the initiating PSP and may be up to **35 characters**. It is the app's handle on the transaction. It is excellent for talking to the app's support team and close to useless for talking to a bank.

The **Retrieval Reference Number**, twelve digits, is generated by the switch and carried through both legs. Consumer apps display it under labels including "UPI Ref No", "UPI transaction reference" and,

most commonly, **UTR**, for unique transaction reference. Strictly, UTR is the term the Indian banking system uses for the reference on a NEFT or RTGS credit and RRN is the term for a switched retail transaction, but in ordinary Indian usage the twelve-digit number on a UPI receipt is called the UTR by customers, banks and NPCI's own support material alike. It is the number that both the remitter bank and the beneficiary bank can search on. It is the number to write down.

Now the procedure, because this is the part a customer actually needs.

Money leaves the account and does not arrive. First, distinguish two cases, because they have completely different remedies. If the payment **failed** — the app showed an error, a pending state or a timeout — this is a system failure and the rules are on the customer's side. If the payment **succeeded but went to the wrong person** — the customer typed the wrong address, or was defrauded — this is not a system failure, there is no automatic reversal, and the remedy is a fraud report to the bank and to the National Cyber Crime Reporting Portal, plus, in practice, the beneficiary's consent.

For the failure case, the governing rule is the Reserve Bank of India's circular **RBI/2019-20/67 of 20 September 2019, "Harmonisation of Turn Around Time (TAT) and customer compensation for failed transactions using authorised Payment Systems"**. Its two UPI rows are precise. Where the account is debited but the beneficiary account is not credited on a funds transfer, the beneficiary bank must auto-reverse **latest on T+1 day**, and beyond that the customer is entitled to **INR 100 per day of delay**. Where the account is debited but transaction confirmation is not received at a merchant location, auto-reversal is **within T+5 days**, with the same INR 100 per day beyond it. The compensation is not discretionary and does not require the customer to ask.

The escalation ladder, in order, is: raise the complaint **inside the app**, because UPI carries a native dispute mechanism and apps are required to expose it; if unresolved, go to the **PSP bank**; if still unresolved after thirty days, escalate to NPCI through the UPI Help portal at upihelp.npci.org.in; and beyond that to the **Reserve Bank - Integrated Ombudsman Scheme** through the RBI's complaint management system. At every rung, the first thing anybody will ask for is the twelve-digit reference.

There is one more state worth knowing. A transaction may sit in **deemed** status, where the switch never received a definitive answer from one leg. To a customer it looks identical to a failure. Internally it is not a failure at all; it is an unresolved fact, and it is settled by reconciliation rather than by reversal. NPCI's response-code extensions for deemed debit on mandate execution, issued as an addendum to circular OC-128 in FY 2026-27, exist because this state kept producing wrong customer outcomes on AutoPay. If a customer's payment sits pending for hours rather than seconds, deemed status is usually why.

■ QR codes

A UPI QR code is not a special data format. It is a text string, encoded as a standard QR symbol, containing a **UPI deep link** of the form `upi://pay?` followed by URL query parameters. The parameters a practitioner meets are `pa` for the payee address, `pn` for the payee name, `am` for amount, `cu` for currency, always `INR`, `tr` for a transaction reference, `tn` for a note, `mc` for the merchant category code, `mode` and `sign` for signed intent. A static merchant QR omits `am`, so the customer types the amount; a dynamic QR, generated per bill by a till or a printer, includes both `am` and `tr`, which is what allows automatic reconciliation against an invoice.

The same string, invoked on a handset rather than printed, is an **intent**: tapping a `upi://pay?` link on an Android device offers the user their installed UPI apps and hands the parameters straight across. That is the mechanism behind every "Pay by UPI" button on an Indian checkout page, and it is why a UPI merchant integration requires no card fields, no 3-D Secure and no redirect.

NPCI publishes brand guidelines for merchant QR, with a circular series on mandatory BHIM UPI merchant QR branding running to an addendum in FY 2026-27, and separately supports EMVCo-compliant QR for interoperability with card schemes and for international acceptance. Circular OC No. 236 of FY 2026-27 covers enhancement of international mobile application payment acceptance through UPI QR in India, which is the mechanism that lets a visitor's foreign wallet scan an Indian merchant's sticker.

The economics of the QR are the whole story of UPI's merchant reach. A card terminal costs a merchant a device, a rental, a telephone or data line and a percentage of every sale. A UPI QR costs the price of printing. The Reserve Bank of India's Annual Report for 2024-25 recorded UPI QR codes rising 91.5 per cent to **65.8 crore**, that is 658 million, against a point-of-sale terminal estate roughly two orders of magnitude smaller. There is no version of the card model that reaches a tea stall. There is an obvious version of the sticker model that does.

■ Collect requests, mandates and UPI AutoPay

UPI AutoPay, launched in 2020, is the recurring-payments layer. It is a genuine e-mandate rather than a stored credential: the customer authorises a mandate once, with additional factor of authentication, and the mandate is registered against their account at the remitter bank with a **Unique Mandate Number**. Subsequent executions do not require the PIN.

The parameters a practitioner needs are these. The customer receives a **pre-debit notification at least 24 hours before any execution**, so that funds can be arranged; NPCI waives the pre-debit notification only for auto-replenishment of NETC FASTag and RuPay NCMC balances. The **additional factor of authentication limit** — the ceiling below which no per-execution authentication is needed — was raised from **INR 15,000 to INR 1,00,000** for credit card bill payment, mutual fund and insurance use cases. The customer can modify, revoke, pause and unpause a mandate from within any UPI app, and that control must be exposed in the app rather than buried with the merchant, which is a materially stronger consumer position than a card-on-file subscription.

Frequency options run from one-time through daily, weekly, fortnightly, monthly, bi-monthly, quarterly, half-yearly, yearly and as-presented, which is why AutoPay swallowed use cases as varied as OTT subscriptions, insurance premiums, SIP mutual fund instalments, loan EMIs and utility bills.

One operational constraint matters for anyone building on AutoPay. From **1 August 2025**, NPCI began rate-limiting UPI APIs to protect the switch at peak. Balance enquiries are capped at **50 per app per day**, non-customer-initiated API calls are restricted during peak hours, and **peak hours are defined as 10:00 to 13:00 and 17:00 to 21:30**. Mandate executions are pushed into off-peak windows. If your batch of mandate executions is scheduled for eleven in the morning, it will not run when you expect it to.

■ Limits

NPCI sets ceilings; member banks set their own limits within them, and often lower. The default per-transaction and per-day ceiling for ordinary person-to-person and person-to-merchant use is **INR 1,00,000**, with banks free to impose tighter caps, and many impose a first-24-hour cap on a newly linked account and a cap on the number of transactions per day.

Higher ceilings apply to specific merchant categories, and the governing document is NPCI circular **NPCI/UPI/2025-26/OC/185B dated 28 August 2025**, an addendum to OC 185A, with a compliance date of **15 September 2025**. The enhanced limits apply only to merchants classified as **Verified Merchant** and compliant with NPCI's merchant onboarding circulars. Person-to-person limits were explicitly left unchanged.

Category	Merchant category code	Per transaction	Cumulative, 24 hours
Capital markets	6211	INR 5 lakh	INR 10 lakh
Insurance	5960	INR 5 lakh	INR 10 lakh
Government e-Marketplace, EMD payments	9311	INR 5 lakh	INR 10 lakh
Travel	4722	INR 5 lakh	INR 10 lakh
Credit card bill payments	5413	INR 5 lakh	INR 6 lakh
Collections	7322	INR 5 lakh	INR 10 lakh
Business or merchant, including pre-approved payments	Payer MCC not 0000, purpose code 20 or 15	INR 5 lakh	not applicable
Jewellery	5944	INR 2 lakh	INR 6 lakh
FX-Retail via BBPS	6012	INR 5 lakh	INR 5 lakh
Digital account opening, term deposits	7410	INR 5 lakh	INR 5 lakh
Digital account opening, initial funding	7409	INR 2 lakh	INR 2 lakh

Two adjacent products carry their own ceilings. **UPI Lite** is an on-device wallet designed for payments **below INR 1,000**, processed without a UPI PIN and without a round trip to the bank's core banking system; it exists to take load off issuers' cores and to keep passbooks uncluttered, and the balance behaves like cash in that losing the device without disabling the wallet means going to the bank for a refund. **UPI Circle**, which lets a primary user delegate spending authority to a secondary user or to a trusted device, carries limits of **INR 5,000 per transaction per device, INR 15,000 per month per device, INR 2,000 in the first 24 hours after linking, and a maximum of five linked secondary devices or software agents**, per NPCI's published product page at the time of writing.

All of these figures are accurate at the time of writing and all of them have moved before. Read the current circular.

■ Why it reached these volumes

It is tempting to attribute UPI's scale to India's population, but population explains the ceiling, not the trajectory. Several deliberate design choices did the work, and they are transferable.

Interoperability was mandatory from the first day. Any app could pay any bank account. There was never a period in which a merchant had to accept four different wallets, and there was never a network effect that a single incumbent could capture. Compare the closed-wallet era that immediately preceded it, in which a customer holding one wallet simply could not pay a merchant who accepted another.

The app layer was separated from the licensed layer. Firms with no balance sheet could compete on user experience alone. That created ferocious product competition without creating prudential risk, because customer money never rested outside a bank.

Authentication was made cheap. Device binding plus PIN removed the OTP round trip, which removes both latency and a large class of failures. A payment that completes in under two seconds gets used for a forty-rupee purchase. A payment that takes twenty seconds and fails one time in twenty does not.

Addressing was abstracted. The VPA removed both the friction of typing an account number and the risk of publishing one. The mandatory name-return closed most of the mis-payment gap that abstraction would otherwise have opened.

Merchant acceptance was made nearly free, at both ends. Zero MDR removed the merchant's financial reason to refuse, and the printed QR removed the merchant's capital cost to accept. Those two together are why acceptance reached the informal economy, which no card scheme has ever managed anywhere.

The prerequisites were already in place. Bank accounts from the Jan Dhan financial inclusion programme, identity from Aadhaar, and mobile connectivity at some of the lowest data prices in the world gave UPI an addressable base it did not have to create. Demonetisation in November 2016, seven months after launch, provided an accelerant that no product plan could have arranged.

The specification kept moving. UPI 2.0 in August 2018 added linked overdraft accounts, one-time mandates, invoice-in-the-inbox and signed intent and QR. Later layers added AutoPay, UPI Lite, UPI 123Pay for feature phones, RuPay credit card on UPI, credit line on UPI, UPI Circle and IPO applications on UPI. Each layer reused the same switch, the same addressing and the same authentication, which is why each shipped in months rather than years.

UPI 123Pay deserves a specific mention because it is the clearest evidence that the design was about reach rather than about smartphones. It offers feature-phone users four routes: calling an interactive voice response number, an app running on the feature phone itself, a missed-call-based flow, and proximity sound-based payments. None of these requires a data connection in the ordinary sense.

The results, in NPCI and government figures published on **30 April 2026** to mark ten years of the system: annual volume grew from **2 crore transactions in FY 2016-17 to 24,161.69 crore in FY 2025-26**, an increase of roughly twelve thousand times; annual value in FY 2025-26 exceeded **INR 314 lakh crore**; the daily average through 2025 was about **60 crore transactions**, rising past 66 crore; person-to-merchant payments made up **63 per cent of volume** while person-to-person made up **71 per cent of value**; and **86 per cent of merchant payments were below INR 500**. The International Monetary Fund, in a June 2025 report, recognised UPI as the world's largest real-time payment system by transaction volume, at roughly **49 per cent of global real-time payment volume**.

■ The honest caveats

Concentration is real. Two apps run about four fifths of the volume, the rule says 30 per cent, and the deadline has moved three times. An outage at one of those apps, or at one of the two dominant sponsor banks, degrades payments for a very large fraction of the country at once. NPCI publishes uptime and downtime statistics per member precisely because this is the live risk.

Economics are unresolved. Zero MDR made acceptance universal and left the ecosystem dependent on a government incentive scheme that has consistently funded only a fraction of industry cost. The August 2026 amendment to Section 10A is the first legislative acknowledgement that this cannot continue indefinitely. Whatever emerges will be a compromise between the merchant reach that free acceptance bought and the revenue that maintaining the rail requires.

Fraud has migrated rather than disappeared. The collect request was closed off; social engineering, mule accounts and fake customer-support numbers were not. NPCI's FY 2026-27 circular OC No. 234 on safeguarding user information in UPI, and the revision of turnaround times for fraud and wrong-credit chargebacks in URCS, are the current front line.

And the system is not finished. Circulars in the current financial year cover a limit management member portal, on-device biometric authentication with an enhanced transaction ceiling, numeric UPI IDs, and international QR acceptance. Anything in this chapter with a number attached should be checked against NPCI's circular list before it is relied upon in production. That is not a weakness of the writing. It is the nature of a rail that is still being built at the same time as it carries eight thousand payments a second.

42.98 Common wrong ideas

Wrong: UPI is an instant payment system end to end. **Right:** the customer experience is instant and the interbank plumbing is deferred net settlement, in ten cycles a day between 09:00 and 21:00 since 3 November 2025, and a practitioner who confuses the two will size their liquidity wrong.

Wrong: NPCI holds the money and shifts it between banks. **Right:** NPCI is a switch that routes messages, applies rules, keeps score and produces settlement files; the money moves across banks' current accounts at the Reserve Bank of India.

Wrong: India moves its money on UPI. **Right:** UPI is 85.5 per cent of digital payment volume and 9.5 per cent of value, while RTGS is 0.1 per cent of volume and 68.6 per cent of value; that claim describes counts, not sums.

Wrong: PhonePe, Google Pay and Paytm are banks. **Right:** they are third-party application providers riding on a sponsor payment service provider bank's licence, which is why the right question during an outage is which sponsor bank is down.

Wrong: a virtual payment address keeps the account holder anonymous. **Right:** it hides the account number from the payer, and nothing at all from NPCI or the payee's bank, and the registered name is deliberately returned to the payer before confirmation.

Wrong: UPI is free. **Right:** zero merchant discount rate is a statutory policy subsidised by member banks, NPCI switching fees and a government incentive scheme, three flows already sit outside the perimeter, and the statutory bar was removed on 6 August 2026.

Wrong: the transaction ID shown in the app is the reference to quote when chasing a payment. **Right:** the up-to-35-character transaction ID is the app's handle; the twelve-digit retrieval reference number, commonly called the UTR, is the number both banks can search on.

Wrong: if money leaves the account and does not arrive, the customer must chase it and hope. **Right:** auto-reversal is due latest on T+1 for a funds transfer and within T+5 at a merchant location, with INR 100 per day of delay beyond that, and the compensation is not discretionary.

Wrong: a collect request is a normal way to ask an individual for money. **Right:** person-to-person collect was progressively restricted and then discontinued entirely from 1 October 2025, because a request for money looks like an offer of money to an inexperienced user.

Wrong: mandate executions and API calls can be scheduled whenever it suits the merchant. **Right:** since 1 August 2025 NPCI rate-limits UPI APIs, caps balance enquiries at fifty per app per day, and pushes non-customer-initiated calls out of the peak windows of 10:00 to 13:00 and 17:00 to 21:30.

42.99 Chapter summary in 20 lines

1. UPI is India's retail payment switch, operated by NPCI, a not-for-profit company promoted by the Reserve Bank of India and the Indian Banks' Association, live since 11 April 2016.
2. NPCI holds no money: it keeps the address book, routes the messages, applies the rules and produces the settlement files.
3. A payment is addressed by virtual payment address, UPI number, QR code or raw account number plus IFSC, and the address is resolved to an account by NPCI's central mapper.
4. Before any debit, the payee's registered name is returned and displayed, which is why mis-payment is rare and why a UPI address cannot be treated as private.
5. Two-factor authentication is satisfied by a bound device plus a PIN, with no OTP round trip, which NPCI calls single-click two-factor authentication.
6. The PIN is encrypted inside NPCI's Common Library and decrypted only in the remitter bank's hardware security module, so neither the app nor the PSP ever sees it.
7. A push payment is two synchronous legs: the switch performs the debit at the remitter bank and only on success performs the credit at the beneficiary bank.
8. Failures come in three shapes: a clean decline, a credit failure that reverses within minutes, and a timeout that leaves the transaction in deemed status until reconciliation resolves it.
9. Settlement is deferred and net across members' current accounts at the Reserve Bank of India, in ten cycles a day between 09:00 and 21:00 since 3 November 2025, with two separate cycles for disputes.
10. Between the customer's instant result and settlement the beneficiary bank carries credit risk on the remitter bank, which is why NPCI runs net debit caps and collateral requirements.
11. The app layer was separated from the licensed layer, so firms with no balance sheet compete on user experience while customer money never rests outside a bank.
12. That separation also produced extreme concentration: two apps run about four fifths of volume against a 30 per cent cap whose deadline has moved three times and now stands at 31 December 2026.
13. In July 2026 UPI carried 23.66 billion transactions worth about INR 29.9 lakh crore across 741 banks, roughly 763 million payments a day.
14. The average payment was INR 1,263, about eleven pounds, which is the number that explains everything: this is a rail for chai and vegetables, not for flats.
15. Zero merchant discount rate, statutory since 1 January 2020, is why a tea stall accepts a laminated sticker that cost the price of printing rather than a terminal with a rental and a percentage.
16. On 6 August 2026 the Lok Sabha passed an amendment removing the statutory bar on charging, so the free-to-use position is now a policy choice rather than a law of nature.
17. When money vanishes, the twelve-digit reference is the only thing that lets a bank find it, and the escalation ladder runs app, PSP bank, NPCI's UPI Help portal, then the Reserve Bank Integrated Ombudsman Scheme.
18. UPI AutoPay is a genuine e-mandate with pre-debit notification, customer-side control and frequency options from daily to as-presented, subject to peak-hour rate limits.
19. The scale came from mandatory interoperability, cheap authentication, abstracted addressing, near-free acceptance at both ends, an existing base of accounts, identity and cheap data, and a specification that never stopped moving.
20. The unresolved problems are concentration, the economics of zero MDR, and fraud that migrated to social engineering rather than disappearing, which is why every figure here should be checked against NPCI's current circular before production use.

Sources: NPCI (UPI product statistics for July 2026, About UPI, UPI AutoPay, UPI Lite, UPI Circle, UPI

123PAY, UPI circular list, and circular NPCI/UIP/2025-26/OC/185B of 28 August 2025); Reserve Bank of India (circular RBI/2019-20/67 of 20 September 2019 on turn around time and customer compensation, Annual Report 2024-25, and the half-yearly Payment Systems Report covering H2 2025); Press Information Bureau releases of 30 April 2026 and on zero MDR; the Payment and Settlement Systems Act 2007 as amended, and reporting on the Taxation and Other Laws (Amendment) Bill, 2026 passed by the Lok Sabha on 6 August 2026.

Instant Rails Elsewhere

43.0 What this chapter gives you

1. You will be able to explain why Brazil, the United States and the euro area reached three different outcomes from the same forty-year-old technology.
2. You will be able to name the design choices that predict adoption, central operation or compulsion, mandated reach, a rule-based price cap and memorable addressing, and test any new system against them.
3. You will be able to separate Pix the scheme from SPI the settlement system and DICT the directory, which everyday usage runs together.
4. You will be able to explain why a Pix payment and a Faster Payment look identical to a customer and put risk and liquidity in completely different places for a treasurer.
5. You will be able to argue that the United States has more instant payment infrastructure than most countries and less instant payment, and say exactly what it is missing.
6. You will be able to quote the three Same Day ACH deadlines and say what happens to a file that misses all of them.
7. You will be able to state what the Instant Payments Regulation actually obliges, and the date each obligation bites for euro-area banks, non-euro-area banks and payment institutions.
8. You will be able to read the TIPS settlement dip of October 2025 as the expected cost of a well-executed mandate rather than as a failure.
9. You will be able to say who is paying for a “free” instant payment in each of the three jurisdictions, and whether they will still be paying in five years.
10. You will be able to diagnose what Britain has, what it lacks, and why the missing piece is alias addressing rather than speed.

On the same Tuesday morning, three people pay the same small amount of money to somebody they can see.

In São Paulo, a woman buys two kilos of tomatoes at a street market. The stallholder turns a laminated card round so she can photograph the square black-and-white pattern on it. She taps twice, the stallholder’s phone chimes, and the money is his, finally, irreversibly, at eleven o’clock in the morning. Neither of them paid anything to move it.

In Columbus, Ohio, a man pays his window cleaner. He has the window cleaner’s bank routing number and account number written on a slip of paper, and he types them into his bank’s website. It tells him the payment will arrive on Thursday. He also has the option to pay a fee to make it arrive today, and there is a third option, which is to write a cheque, and this is 2026.

In Rotterdam, a woman pays her plumber from her phone. Before she confirms, the screen shows her the name registered to the account she has typed, and it matches the name on the invoice. She presses send. The plumber's bank credits him within ten seconds, and by law her bank may not charge her more for this than it would have charged for a payment that took a day.

Three countries, three answers, one question. The question is the same everywhere and it is not a technical question. Moving a number from one ledger to another has been technically trivial for forty years. The question is who decides, who must join, and who pays. Brazil answered it one way, the United States another, and the European Union a third, and the differences between the three explain almost everything about why instant payments are ubiquitous in one place, legislated into existence in another, and still an optional extra in the third.

All figures, dates, limits and cut-off times in this chapter are accurate at the time of writing, August 2026. Instant payment systems are among the fastest-moving infrastructures in finance, and several of the numbers below changed within the last eighteen months.

The plain version

Imagine three towns, each with a river running through it and no bridge.

In the first town, the council decides that a bridge is not something you leave to chance. It builds the bridge itself, out of public money. Then it passes a rule: every haulage company in town above a certain size must use the bridge, and none of them may charge ordinary residents a penny to walk across. Businesses pay a small toll, but people do not. The bridge is open every hour of every day.

In the second town, the council does not build anything. It says that bridges are a matter for private enterprise. Two consortia eventually build two bridges, some distance apart. Neither bridge connects to every road. A haulier who wants to reach every address in town must have an arrangement with both, and many hauliers have an arrangement with neither, because the old ferry still runs and the ferry is cheap and everybody is used to the ferry. Some residents cross the river every day and have never used either bridge.

In the third town, the council also does not build the bridge. But it does something else. It passes a law: within eighteen months, every haulage company must be able to receive goods from every other haulage company across the river, within ten seconds, and may not charge more for the fast crossing than for the slow one. It then says: here is a piece of publicly owned road at the centre of the town that all the bridges must connect to. How you build your end is your business. That the connection exists is not negotiable.

The first town is Brazil. The second is the United States. The third is the European Union. And the striking thing, when you look at the traffic figures, is how well the outcome matches the design.

■ Brazil, and what happens when the state builds the road

Brazil's system is called Pix. It is run by the country's central bank, the Banco Central do Brasil, and it opened in November 2020.

Two things about Pix matter more than everything else combined.

The first is that if you are a bank or a payment company in Brazil with more than five hundred thousand active customer accounts, you are not asked whether you would like to join. You are required to. There was no launch period during which some banks had it and others did not, no competitive advantage for early movers, no customer ringing up to ask why her bank did not support it. From very close to the beginning, if you had a bank account in Brazil, you could send and receive.

The second is that banks are legally forbidden from charging ordinary people to use it. Not discouraged. Forbidden, in a written rule, with exceptions so narrow that they prove the point. A shop pays to receive. A person does not pay to send.

Then there is a third thing, less dramatic but arguably the cleverest. To pay someone in Brazil you do not need their bank details. You need a *key*: their mobile number, their email address, their national tax number, or a long random string their bank generated for them. The central bank runs a directory that maps every key to an account. You type the key, the directory returns the account, and it also returns the name on that account so you can check you are paying who you think you are paying. Businesses print a square barcode instead, which contains the same thing.

The numbers that follow from those three decisions are difficult to argue with. In 2021, its first full year, Pix carried 9.4 billion transactions. In 2024 it carried 63 billion, moving BRL 26.4 trillion, which the central bank's own deputy governor put at around two and a half times Brazil's entire economic output for that year. At the five-year mark, in November 2025, close to 170 million people were using it, in a country of roughly 215 million, along with more than 20 million businesses. As at 15 August 2026, 889 institutions took part.

For a shop, the cost of taking a Pix payment has been measured at about 0.22 per cent of the sale, against roughly 1.1 per cent for a debit card and 2.2 per cent for a credit card. That is not a marginal saving. That is the difference between a business model working and not working.

■ The United States, and what happens when nobody is in charge

The United States has the largest economy in the world and, for most of the last decade, some of the slowest ordinary payments in the developed world. The reason is not technical incompetence. It is that the United States has roughly nine thousand banks and credit unions, no legal power to make them all join anything, and a payment system that was already good enough that nobody had to.

The old system is called ACH, for Automated Clearing House. It is a batch system, in the exact sense of the Bacs cycle familiar to British readers: files of instructions are collected, sorted, exchanged and applied, with settlement in between. It carries wages, benefits, tax refunds, mortgage payments and most business-to-business invoicing. It is enormous. In 2025 it moved 35.2 billion payments worth USD 93 trillion.

It also got faster. Since 2016 there has been a same-day option, with three deadlines during the working day. Submit a file by 10:30 in the morning, New York time, and it settles at one o'clock the same afternoon. Submit by 2:45 and it settles at five. Submit by 4:45 and it settles at six. Miss all three and you are on tomorrow. That is not instant, but for a great many purposes it is close enough, and it costs a fraction of a wire transfer, and this is precisely why the United States was slow to build anything faster. The pressure was never quite unbearable.

When faster things were built, they were built twice. A group of the largest banks built one, called RTP, which went live in 2017. The Federal Reserve built another, called FedNow, which went live in July 2023. Both run around the clock. Both are genuinely instant. And there are two of them, which means that a bank must join both to be sure it can reach everyone, and many banks have joined neither.

The volumes tell the story plainly. In its first full year, 2024, FedNow settled 1.5 million payments. In 2025 it settled 8.4 million. That is fast growth from a small base, and it is a rounding error next to 35.2 billion ACH payments. Meanwhile, Americans do in fact send each other money instantly, in very large quantities, using a service called Zelle, which is not a payment rail at all but a directory and a set of rules sitting on top of the banks. In 2024 Zelle carried 3.6 billion payments worth more than a

trillion dollars. It solved the addressing problem, which is the problem people actually experience, and left the plumbing alone.

■ Europe, and what happens when you legislate

The euro area had the technology from 2017 and almost nobody used it. A pan-European instant payment scheme existed, banks could join it, and joining was voluntary, so a large proportion of banks did not. Money crossing Europe in ten seconds was possible in theory and unavailable in practice, because the bank at the other end had not signed up.

So the European Union passed a law. Under the Instant Payments Regulation, from 9 January 2025 every bank in the euro area had to be able to *receive* an instant euro payment. From 9 October 2025 every bank had to be able to *send* one, and had to offer a free check of the payee's name before you press the button, and was forbidden to charge more for the instant version than for the slow version.

The effect was immediate and visible in the settlement figures. The Eurosystem's instant settlement platform, TIPS, handled 1.35 billion transactions in 2024. In 2025 it handled 2.47 billion, a rise of 82.5 per cent, with the sharpest growth in the closing months of the year, exactly as the sending obligation bit. On 1 December 2025 it settled 5,759,570 payments in a single day.

Note what Europe did *not* do. It did not build a consumer product; there is no European Pix app. It legislated the two things Brazil got by other means, universal reach and price parity, and let the market do the rest.

■ The pattern

Put the three next to each other and the pattern is not subtle.

	Brazil	United States	Euro area
Who runs the instant rail	Central bank	Federal Reserve and a private bank consortium, in parallel	Eurosystem platform, private schemes on top
Must banks join	Yes, above a size threshold	No	Effectively yes, by law, since 2025
Can consumers be charged	No, by regulation	Yes	Not more than for a slow payment
How you address a payment	Phone, email, tax number or random key	Routing number and account number	IBAN, with a name check
Adoption	Near-universal	Growing from a small base	Rising steeply since the deadlines

Central operation, or at least central compulsion. Mandated reach. A cap on what the customer can be charged. Addressing a human being can remember. Where those four are present, instant payments take over. Where any of them is missing, they do not.

Where the plain version stops being true

The story above is true in outline and misleading in four specific places, and each of them matters if you are building anything.

■ “Instant” describes the customer’s experience, not the banks’

When a Pix payment lands, the two banks involved have genuinely settled with each other, individually, transaction by transaction, in central bank money, at that instant. When a FedNow payment lands, the same is true. That is unusual.

For a great many fast payment systems it is not true at all. In the United Kingdom, a Faster Payment reaches the recipient in seconds, but the obligation between the two banks is settled afterwards, on a net basis, across accounts at the Bank of England, against balances the banks have prefunded. India’s UPI, which moves more transactions than any system on earth, likewise gives the user an instant result over a deferred net settlement arrangement underneath. The customer cannot tell the difference and should not have to. The bank’s treasurer very much can, because the two models put risk and liquidity in completely different places. Real-time gross settlement means every payment consumes central bank money at the moment it is made, all night and all weekend, which means a bank must fund an account it cannot easily top up out of hours. Deferred net settlement means the banks owe each other for a while, which is cheaper and introduces exposure between them.

So “instant” is a statement about the payee’s balance, not about the plumbing. Two systems that look identical to a customer can be entirely different animals.

■ Public ownership correlates with adoption; it does not cause it

It is tempting to read Brazil’s numbers and conclude that the state should build the rail. The evidence for that is real but weaker than the story suggests.

The Bank for International Settlements ran the comparison properly: thirteen jurisdictions, monthly data from April 2001 to December 2023, just over a thousand country-month observations, controlling for country and year. Public ownership of a fast payment system is associated with an increase of about 1.8 per cent in transactions per head. That is a genuine effect and it is the *smallest* of the four they found. Allowing non-bank payment firms to participate is associated with 3.5 per cent. Each additional use case is worth about 2 per cent, and so is each additional cross-border connection.

And there are counter-examples in plain sight. Sweden’s Swish, one of the most thoroughly adopted instant payment products anywhere, is owned and operated entirely by private banks. India’s UPI is governed jointly by private banks and the central bank through a purpose-built body. What Brazil got from public ownership was not superior engineering. It was the political authority to compel participation and to forbid consumer fees, and those are the levers that moved. A private consortium with the same two powers would have had much the same result. A central bank without them would not.

■ “Free” means somebody else is paying

Pix is free to individuals because a rule says it must be. It is not free to operate. The central bank charges participating institutions for the settlement infrastructure and for the addressing directory. Merchants pay to receive. The 0.22 per cent that a Brazilian shop pays is small next to a card fee, but it is not zero, and the entire consumer-facing economics of the system rest on it.

The same is true everywhere. The Federal Reserve charges banks 4.5 cents to originate a FedNow customer credit transfer, a dollar for a liquidity transfer and a cent for a request for payment, and in 2026 waived the twenty-five dollar monthly participation fee and rebated the first 2,500 credit transfers a month, which is a subsidy with a purpose. The Eurosystem charges 0.2 of a euro cent per instant payment, split between the sending and receiving bank. In Europe, the Instant Payments Regulation does not say instant payments must be free. It says they may not cost the customer *more* than a non-

instant one, which in most euro-area retail banking means free in practice and is a different legal statement.

If you are modelling a business on a rail because “it’s free”, find out who is paying and whether they are still going to be paying in five years.

■ The United States is not behind; it is differently arranged

The comparison that makes the United States look primitive is unit volume on instant rails, and by that measure it is far behind Brazil and India. But the United States has more instant payment infrastructure than most countries, not less. It has two independent real-time networks, one public and one private, both settling in central bank money, both with transaction ceilings far above anything in Europe or the United Kingdom: RTP raised its per-transaction limit to USD 10 million on 9 February 2025, and FedNow raised its network limit from USD 1 million to USD 10 million on 12 November 2025. British readers should sit with that. The Faster Payments scheme maximum is GBP 1 million.

What the United States lacks is not rails. It lacks the three things Brazil legislated: compulsory reach, a price cap at the consumer end, and one directory. It also has something Brazil did not have, which is a card industry so profitable and so entrenched that account-to-account payment at the point of sale has no natural sponsor. Nobody in the American payments industry makes more money when a customer pays a shop directly from their bank account. In Brazil, the central bank did, in the sense that it was pursuing a public objective rather than a margin.

And the volume comparison flatters the wrong number in any case. Same Day ACH alone carried 1.4 billion payments worth USD 3.9 trillion in 2025, up 16.7 per cent on the year, and Zelle carried 2 billion payments worth nearly USD 600 billion in the first half of 2025 alone. Americans are not waiting three days to pay each other. They are doing it over three different systems, none of which reaches everybody.

The technical version

■ Brazil: Pix, SPI and DICT

Pix is a payment scheme owned by the Banco Central do Brasil, established by BCB Resolution 1 of 12 August 2020 and launched in November 2020. It is important to separate three things that colloquial usage runs together.

Pix is the scheme: the rulebook, the branding, the user experience standards, the participation criteria.

SPI, the *Sistema de Pagamentos Instantâneos*, is the settlement infrastructure. It is a real-time gross settlement system built and operated by the BCB specifically for instant payments, running 24 hours a day, every day of the year. Direct participants hold dedicated Instant Payment accounts, known as PI accounts, at the central bank. Every Pix transaction settles individually across those accounts in central bank money at the moment it is made. Brazil already had a conventional RTGS system, the STR, when Pix was designed; the decision to build a separate platform rather than extend the existing one was a deliberate one about service levels and technology. Legally, SPI was built for Pix, and other payment schemes cannot connect to it for instant settlement.

DICT, the *Diretório de Identificadores de Contas Transacionais*, is the addressing directory. It maps aliases, known as Pix keys, to transaction accounts. Four key types are permitted: the CPF or CNPJ tax identifier, a mobile telephone number, an email address, and a randomly generated key, the EVP. Participants that are direct participants in SPI must access DICT directly.

Participation criteria are set out publicly. An institution licensed by the BCB that holds more than 500,000 active customer accounts, counting demand, savings and prepaid payment accounts together, is obliged to join, and once it crosses that threshold it has 90 days to apply for membership in the transaction account provider category. Smaller institutions may join voluntarily; institutions outside the regulatory perimeter may participate indirectly through a direct participant, subject to a contract, demonstrated technical capacity and a minimum equity requirement.

Pricing to end users is governed by BCB Resolution 19 of 1 October 2020, which is worth quoting because its structure is the model other jurisdictions keep reaching for. Article 3 prohibits an account-holding institution from charging a natural person customer, including individual entrepreneurs, for the remittance of funds for transfer or purchase, or for the receipt of funds for the purpose of transfer. The single paragraph carves out transactions conducted through in-person channels or the call centre where an electronic means was available. Article 4 permits charging a natural person for the receipt of funds *for the purpose of purchase*, which is how a sole trader taking payment for goods ends up paying, and permits charging legal entities for both sending and receiving.

The functional set has widened steadily: billing (Pix Cobrança), cash withdrawal at merchants (Pix Saque), cash back (Pix Troco), scheduled payments (Pix Agendado), contactless payment over NFC (Pix por Aproximação, launched 28 February 2025) and direct debit (Pix Automático, launched June 2025). Pix Automático is the structurally interesting one: any participant may offer it to corporate customers regardless of where the payer's account is held, which breaks the historic pattern in which recurring collection was the preserve of a handful of large banks that the biller had to bank with. Fraud handling runs through a Special Return Mechanism, the MED, which permits funds to be blocked and returned despite transactions being nominally irrevocable, with a further revision announced by the BCB.

Adoption, on the central bank's figures: 9.4 billion transactions worth around BRL 5 trillion in 2021; 63 billion transactions worth BRL 26.4 trillion in 2024; close to 170 million users and more than 20 million businesses at the five-year mark in November 2025. As at 15 August 2026 the BCB listed 889 participants, of which 560 were credit cooperatives and 173 payment institutions, which is itself a comment on how wide the participation rules are drawn.

■ The United States: ACH, Same Day ACH, FedNow and RTP

The ACH network is governed by the Nacha Operating Rules and cleared by two operators: FedACH, run by the Federal Reserve, and the Electronic Payments Network, run by The Clearing House. It carries credits and debits in batch, with deferred net settlement across Federal Reserve accounts.

In 2025 the network processed 35.2 billion payments valued at USD 93 trillion, the thirteenth consecutive year in which total value rose by at least a trillion dollars. December 2025 was the highest-volume month recorded, at 3.22 billion payments, and November 2025 set an average daily record of 151 million payments.

Same Day ACH is an expedited service within the same network, not a separate rail. The FedACH processing schedule, effective 12 September 2022, defines three same-day windows.

Transmission deadline	Target distribution	Settlement
10:30 ET	12:00 ET	13:00 ET, current day
14:45 ET	16:00 ET	17:00 ET, current day
16:45 ET	17:30 ET	18:00 ET, current day

Items not eligible for same-day settlement, including anything submitted in the later overnight windows, settle at 08:30 ET on the next banking day. The per-transaction limit for same-day eligibility is USD 1,000,000, effective 18 March 2022. Items above that limit submitted in a same-day window are not rejected; the operators assign them to the next available window for next-day settlement. Nacha has approved an increase of that limit to USD 10,000,000 with an effective date of 17 September 2027. In 2025, Same Day ACH carried 1.4 billion payments worth USD 3.9 trillion, up 16.7 per cent by volume and 21.4 per cent by value, averaging 5.8 million payments per day across the year and 7.8 million per day in December.

The FedNow Service launched in July 2023. It settles customer credit transfers individually in real time across participants' Federal Reserve accounts, using ISO 20022 messages: pacs.008 for the credit transfer, pacs.004 for the return, pacs.009 for a liquidity management transfer, pain.013 for a request for payment.

Operating hours are worth stating precisely, because “24/7” conceals a business-day convention. Fed-Now operates on a 24-hour funds transfer business day every day of the week including weekends and Federal Reserve holidays. The business day rolls over at approximately 19:01 ET each calendar day, with the new day beginning immediately as the previous one ends and no interruption in processing. Liquidity Management Transfers, which let a participant move funds between its Federal Reserve account and the balances supporting instant payments, run 19:00 to 07:00 ET on weekdays and continuously at weekends and on Federal Reserve holidays.

Value limits: the default per-transaction limit is USD 100,000, which a participant may adjust, and the network limit rose from USD 1 million to USD 10 million effective 12 November 2025.

The 2026 fee schedule, effective 1 January 2026:

Item	Fee
Customer credit transfer origination (pacs.008)	USD 0.045 per item
Customer credit transfer origination, on-us	USD 0.000
Return origination (pacs.004)	USD 0.045 per item
Liquidity Management Transfer (pacs.009)	USD 1.00 per item
Request for Payment (pain.013)	USD 0.01 per item
Participation fee	USD 25.00 per routing number per month, discounted to USD 0.00 in 2026
Origination discount	Minus USD 0.045 per item for up to 2,500 credit transfers per month in 2026

Settled volumes, from the Federal Reserve's own statistics, last updated 6 July 2026:

Period	Settled payments	Value (USD)	Average value (USD)	Average daily volume
2023 (part year)	47,262	18,433,363	390	285
2024	1,505,250	38,196,907,431	25,376	4,113
2025	8,413,402	853,411,108,511	101,435	23,050
2026 Q1	2,728,510	271,252,920,121	99,414	30,317
2026 Q2	4,997,811	274,663,803,964	54,957	54,921

The 2026 second quarter is the interesting line: volume up 83.2 per cent on the quarter while value was almost flat, and average payment value more than halved. That is the signature of a wholesale-ish

network beginning to carry retail traffic. Participation was reported by the Federal Reserve at more than 1,500 financial institutions across all fifty states in autumn 2025, with 97 more institutions and four certified service providers joining in the first quarter of 2026.

The RTP network, operated by The Clearing House, launched in 2017 and predates FedNow by six years. It settles from a prefunded joint balance held at the Federal Reserve. Its per-transaction limit rose to USD 10 million effective 9 February 2025, a tenfold increase on the previous USD 1 million cap. At the end of 2024 The Clearing House reported the network averaging about one million payments a day, with some 285,000 businesses a month originating payments through participating institutions; by mid-2025 it reported more than 1,000 banks and credit unions live, a 51 per cent increase year on year.

Zelle is neither of these. It is an overlay operated by Early Warning Services, owned by a group of large banks, which resolves an email address or mobile number to an account and moves the resulting instruction over whichever underlying rail the participating banks use. It is the closest thing the United States has to a DICT. It carried 3.6 billion payments and more than USD 1 trillion in 2024, and 2 billion payments worth nearly USD 600 billion in the first half of 2025.

Why the United States is unusual, stated plainly. There is no legal power to compel roughly nine thousand depository institutions to join a payment system. There are two competing instant networks, so reach is a commercial question rather than a given. Addressing is by routing transit number and account number, credentials which are both hard to remember and, being reusable, unsafe to publish. Consumer pricing is unregulated. Cheques remain in use. And the existing batch rail is cheap, extremely reliable and, since 2016, same-day capable, which removed most of the urgency. The Bank for International Settlements notes a further constraint: FedNow, like Mexico's CoDi and Sweden's Swish, is open only to banks, in some cases for legislative reasons, and non-bank participation is one of the strongest correlates of adoption they measured.

■ Europe: T2, TIPS and the Instant Payments Regulation

Terminology first, because the old name persists and is now wrong. TARGET2 was replaced on 20 March 2023 by **T2**, following consolidation with the securities settlement platform T2S. TARGET Services today comprise T2 for high-value and wholesale euro payments, T2S for securities, TIPS for instant payments and, since 16 June 2025, the Eurosystem Collateral Management System.

T2 is the euro area's real-time gross settlement system: the direct analogue of CHAPS within the Bank of England's RTGS, at continental scale. In 2025 it settled an average of 431,067 payments per day in euro, worth EUR 1,932.8 billion, across 255 operating days, for a yearly total of 111,910,103 transactions worth EUR 492,859.0 billion. That total exceeded the previous peak of 107,999,982 transactions set in 2024, and the value grew 6.3 per cent, ahead of euro-area nominal GDP growth of 3.8 per cent.

The business day matters for anyone building payment operations against it. The real-time settlement window concludes at 18:00 CET and may be interrupted by an optional maintenance window from 03:00 to 05:00 CET. The cut-off for customer payments is 17:00 CET; all other categories may settle until 18:00 CET. In 2025, 99.7 per cent of volume and 96.3 per cent of value had settled by the 17:00 customer cut-off, with the remaining payments in the final hour typically few in number and very large in value, as banks square positions and refinance on the money market.

TIPS, TARGET Instant Payment Settlement, launched in November 2018. It settles instant credit transfers in central bank money, 24 hours a day, 365 days a year. Participants hold a TIPS Dedicated Cash Account, a DCA, funded from their RTGS position; the alternative structure is an Ancillary System Technical Account, an ASTA, used where a clearing mechanism sits between the participant and the

platform. Pricing is EUR 0.002 per instant payment, split equally between originator and beneficiary at EUR 0.001 each.

TIPS is a settlement platform, not a scheme. The scheme is **SCT Inst**, the SEPA Instant Credit Transfer scheme of the European Payments Council, launched in November 2017. Participation in TIPS, directly or indirectly, is mandatory for payment service providers that have signed up to SCT Inst and are reachable via T2. Until 9 October 2025 the scheme-level default maximum amount was EUR 100,000; that ceiling has been removed.

The **Instant Payments Regulation**, Regulation (EU) 2024/886, was adopted on 13 March 2024. It amends the SEPA Regulation, the cross-border payments regulation, the Settlement Finality Directive and PSD2. Its four substantive obligations are that a payment service provider offering ordinary credit transfers must also offer instant ones (Article 5a); that charges for instant credit transfers may not exceed charges for other credit transfers of corresponding type (Article 5b); that providers must offer a free verification of payee service before the payer authorises, returning a match, close match, no match or other result (Article 5c); and that providers must screen their own customer base against targeted financial sanctions at least daily rather than screening each transaction in flight (Article 5d).

The implementation timetable is staggered and precise.

Requirement	Who	Deadline
Receive instant payments	Euro area Member States	9 January 2025
Receive instant payments	Non-euro area Member States	9 January 2027
Receive instant payments	EMIs and payment institutions, euro and non-euro area	9 April 2027
Send instant payments	Euro area Member States	9 October 2025
Send instant payments	Non-euro area Member States	9 July 2027
Send instant payments	EMIs and payment institutions, euro area	9 April 2027
Send instant payments	EMIs and payment institutions, non-euro area	9 July 2027
Send outside business hours from accounts in national currency	Non-euro area Member States	9 June 2028
Equality of charges	Euro area Member States	9 January 2025
Equality of charges	Non-euro area Member States	9 January 2027
Verification of payee	Euro area Member States	9 October 2025
Verification of payee	Non-euro area Member States	9 July 2027

The Eurosystem offers a verification of payee service built on solutions developed by Banco de Portugal and Latvijas Banka, conforming to the European Payments Council's VoP scheme rulebook, effective 5 October 2025. Separately, from October 2025, non-bank payment service providers meeting the conditions in the TARGET Guideline may access TARGET Services directly, including T2 and TIPS, a change with the same intent as Brazil's wide participation categories.

The effect on volumes was sharp and unambiguous. TIPS transactions rose from 1.35 billion in 2024 to 2.47 billion in 2025, growth of 82.5 per cent. Euro-denominated TIPS activity averaged 2,735,053 payments per day in 2025, worth EUR 2.2 billion, up 110.8 per cent from a daily average of 1,297,241 in 2024. December 2025 averaged 3,845,376 payments a day and the maximum single day was 5,759,570 on 1 December 2025. Germany overtook Italy as the largest contributor by volume, rising from 15.2 to 19.8 per cent of the total, and contributed 27.4 per cent of value.

The stress shows in the settlement ratio. Across 2025, 98.0 per cent of euro traffic settled by volume and 95.8 per cent by value. The value ratio fell to 93.1 per cent in October, coinciding with the sending obligation, as participants adapted to the shorter processing timeline, newly connected participants struggled with genuine 24/7 operation, and much larger payments arrived following the removal of the EUR 100,000 ceiling. By December it had recovered to 96.5 per cent. That is what a well-executed mandate looks like from the inside: a measurable two-month degradation, then recovery.

TARGET Services are also becoming multi-currency. The Danish krone joined T2 and TIPS in April 2025, with T2 settling an average of 11,425 krone payments per day worth DKK 1,641.9 billion in 2025, and the migration of Danish instant payments contributed materially to TIPS growth. Cross-currency capability was implemented in TIPS in June 2025, and the ECB is exploring interlinking TIPS with fast payment systems elsewhere.

■ What the successful systems have in common

Set Brazil, the euro area, India, Thailand and the United Kingdom side by side and five design choices recur wherever adoption is high. They are not all necessary, but the more of them you have, the steeper the curve.

Central operation or central compulsion. Someone with authority has to be able to say that this exists and everyone connects to it. In Brazil that authority is ownership of the scheme and the infrastructure. In the euro area it is a regulation with dates in it. What matters is not the flag on the building; it is that reachability stops being a commercial decision.

Mandated participation, and therefore guaranteed reach. This is the single largest determinant, and it is upstream of everything else. A payment system a customer cannot rely on reaching the other party is not a payment system; it is a feature. Brazil compelled institutions above 500,000 active accounts. Europe compelled everyone, in two stages. The United States compelled nobody, and the outcome is exactly what you would predict.

Low or zero cost at the consumer end, established by rule. Pix forbids charging individuals. The Instant Payments Regulation forbids charging more than for a slow transfer. Neither leaves it to competition, because competition in retail banking does not reliably drive a per-transaction fee to zero, particularly where the incumbent alternative pays the bank a share of interchange.

Simple addressing, backed by a directory, with a name check. A payment instrument that requires the payer to obtain and correctly transcribe an account credential has a ceiling on its use. Aliases remove that ceiling. Brazil built DICT. India built the virtual payment address. Europe legislated verification of payee, which is a name check rather than an alias directory, and is a weaker version of the same idea. The United States has aliases, but they live in a private overlay rather than in the rail.

A settlement account that is open when the rail is open. An instant system running at three in the morning on a Sunday needs participants to hold funded positions at the central bank outside banking hours, and needs a way to top them up. The Eurosystem gives participants TIPS DCAs. The Federal Reserve gives them Liquidity Management Transfers with defined overnight and weekend windows. The BCB gives direct participants remunerated PI account balances, liquidity lines, and a window after the close of the RTGS system to move funds across. Skip this and the system works beautifully until a large payment arrives on a bank holiday.

There is a sixth, less often stated: **one scheme, not several.** Two competing instant networks in one currency area do not double the reach; they halve the certainty. The United States is the demonstration.

■ What that predicts for other markets

For the United Kingdom, the diagnosis is uncomfortable, because Britain was early. Faster Payments arrived in 2008, before almost anyone, and being early froze in the design assumptions of 2008. The country has central operation, effective reach, and no consumer charge for retail Faster Payments in practice. What it does not have is alias addressing. A British payment is still addressed by six-digit sort code and eight-digit account number, exactly as it was in 1968, and Confirmation of Payee is a check on that credential rather than a replacement for it. Britain has the Europe answer to fraud without the Brazil answer to addressing. Open banking has partially filled the gap, but as an API layer that initiates payments over Faster Payments it inherits the addressing model rather than replacing it. Any future British improvement will have to solve identity resolution, and the scheme maximum of GBP 1 million per Faster Payments transaction will look increasingly conservative next to the USD 10 million now permitted on both American instant networks.

For the United States, expect convergence on reach without convergence on scheme. Both networks are growing, both have raised limits far beyond international norms, and the FedNow figures for the second quarter of 2026 show average payment value falling as volume climbs, which is retail adoption beginning. What will not happen without legislation is a single directory or a consumer price cap, so Zelle will continue to be the alias layer and card interchange will continue to fund the resistance at the point of sale.

For the euro area, the reach problem is solved and the product problem is not. TIPS settles the payments; it does not give a consumer a reason to use one. That gap is where private schemes layered on SCT Inst are competing, and their success or failure over the next few years will show whether legislated infrastructure is sufficient without a consumer-facing brand. Brazil's answer suggests it is not: Pix succeeded partly because the central bank specified minimum user experience requirements and a single recognisable name rather than leaving presentation to 889 institutions.

For markets that have not yet built, the BIS evidence gives a clear order of operations. Admit non-bank payment service providers, because that was the strongest single correlate they measured. Launch with several use cases rather than one, because each additional use case is worth about as much as public ownership. Plan cross-border connections early. And, though the data were not available to test it directly, mandate participation, because that is the mechanism by which Brazil obtained in one year what voluntary schemes in Europe failed to obtain in seven.

The counter-pressure is fraud, and it is the one thing every fast system has had to retrofit. Irrevocable, instant, push payments to an account you cannot inspect are an excellent instrument for a confidence trick, which is why Brazil built the MED into an otherwise irrevocable scheme, why Europe legislated verification of payee alongside the sending obligation rather than afterwards, and why the United Kingdom has operated a mandatory reimbursement regime for authorised push payment fraud since October 2024. Speed and reversibility trade against each other, and the trade cannot be designed away, only priced and allocated. Every jurisdiction that has built a fast rail has ended up building a slow, human, contested process alongside it for the payments that should not have happened. That process is not a failure of the design. It is part of the design, and the ones that planned for it from the start have had a considerably easier time than the ones that did not.

43.98 Common wrong ideas

Wrong: if a payment lands instantly, the two banks have settled. **Right:** that is true of Pix over SPI and of FedNow, and false of Faster Payments and UPI, which give an instant customer result over deferred net settlement.

Wrong: public ownership is what makes a fast payment system succeed. **Right:** the BIS comparison found public ownership the smallest of four effects at about 1.8 per cent, behind non-bank participation at 3.5 per cent; what Brazil actually used was the authority to compel participation and forbid consumer fees.

Wrong: Pix is free. **Right:** it is free to individuals because BCB Resolution 19 says it must be, while merchants pay about 0.22 per cent of the sale and participating institutions pay the central bank for settlement and for the directory.

Wrong: the Instant Payments Regulation made instant euro payments free. **Right:** it forbids charging more for an instant credit transfer than for a corresponding non-instant one, which usually means free in practice and is a different legal statement.

Wrong: the United States is technologically behind on instant payments. **Right:** it runs two independent real-time networks settling in central bank money, both now with USD 10 million ceilings against the Faster Payments maximum of GBP 1 million; what it lacks is compulsory reach, a consumer price cap and one directory.

Wrong: Americans wait days to pay each other. **Right:** Same Day ACH carried 1.4 billion payments worth USD 3.9 trillion in 2025 and Zelle carried 2 billion in the first half of that year; the traffic is spread across systems, none of which reaches everybody.

Wrong: Zelle is an instant payment rail. **Right:** it is an overlay operated by Early Warning Services that resolves an email address or mobile number to an account and rides whichever underlying rail the banks use; it is the closest thing the United States has to DICT.

Wrong: two competing instant networks give a country better coverage than one. **Right:** they do not double the reach, they halve the certainty, because a bank must join both to be sure of reaching everyone and many have joined neither.

Wrong: TARGET2 is the euro area's real-time gross settlement system. **Right:** TARGET2 was replaced by T2 on 20 March 2023; TIPS is the instant settlement platform, and SCT Inst is the scheme that runs over it.

Wrong: an instant, irrevocable rail removes the need for a dispute process. **Right:** every jurisdiction that has built one has retrofitted a slow, human, contested process alongside it, from Brazil's MED to Europe's verification of payee to Britain's mandatory reimbursement regime.

43.99 Chapter summary in 20 lines

1. Moving a number from one ledger to another has been technically trivial for forty years, so the interesting questions are who decides, who must join and who pays.
2. Brazil built the rail publicly, the United States left it to the market, and the European Union legislated, and the traffic figures match the designs.
3. Pix is a scheme owned by the Banco Central do Brasil, launched in November 2020, settling individually and in real time across dedicated PI accounts on SPI, every hour of every day.
4. Institutions licensed by the BCB with more than 500,000 active accounts are obliged to join, so reachability was never a commercial decision.
5. BCB Resolution 19 forbids charging a natural person to send, while a shop pays about 0.22 per cent against roughly 1.1 per cent for debit and 2.2 per cent for credit cards.
6. DICT maps four kinds of key, tax identifier, mobile number, email address and a random string, to accounts, and returns the account name so the payer can check.

7. Those three decisions produced 63 billion transactions worth BRL 26.4 trillion in 2024 and close to 170 million users by the five-year mark.
8. The United States has roughly nine thousand depository institutions, no legal power to compel any of them to join anything, and a batch rail that was already good enough.
9. ACH carried 35.2 billion payments worth USD 93 trillion in 2025, and has had three same-day windows since 2016, which removed most of the urgency to build anything faster.
10. When faster rails were built they were built twice, RTP in 2017 and FedNow in 2023, so reach became a commercial question rather than a given.
11. FedNow settled 8.4 million payments in 2025, a rounding error next to ACH, while Zelle carried 3.6 billion in 2024 by solving the addressing problem and leaving the plumbing alone.
12. The euro area had SCT Inst from 2017 and almost nobody used it, because joining was voluntary and the bank at the other end had not joined.
13. Regulation (EU) 2024/886 required euro-area providers to receive instant payments from 9 January 2025 and to send them, check payee names and match prices from 9 October 2025.
14. TIPS volumes rose 82.5 per cent to 2.47 billion in 2025, with a measurable settlement-ratio dip in October and recovery by December, which is what a well-executed mandate looks like from the inside.
15. “Instant” is a statement about the payee’s balance, not about the plumbing, and real-time gross and deferred net settlement put risk and liquidity in entirely different places.
16. Nothing is free: Brazilian merchants, FedNow’s 4.5 cents per credit transfer and the Eurosystem’s 0.2 of a euro cent are all somebody paying for the rail.
17. The BIS evidence ranks non-bank participation, additional use cases and cross-border links above public ownership as correlates of adoption.
18. Five things recur wherever adoption is high: central operation or compulsion, mandated reach, a rule-based price cap at the consumer end, alias addressing with a name check, and a settlement account open whenever the rail is open, with a sixth being one scheme rather than several.
19. Britain was early, and being early froze in the assumptions of 2008: it has reach and no consumer charge, but still addresses payments by sort code and account number, with Confirmation of Payee as a check on that credential rather than a replacement for it.
20. Speed and reversibility trade against each other and the trade cannot be designed away, only priced and allocated, which is why the slow human fraud process beside a fast rail is part of the design rather than a failure of it.

Sources: Banco Central do Brasil, on Pix participation criteria, SPI and DICT, BCB Resolution 1 of 12 August 2020, BCB Resolution 19 of 1 October 2020, the Pix Statistics participant tables as at 15 August 2026, and the “Pix at 5” press article of November 2025; Federal Reserve Financial Services, on the FedNow Service volume and value statistics updated 6 July 2026, operating hours, the 2026 fee schedule, the transaction limit increase effective 12 November 2025, and the FedACH processing schedule effective 12 September 2022; Nacha, on 2025 ACH Network volume and value, Same Day ACH limits and the increase to USD 10 million effective 17 September 2027; The Clearing House, on the RTP network transaction limit effective 9 February 2025 and participation figures; Early Warning Services, on Zelle volumes; European Central Bank, TARGET Services Annual Report 2025, and the ECB pages on the Instant Payments Regulation, verification of payee and TIPS; European Payments Council, on the SCT Inst and VoP scheme rulebooks; Regulation (EU) 2024/886; Bank for International Settlements, “Fast payments: design and adoption”, BIS Quarterly Review, March 2024, and BIS Bulletin no 52 on Brazil’s Pix; International Monetary Fund, “Pix: Brazil’s Successful Instant Payment System”, 2023.

Open Banking

44.0 What this chapter gives you

1. You will be able to explain why open banking is not a payment system, and name the rail that actually carried the £340.
2. You will be able to say what a payment initiation service provider may and may not do, and why only the account holder can authorise money to move.
3. You will be able to price an open banking payment against a card payment and say why the saving grows with the average transaction value rather than staying constant.
4. You will be able to tell a merchant exactly what they gain and what their customer gives up when a checkout moves from cards to pay-by-bank.
5. You will be able to distinguish the Payment Services Regulations 2017 from the CMA's Retail Banking Market Investigation Order 2017, and say which binds every bank and which produced the specification and binds nine.
6. You will be able to explain why "open banking works with your bank" is a claim about one institution on one day, and build a coverage matrix instead of relying on a principle.
7. You will be able to describe a variable recurring payment's control parameters and say why sweeping is free at the point of use and commercial VRP is not.
8. You will be able to state what the ninety-day rule now requires and why most published explanations of it are out of date.
9. You will be able to trace a payment initiation from consent creation to a terminal status, and say why building a checkout that waits for one is a mistake.
10. You will be able to explain what regulation 79 does, why its application to commercial VRP is unsettled, and why the Direct Debit Guarantee is the standard VRP will be judged against.

At twenty past nine on a Wednesday evening, a woman in Bristol pays a quarterly energy bill of £340. She is on the supplier's website. Next to the card fields there is a button that says "Pay by bank". She presses it, picks her bank from a list of logos, and her phone opens her banking app. The app shows her three lines: the amount, £340.00, the name of the energy supplier, and the account the money will leave. She holds her thumb to the sensor. The app closes, the supplier's page refreshes, and the bill is marked paid. The whole thing takes about eleven seconds. The supplier has the money, finally and completely, before she has put her phone down.

At two o'clock the following morning, while she is asleep, £85 moves out of her current account and into a savings account she holds with a different firm. She set this up seven months ago and has not thought about it since. It happens because the balance in her current account crossed a threshold she chose. No standing order was involved. No direct debit was involved. Nobody at either institution touched anything.

Both of those movements used the same machinery, and the machinery has a name that misleads almost everybody who hears it. It is called open banking, and the most important thing to understand about it is a negative: it is not a payment system. No money moves through it. It is a way of asking a bank to do something the bank could already do, and of asking in a form the bank is legally obliged to answer. The £340 travelled down the Faster Payments Service, the same pipe that has carried British instant retail payments since 2008. The £85 travelled down it too. Open banking supplied the instruction and the permission. It did not supply the road.

That distinction sounds pedantic. It is in fact the source of nearly every important consequence in this chapter, including the one that matters most to anyone building a business on it, which is what happens when the customer wants their money back.

All figures, dates, versions and limits in this chapter are accurate at the time of writing, August 2026. This is the fastest-moving area in the entire book. The commercial arrangement that makes open banking payments viable for ordinary businesses went live ten weeks before this was written, the legal framework that will govern the whole thing from the late 2020s has not yet been made, and the body that will run the standard does not yet exist. Check every date below against a primary source before you rely on it.

The plain version

Think of your bank as a large office building with a counter at the front. You go to the counter, you prove who you are, and the clerk does things for you: shows you your statements, moves your money.

Now suppose you want somebody to help you. Maybe an app that adds up all your accounts across four different banks and tells you what you actually spend on food. Maybe a shop that wants to be paid without the faff of you typing an account number off an invoice.

Before 2018, there was only one way to let a helper do either of those jobs. You gave them your keys. You told the app your online banking username and password, and it logged in pretending to be you, walked around inside your account, and copied down what it found. This was, and this is not too strong a word, a disaster. Your bank could not tell the helper apart from a thief, because there was nothing to tell apart: they were using the same credentials in the same way. If something went wrong nobody could say whose fault it was. And it broke every time the bank redesigned its login page.

Open banking replaced the keys with two service hatches cut into the wall of the building, and a rule that the bank has to keep them open.

Through the first hatch, a helper who holds a **reader's pass** can ask for a copy of your statements. Not your password, not the ability to move anything, just the numbers: balances, transactions, standing orders, the name on the account. That reader's pass is a real licence, issued by the financial regulator, and the helper's name goes on a public list.

Through the second hatch, a helper who holds a **form-filler's pass** can hand in a payment instruction on your behalf. And here is the part that carries all the weight: the form-filler cannot sign the form. They fill in the payee, the amount, the reference. Then they hand it back to you, and you walk to the counter yourself and sign it, with your own thumbprint or your own passcode, in your own bank's app. The bank checks the signature, not the form-filler's word. Then the bank pays.

That is the whole idea. Two hatches, a list of who is allowed to knock, and a hard rule that only the account holder can authorise money to move.

■ What it costs, and who it costs

Here is why a business cares, using the £340 energy bill.

If our Bristol customer had paid by debit card, a slice of that £340 would have gone to the bank that issued her card. That slice is capped by law at 0.2 per cent for a consumer debit card, which is 68 pence. If she had used a credit card the cap is 0.3 per cent, which is £1.02. Those caps only cover one slice; the card network takes a fee and the supplier's own payment provider takes a margin, and the all-in cost is meaningfully higher than the cap.

The key thing is not that the number is big. The key thing is that it is a *percentage*. Doubling the bill doubles the fee. A business collecting £3,400 pays ten times what a business collecting £340 pays, for exactly the same amount of work.

An open banking payment is not priced that way. The form-filler charges the business a fee per payment, quoted in pence, and it does not move when the amount moves. Collecting £3,400 costs the same as collecting £34. For a business with large average transaction values, an energy supplier, a letting agent, an investment platform, a car dealer, that is not a discount. That is a different order of magnitude.

The second difference is timing. The card payment showed up as approved instantly, but the supplier would not actually have the cash for a day or two. The open banking payment went over Faster Payments, which runs every hour of every day including Christmas morning, and the money was in the supplier's account in seconds. On a Wednesday evening. Which is when people pay bills.

The third difference is the one nobody puts on the marketing page. With the card, the supplier's money can be taken back off them, months later, if the customer complains to her bank and the bank agrees. With the open banking payment, it cannot. Once a Faster Payment has gone, it has gone. If the customer wants a refund, the supplier has to choose to send her one.

That is the trade, in one sentence: **the merchant gets the money faster, cheaper and for keeps, and the customer gives up the right to reach back and take it.**

■ Standing orders that can change their mind

The £85 that moved at two in the morning is the other half of open banking, and it has a clumsy name: a variable recurring payment, or VRP.

A standing order is a fixed instruction: £85 on the 3rd of every month, until I say stop. It cannot adapt. A direct debit is the opposite: the company you are paying decides the amount, and tells your bank to take it. That is flexible, but you have handed the steering wheel to somebody else.

A VRP sits between the two. You sign one form, once, that says something like: *this named helper may take money from my account, but never more than £200 in a single payment, never more than £600 in a calendar month, only to this one named destination, and only until 5 April next year*. After that, the helper can make payments inside those walls without waking you up, and cannot make one that steps outside them by a penny. Your bank enforces the walls, not the helper.

Right now the overwhelming majority of live VRPs do one job, and it has a name too: **sweeping**. Sweeping means moving your own money between your own accounts. Current account into savings when the balance goes above a threshold. Savings into current account when the balance goes below one, so you do not slide into an overdraft. Spare cash into an account that repays a loan early. The nine largest British banks were required to build sweeping, and they have all built it.

The interesting frontier, and it is only just open, is letting a VRP pay somebody else: your energy company, your council, a charity, a subscription. That is called commercial VRP, and it went live in Britain on 2 June 2026 under a new industry scheme. It is the first serious attempt to build something that does the direct debit's job without the direct debit's three-day cycle and without the direct debit's paper-era limitations.

■ How big is this

Open banking payments made in the UK in 2025	351 million	up 57 per cent on 2024
Successful requests answered by banks in 2025	24.0 billion	up 27 per cent
User connections at December 2025	16.5 million	up 36 per cent
Successful requests in June 2026 alone	2.8055 billion	99.50 per cent success rate
Average time a bank took to answer, June 2026	349 milliseconds	

Those are Open Banking Limited's own published figures. In mid-2026 the ecosystem passed one billion open banking payments and one hundred billion requests since it started in 2018.

For scale: 31 million open banking payments were made in March 2025, which Open Banking Limited put at roughly one in every thirteen Faster Payments. So it is real, it is growing quickly, and it is still a minority of a minority of how Britain moves money.

■ The one-line summary

Open banking is a set of doors cut into the wall of your bank, a licensed list of who may knock, a rule that the bank must answer, and an absolute requirement that only you can authorise money to leave. Everything behind the door is the old plumbing.

Where the plain version stops being true

■ The hatch is not the same size at every bank

The plain version says "the bank has to keep the hatch open", and that is true, but it hides that two entirely different pieces of law are doing two entirely different jobs, and they cover different institutions.

The first is the Payment Services Regulations 2017, which apply to every account servicing payment service provider in the country. They require access. They say almost nothing about what that access looks like. Regulation 69 says a bank must communicate securely with a payment initiation service provider and treat its payment orders exactly as it treats orders coming direct from the customer. Regulation 70 says the same for account information. Neither regulation specifies a data format, a field name, an uptime target, or a response time.

The second is the Competition and Markets Authority's Retail Banking Market Investigation Order 2017, made under the Enterprise Act 2002. That is where the actual specification came from, and it binds nine institutions: Barclays, HSBC, Lloyds Banking Group, Nationwide, NatWest Group, Santander, Danske Bank, Bank of Ireland and Allied Irish Bank. Those nine, universally called the CMA9, were required to fund and implement a common standard against a published roadmap, and to submit to monitoring.

Everybody else built whatever they liked, or built to the standard voluntarily, or built something that technically satisfies the regulation and is miserable to integrate with. Sweeping VRP was mandated on the CMA9 alone. So the sentence “open banking works with your bank” is a claim about one specific institution on one specific day, not a claim about the United Kingdom. Anyone building a product needs a coverage matrix, not a principle.

■ **There is not one consent, there are at least two, and they drift**

The plain version has you signing one form. In reality two separate permissions exist, they are held by two different parties, they have different lifetimes, and cancelling one does not reliably cancel the other.

The first is the consent you give the third party. That is a contractual and data-protection permission: you told this app it may look at your accounts, for these purposes, for this long.

The second is the authorisation you give your bank. That is a strong customer authentication event which produces an access token and a refresh token, held by the third party, scoped to a consent resource that lives on the bank’s servers.

They come apart routinely. Delete the app from your phone and the tokens keep working. Cancel in the app and the bank may still show a live consent in its “connected services” screen. Cancel at the bank and the app’s next call fails, but the app may not tell you why.

The famous ninety-day rule has also changed, and most explanations of it are now wrong. Originally the bank had to re-authenticate you every ninety days for data access to continue, which meant a monthly-budgeting app dragged you through your bank’s login four times a year for every bank you had connected. The Financial Conduct Authority changed this in policy statement PS21/19, adding a new Article 10A to the UK version of the technical standards. Since 26 March 2022, the obligation is a *reconfirmation of consent* that the account information provider obtains from you, at least every ninety days, rather than a *re-authentication* the bank performs. The FCA expected banks to adopt the exemption by 30 September 2022. The friction moved from the bank’s app to the third party’s app, which was the point. The ninety days did not go away.

■ **The money arrives instantly, but that is not the same as knowing it arrived**

The plain version says the supplier has the money in seconds. Usually true. What is not true is that the supplier’s software has been told so with certainty.

A payment initiation service provider does not receive a settlement confirmation from Pay.UK. It receives a status from the bank about the payment order, which is a different object in a different system. Statuses lag, sometimes by minutes, and a number of banks stop at an intermediate status and never emit the terminal one. Building a checkout that waits for final confirmation before releasing goods is a good way to build a checkout that hangs.

The limits are similarly misleading. The Faster Payments scheme maximum is £1 million per transaction, and you will almost never meet it. Individual banks set their own lower ceilings, and the ceiling applied to payments initiated through open banking is frequently lower than the ceiling the same bank applies inside its own app for the same customer. Some banks step the limit by authentication method, some by account age, some by whether the payee is already known. Never size a product against the scheme limit. Size it against the worst bank on your coverage matrix, and find out what that number is by testing rather than by asking.

And “instant” has a legal floor rather than a legal guarantee. Pay.UK’s own description is that payments between direct participants normally arrive almost immediately but can occasionally take up to two hours, with the statutory backstop being end of the following business day.

■ It is not a cheaper card, and treating it as one is expensive

This is the correction that costs real money.

A card payment is a *pull* with a dispute system bolted to it. Two dispute systems, in fact. Chargeback is a set of private rules written by the card schemes: the customer complains, the issuer reverses the transaction, the merchant has to defend it. It is not law, and it exists because the schemes decided it should. On top of that, for credit cards, section 75 of the Consumer Credit Act 1974 makes the card issuer jointly and severally liable with the retailer for breach of contract or misrepresentation, where the cash price of the item is over £100 and no more than £30,000. Section 75A extends a narrower right above that.

An open banking payment is a *push* with neither. There is no chargeback, because there is no scheme sitting between the parties writing dispute rules. There is no section 75, because there is no credit agreement. A refund is a brand new payment, in the opposite direction, that the merchant elects to send.

The statutory refund right that does exist, regulation 76 of the Payment Services Regulations 2017, requires a payment service provider to refund an *unauthorised* transaction and restore the account, by the end of the business day following the day it becomes aware. But an open banking payment that the customer approved with strong customer authentication is, by definition and by design, authorised. The regulation is not available for a sofa that never arrived.

The authorised push payment reimbursement regime, live since 7 October 2024, is also not the answer. It covers being deceived into sending money to a criminal, up to £85,000 per claim, with the cost shared equally between the sending and receiving firms. It does not cover a merchant dispute. A supplier who took your money and then went quiet is a civil matter, not a scam claim.

So the honest framing is that open banking payments are a *bank transfer with a good user interface*, and they carry exactly the protections of a bank transfer, which is to say almost none beyond the fraud regime. That is fine, and often better than fine, for paying your own energy bill or moving your own money. It is a materially worse deal for the consumer than a credit card when buying a mattress from a company they have never heard of. Any business presenting the two side by side and describing them as equivalent is misrepresenting the product.

The technical version

■ Two legal instruments, one ecosystem

UK open banking sits on two foundations that are frequently conflated.

The Payment Services Regulations 2017 (SI 2017/752) transposed the second Payment Services Directive, Directive (EU) 2015/2366, and came into force on 13 January 2018. The relevant provisions:

- Regulation 68, confirmation of availability of funds, for card-based payment instrument issuers.
- Regulation 69, access to payment accounts for payment initiation services. The ASPSP must communicate securely with the PISP and treat its payment orders identically to orders received directly from the payer. The PISP must not hold the payer’s funds at any time, must not store sensitive payment data, and must not alter the amount, the payee or any other feature of the transaction.

- Regulation 70, access to payment accounts for account information services. The AISP must have the user's explicit consent, must identify itself in every communication session, and must access only the designated accounts and their associated transactions.
- Regulation 71, limits on the use of payment instruments and access to payment accounts. An ASPSP may deny a TPP access for reasonably justified and duly evidenced reasons relating to unauthorised or fraudulent access, but must notify the customer, report the incident to the FCA, and restore access once the justification lapses.
- Regulation 100, authentication, which imposes strong customer authentication.

The Retail Banking Market Investigation Order 2017, made by the CMA under the Enterprise Act 2002, produced the standard. Article 10.5 of Part 2 provided for an Implementation Trustee to drive delivery against an agreed roadmap. Open Banking Limited, originally the Open Banking Implementation Entity, was established by the CMA9 in October 2016 to do that work. The CMA declared the roadmap substantially complete on 12 January 2023, with six providers finished, and confirmed full completion in September 2024 once Danske Bank, Bank of Ireland and Allied Irish Bank had delivered their remaining obligations, including sweeping VRP.

The forward-looking picture is different again. The Joint Regulatory Oversight Committee has been wound up and replaced by a department inside the FCA, which is now lead regulator. HM Treasury is expected to legislate under the **Data (Use and Access) Act 2025**, whose Part 1 provides for smart data schemes, to give the FCA rulemaking powers over open banking. The FCA set out its expectations for a successor to Open Banking Limited, called the Future Entity, in Feedback Statement **FS25/4** (August 2025), and wrote to trade associations on 19 December 2025, published 16 January 2026, inviting industry to coalesce behind a candidate body and offering an independent consultant-led assessment. The FCA expects the Future Entity and commercial scheme operators to be regulated as **interface bodies** under the DUAA, and intends to consult on the long-term regulatory framework by the end of 2026.

■ The role taxonomy

Term	Meaning
ASPSP	Account servicing payment service provider. The bank or building society holding the account.
AISP	Account information service provider. Reads account data.
PISP	Payment initiation service provider. Constructs and submits payment orders for the customer to authorise.
CBPII	Card-based payment instrument issuer. May ask a yes or no question about whether funds are available.
TPP	Third party provider. Umbrella term for AISPs, PISPs and CBPIIs.
PSU	Payment service user. The customer.

A firm providing account information services only may be a registered account information service provider, a lighter regulatory status. Payment initiation requires full authorisation as a payment institution, or the firm must already be a credit institution or e-money institution, and it must hold professional indemnity insurance or a comparable guarantee. In practice a large proportion of merchants do not hold either permission and reach the rails through a regulated PISP acting as their provider, or as an agent.

Identity in the ecosystem is managed through the **Open Banking Directory**. Participants enrol, obtain certificates, and publish software statements. Third parties obtain a Software Statement Assertion signed by Open Banking, which can be used for dynamic client registration at each ASPSP's authorisation server, avoiding a manual onboarding per bank.

Certificates deserve a paragraph because Brexit changed them. Article 34 of the SCA-RTS originally required eIDAS qualified certificates: a QWAC for transport, a QSeal for signing. UK TPPs' eIDAS certificates were revoked at the end of the transition period. The UK version of the RTS was amended to permit alternatives, and the Open Banking Certificate Authority issues ETSI-profile equivalents: **OBWAC** for website authentication and **OBSeal** for electronic seals. EEA TPPs accessing UK ASPSPs may still present QWACs and QSeals, which can be uploaded to the Directory.

As at June 2026 the Directory carried 327 regulated providers: 88 account providers and 239 third party providers, of which 143 regulated entities had at least one proposition live with customers.

■ The security profile

The current specification is the Open Banking Standard, currently v4.0.1. Version 4.0 was published on 28 June 2024 and was the first major release since 2018, developed with over forty ecosystem participants.

Its headline change was the security uplift from FAPI 1.0 Implementer's Draft 2 to **FAPI 1.0 Advanced Final**, made mandatory, with the earlier profile deprecated at the end of 2024. Version 4.0 also aligned message content with ISO 20022, including the fields required for Bank of England CHAPS compliance from 1 May 2025, centralised the codeset repository, and improved payment status reporting and error messaging.

The concrete requirements a practitioner will implement:

Transport. TLS 1.2 with mutual authentication. The client presents its transport certificate; the ASPSP validates it against the Directory.

Client authentication. `private_key_jwt`. The TPP signs a JWT assertion with its private key; the ASPSP validates against the registered public key. Shared secrets are not used.

Authorisation. OAuth 2.0 and OpenID Connect. The specification supports the authorisation code grant and identifies the **hybrid grant** as the preferred redirect mechanism under the FAPI profile. Authorisation parameters are packaged in a signed JWT **request object**, which prevents tampering during the redirect.

Intent model. Rather than relying on coarse OAuth scopes, the standard uses dedicated consent resources: `account-access-consent`, `domestic-payment-consent`, `domestic-vrp-consent` and their siblings. An access token binds to exactly one PSU and one intent.

Message signing. Detached JWS per RFC 7515 Appendix F, algorithm PS256, carried in the `x-jws-signature` header. The JOSE header carries critical claims `http://openbanking.org.uk/iat`, `http://openbanking.org.uk/iss` and `http://openbanking.org.uk/tan`, the last identifying the trust anchor domain. This gives non-repudiation independently of the TLS session. Message encryption via JWE is optional.

Correlation headers. `x-fapi-interaction-id`, an RFC 4122 UUID that the ASPSP must echo in its response, is the single most useful field in the entire specification when you are on the phone to a bank's support desk about a failed payment. `x-fapi-auth-date` carries the time of the PSU's last

login with the TPP. `x-fapi-customer-ip-address` carries the PSU's IP when the PSU is present, and its absence signals an unattended call.

Idempotency. POST endpoints that create resources take `x-idempotency-key`, maximum 40 characters. The ASPSP must treat an identical key from the same TPP within 24 hours as a duplicate. The TPP must not modify the request body while reusing a key, and must leave at least one second between a submission and its retry.

Authentication models. The redirect model sends the PSU to the ASPSP's authorisation endpoint, in practice app-to-app on mobile, and back again. The decoupled model uses CIBA, Client Initiated Backchannel Authentication: the TPP initiates, the PSU authenticates on a separate device or channel, and the TPP polls. Decoupled matters for call centres, in-branch journeys and any flow where the PSU is not holding the device that displays the merchant.

■ A payment initiation, step by step

1. The PISP POSTs to `/domestic-payment-consents` with the payee, amount, currency, reference and, optionally, the debtor account. The ASPSP returns a `ConsentId`. Nothing has been authorised and no money has moved.
2. The PISP constructs a signed request object carrying the `ConsentId` as an intent claim and redirects the PSU to the ASPSP's authorisation endpoint.
3. The ASPSP authenticates the PSU with strong customer authentication. If the debtor account was not specified, the PSU selects it here. The ASPSP plays the consent back to the PSU: amount, payee, reference. The PSU accepts or rejects.
4. The consent resource transitions to authorised. The ASPSP returns an authorisation code to the PISP's redirect URI. The PISP exchanges it at the token endpoint for an access token scoped to that consent.
5. The PISP POSTs to `/domestic-payments`, referencing the `ConsentId`, with a request body that must match the consented payment exactly.
6. The ASPSP submits the payment to a clearing system. For a domestic single immediate payment in sterling this is, in practice, almost always Faster Payments.
7. The PISP polls GET `/domestic-payments/{DomesticPaymentId}` for status.

Version 4.0 aligned payment order statuses to the ISO 20022 external code set. The values a practitioner will encounter include RCVD received, PDNG pending, ACTC accepted technical validation, ACCP accepted customer profile, ACFC accepted fraud checking, ACSP accepted settlement in process, ACWC accepted with change, ACSC accepted settlement completed, ACWP accepted with partial execution, ACCC accepted and closed, BLCK blocked, RJCT rejected, and CANC cancelled for future-dated payments and standing orders.

The specification also supports **multi-authorisation** for accounts requiring more than one signatory, through the `MultiAuthorisation` block carrying `Status`, `NumberRequired`, `NumberReceived`, `LastUpdateDateTime` and `ExpirationDateTime`, with the PISP able to request an authorisation type of `Single` or `Any` and an optional `CompletionDateTime`. Corporate journeys live or die on this.

There is a `RefundAccount` structure in the domestic and international payment models, which returns the payer's account details to the PISP so that a merchant can send a refund. Note carefully what this is: it is an address, so that the merchant can make a new payment. It is not a reversal mechanism.

■ Strong customer authentication in the consent flow

Regulation 100 of the PSRs 2017 requires a payment service provider to apply strong customer authentication where a user accesses a payment account online, whether directly or through an AISP; initiates an electronic payment transaction; or carries out any action through a remote channel which may imply a risk of payment fraud or other abuses.

Strong customer authentication means two or more elements from independent categories: knowledge, something only the user knows; possession, something only the user has; and inherence, something the user is. Independence means the breach of one does not compromise the others.

Regulation 100(2) adds **dynamic linking**: where a payer initiates an electronic remote payment transaction, directly or through a PISP, the authentication must include elements which dynamically link the transaction to a specific amount and a specific payee. This is why your bank's app shows you the amount and the payee name on the approval screen and why a PISP cannot alter either after authorisation. It is also why the payment request body must match the consent.

The exemptions live in the SCA-RTS, originally Commission Delegated Regulation (EU) 2018/389, onshored and now maintained by the FCA as technical standards in the Handbook. The ones that matter here:

Article 10 and Article 10A, account information. Article 10 permitted an ASPSP not to apply SCA where a user accesses balance or transaction data, subject to a ninety-day limit. Article 10A, introduced by PS21/19, removed the periodic re-authentication at the ASPSP for AIS access through a TPP and shifted the obligation onto the AISP to obtain the customer's explicit reconfirmation of consent at least every ninety days. The FCA strongly encouraged adoption by 30 September 2022.

Article 13, trusted beneficiaries. An ASPSP need not apply SCA where the payee is on a list of trusted beneficiaries previously created by the payer. This is the mechanism that makes VRP work at all.

Article 11, contactless, and Article 16, low value remote, are card-side and rarely relevant to open banking, though the FCA has publicly encouraged continued support for the contactless exemption for charitable donations.

The FCA has also been explicit that firms must develop SCA solutions that work for all groups of consumers, including those without mobile telephones. A journey that assumes app-to-app on a smartphone is not, on its own, an adequate SCA implementation for a whole customer base.

■ Variable recurring payments

The VRP API Profile, at v3.1.10 of the read-write specification, defines two resources:

- **domestic-vrp-consents**, a consent between a PSU and a TPP permitting the TPP to create payments subject to **control parameters**.
- **domestic-vrps**, individual payment orders that must fall within those parameters.

The control parameters set the walls. They include the validity window, `ValidFromDateTime` and `ValidToDateTime`; a maximum individual amount; and periodic limits, each with an amount, a period type, and a period alignment. The period type enumerations are Day, Week, Fortnight, Month, Half-year and Year, and an ASPSP must support all of them. Period alignment is either Consent, meaning the window starts when the consent was created, or Calendar. ASPSPs may impose additional control parameters, limits and restrictions on non-sweeping VRPs.

The flow is: the TPP POSTs a `domestic-vrp-consents` resource, the PSU authorises it once with SCA, and thereafter the TPP POSTs `domestic-vrps` against it without further SCA, provided each payment sits inside the parameters. GET on `vrp-details` returns funds availability and remaining allowance.

Sweeping is a regulatory category, not a technical one. The CMA approved VRP as the mechanism for implementing sweeping under item A10 of the revised roadmap on 27 July 2021, extended the CMA9 implementation deadline to July 2022 by letter of 15 November 2021, and published a clarification of the definition of sweeping in March 2022. On 3 December 2025 the CMA published a letter responding to an OBL recommendation regarding the sweeping use case of alternative forms of credit that closely compete with overdrafts.

The requirements the CMA9 had to meet for sweeping, as set out in OBL's standards proposition, are worth listing because they explain the consent screens you actually see:

- The sweeping consent must specify the payee account name, the payee account identification, the maximum cumulative amount per time window, and the maximum amount of each individual sweeping payment.
- The ASPSP must tell the PSU that the payee will be added to their trusted beneficiary list, and must add it once consent setup completes.
- The ASPSP must then apply the trusted beneficiary SCA exemption when the PISP initiates a payment inside the consent parameters, or the transfer-to-self exemption where the destination is at the same ASPSP. It may decline to apply the exemption in exceptional circumstances, taking an objective approach consistent with the proportionality requirements of the PSRs.
- The ASPSP must provide a specific error indicator if the payee has been removed from the trusted beneficiary list and the exemption can no longer be applied.
- The PISP must attest to the ASPSP that the consent is a sweeping consent.
- The ASPSP must provide non-discriminatory access both to sweeping consent setup and to the initiation of sweeping payments.

That last requirement is the reason sweeping is free at the point of use and commercial VRP is not. Non-discriminatory access, in the CMA's sense, meant the CMA9 could not charge for it.

■ Commercial VRP and the UK Payments Initiative

Extending VRP beyond me-to-me sweeping required something the CMA Order did not provide: a commercial model. A bank has no obligation to offer VRP to a merchant and no reason to build the capacity for free.

Under the JROC, the Payment Systems Regulator led work on a commercial model for so-called phase 1 use cases. The FCA took over as lead regulator and continued the work jointly with the PSR, publishing a summary report on development and rollout in December 2025.

In late 2025, 31 firms came together to establish **UK Payments Initiative Ltd (UKPI)**, an independent industry-owned company to own and operate a commercial VRP scheme with its own rulebook, commercial model and operational standards. The FCA described the first live payments under the scheme as the start of a new era for payments and open banking in the UK, and expected them in the first quarter of 2026.

The scheme was announced as launched on **2 June 2026**, at Money20/20 Europe in Amsterdam. UKPI's own announcement names 24 founding shareholders, including Barclays, HSBC, Lloyds Banking Group, NatWest, Nationwide and Santander, alongside Monzo, Revolut, Starling, GoCardless, Plaid, TrueLayer, Yapily and Token.io. The scheme's first wave covers government payments, utilities,

charities and financial services, with subscriptions and wider e-commerce planned to follow. A launch coverage target of around 75 per cent of UK current accounts is quoted consistently in the guidance scheme participants have issued to merchants, and it is the figure the industry has settled on. It is worth being clear about its standing: it does not appear in the FCA's statement on the launch, nor in the regulators' correspondence with the CMA, and no public scheme document confirms it, so the number rests on industry restatement rather than on a published primary source.

The commercial model itself required a competition-law answer, because a group of competitors agreeing a central price is exactly the arrangement competition law exists to police. On 15 January 2026 the FCA and the PSR wrote to the CMA setting out their position; on 16 January the CMA confirmed it did not intend to take a different view on prioritisation; and on 20 January 2026 the regulators published a joint statement confirming they would not at this stage prioritise a Competition Act 1998 investigation into UKPI's centralised access fee model. That forbearance runs until the government's anticipated legislative framework or July 2027, whichever comes first. The FCA has said it will assess industry-led cVRP growth by the end of 2026 and fold the lessons of the first phase into the long-term framework.

VRPs accounted for roughly 15 to 16 per cent of open banking transaction volume at the time of writing, with sweeping VRP volumes having nearly doubled year on year in 2025.

■ Why it settles over Faster Payments, and what follows from that

Nothing in the specification requires Faster Payments. The standard describes a payment order; the ASPSP decides how to execute it. LocalInstrument can carry a scheme preference and some ASPSPs support CHAPS or scheduled Bacs credits for particular payment types.

In practice, the domestic single immediate payment, which is the overwhelming majority of open banking payment volume, executes over the Faster Payments Service, for the obvious reason: it is the only sterling retail rail that clears in seconds, around the clock, seven days a week. Bacs takes three working days on a submission, processing and entry cycle, and is deferred net settlement, which is fine for payroll and useless for a checkout. CHAPS is same-day, irrevocable real-time gross settlement inside the Bank of England's RTGS system, carrying the large majority of sterling value, with an early-afternoon cut-off for the same-day guarantee, and it is not priced for a £340 energy bill.

Four consequences follow, and they are the practical heart of this chapter.

Irrevocability. Pay.UK's own description is unambiguous: because of their real-time nature, once sent, Faster Payments cannot be cancelled. There is a Credit Payment Recovery process for misdirected payments, which is a best-efforts request to the receiving institution to return funds it has not yet released, not a right. Confirmation of Payee, launched in 2020, checks the payee name against the account before you send, which prevents some errors and some frauds but changes nothing about finality once the payment has gone.

Limits. The scheme maximum is £1 million per transaction. Individual PSPs set their own lower ceilings, and channel-specific ceilings for open banking are frequently lower still.

No chargeback. There is no scheme adjudication layer, because the scheme is Faster Payments and Faster Payments does not adjudicate merchant disputes. Compare the card world, where chargeback rights are written by the card networks into their operating rules and where section 75 of the Consumer Credit Act 1974 imposes joint and several liability on a credit provider for purchases with a cash price over £100 and up to £30,000.

Refunds are new payments. A merchant refunding an open banking payment initiates a fresh credit transfer to the payer. That is why the specification carries a `RefundAccount` structure: not to reverse anything, but to tell the merchant where to send the money. Operationally this means refunds need their own payout capability, their own reconciliation, their own fraud controls and their own audit trail, none of which a merchant migrating off cards will have.

■ The protection map

	Card (credit)	Card (debit)	Direct Debit	Open banking payment
Direction	Pull	Pull	Pull	Push
Rail	Card scheme	Card scheme	Bacs	Faster Payments
Merchant has funds	After settlement, days	After settlement, days	Three working day cycle	Seconds
Reversal by customer	Chargeback, scheme rules	Chargeback, scheme rules	Direct Debit Guarantee, immediate refund on request	None
Statutory purchase protection	Section 75 CCA 1974, £100 to £30,000	None	None	None
Fraud reimbursement	Unauthorised transaction rules, PSRs reg 76	Unauthorised transaction rules, PSRs reg 76	Guarantee, plus reg 76	APP regime since 7 Oct 2024, up to £85,000; reg 76 for unauthorised
Cost to merchant	Percentage	Percentage	Low, per item	Flat fee, per item

The statutory refund provisions in the PSRs are worth stating precisely, because they are routinely misdescribed:

- **Regulation 74** requires notification without undue delay and in any event within thirteen months of the debit date, otherwise the redress provisions are unavailable.
- **Regulation 76** requires the PSP to refund an unauthorised transaction and restore the account, as soon as practicable and no later than the end of the business day following the day it becomes aware. Where a PISP is liable, regulation 75(2) applies and the PISP must compensate the ASPSP immediately on request.
- **Regulation 77** caps the payer's own liability at £35 for losses from a lost, stolen or misappropriated payment instrument, with the cap disapplied where the payer acted fraudulently or with intent or gross negligence, and with no liability at all where the ASPSP failed to apply SCA that regulation 100 required.
- **Regulation 79** gives a refund right for authorised transactions *initiated by or through the payee*, where the authorisation did not specify the exact amount and the amount exceeded what the payer could reasonably have expected. Regulation 80 requires the request within eight weeks of the debit and a response within ten business days.
- **Regulation 95** gives a right of recourse between PSPs where liability is attributable to another provider or intermediary, including for failure to apply SCA.

Regulation 79 is the interesting one for VRP, and the position is genuinely unsettled. Its trigger is a transaction initiated by or through the payee, which is the direct debit shape. A VRP is initiated by a PISP acting for the payer, within limits the payer set, which is a different shape. Whether and how

the regulation 79 refund right attaches to a commercial VRP is one of the questions the FCA's long-term framework will need to answer, and it is the single largest open consumer-protection question in the area. The direct debit's competitive advantage over VRP is not technical. It is the Direct Debit Guarantee, which promises an immediate refund from your own bank, on request, with no time limit and no argument. VRP as currently specified has no equivalent, and the UKPI rulebook's treatment of disputes will determine whether it needs one.

■ What changes next

Three things are moving, and a chapter written eighteen months from now will read differently.

In the UK, the Data (Use and Access) Act 2025 provides the vehicle. HM Treasury is expected to legislate to give the FCA rulemaking powers over open banking; the FCA intends to consult on the long-term regulatory framework by the end of 2026; and industry is expected to stand up a Future Entity to succeed Open Banking Limited as the primary API standard-setting body, regulated as an interface body under the DUAA. Beyond payments, the same Act provides for smart data schemes in energy, telecoms, transport and retail, which is what "open finance" and eventually "open data" mean in a British context.

In the EU, the successor to PSD2 has landed. The Parliament and Council reached provisional political agreement on PSD3 and the accompanying Payment Services Regulation on 27 November 2025; the Council published final compromise texts on 23 April 2026 following COREPER endorsement on 22 April; the ECON committee approved the agreed text on 5 May 2026. Formal adoption by both institutions was still outstanding at the time of writing. A separate Regulation on a Framework for Financial Data Access, which is the EU's open finance instrument, remained in trilogue. Among the changes relevant here, the package clarifies that SCA for account information access should be performed by the account provider once, at the outset.

In the rails beneath, the EU's Instant Payments Regulation has made instant euro transfers universal and price-capped, which removes the main structural reason European open banking payments lagged British ones. Britain's advantage in this area was always the early availability of a genuinely instant, genuinely ubiquitous retail rail. That advantage is now shared.

The through-line is worth restating, because it is the thing most easily lost. Open banking did not build a payment system. It built a legally mandated, cryptographically enforced way of asking a bank to use the payment system it already had, and a licensed register of who may ask. Everything good about it, the speed, the price, the finality, comes from Faster Payments. Everything difficult about it, the absence of chargeback, the absence of section 75, the fact that a refund is a fresh payment somebody has to choose to make, comes from Faster Payments too. It is the same fact wearing two faces, and any product built on this layer has to be designed for both of them at once.

44.98 Common wrong ideas

Wrong: open banking is a payment system and money moves through it. **Right:** it is a legally mandated, cryptographically enforced way of asking a bank to use the payment system it already had, and the money travels down Faster Payments.

Wrong: an open banking payment is a cheaper card payment. **Right:** it is a push with no chargeback and no section 75, and a refund is a brand new payment in the opposite direction that the merchant elects to send.

Wrong: regulation 76 will get the customer's money back when the goods never arrive. **Right:** regulation 76 covers unauthorised transactions, and a payment the customer approved with strong customer authentication is authorised by definition and by design.

Wrong: the authorised push payment reimbursement regime covers a merchant dispute. **Right:** it covers being deceived into sending money to a criminal, up to £85,000 per claim; a supplier who took the money and went quiet is a civil matter.

Wrong: every bank must keep the hatch open to the same standard. **Right:** the Payment Services Regulations require access without specifying a data format, an uptime target or a response time, and only the CMA9 were required to implement the common standard and sweeping VRP.

Wrong: deleting the app or cancelling in it ends the connection. **Right:** there are at least two permissions, held by two different parties with different lifetimes, and they come apart routinely.

Wrong: the ninety-day rule drags a customer through their bank's login four times a year. **Right:** since 26 March 2022, Article 10A has made it a reconfirmation of consent that the account information provider obtains, not a re-authentication the bank performs.

Wrong: the Faster Payments scheme maximum of £1 million is the ceiling to build against. **Right:** individual banks set lower limits, and the limit applied to open-banking-initiated payments is frequently lower than the same bank applies in its own app, so size against the worst bank on the coverage matrix and find the number by testing.

Wrong: an accepted payment status means the money has settled. **Right:** a PISP receives a status about the payment order from the bank, not a settlement confirmation from Pay.UK, statuses lag, and some banks stop at an intermediate status and never emit the terminal one.

Wrong: the RefundAccount structure is a reversal mechanism. **Right:** it is an address, returned so that the merchant can make a fresh payment, with its own payout capability, reconciliation, fraud controls and audit trail.

44.99 Chapter summary in 20 lines

1. Open banking is not a payment system: it is a way of asking a bank to do something it could already do, in a form the bank is legally obliged to answer.
2. Both the £340 energy bill and the £85 overnight sweep travelled down the Faster Payments Service, which open banking did not build and does not own.
3. It replaced shared passwords and screen-scraping with two regulated interfaces: one that reads account data and one that submits a payment order for the customer to authorise.
4. The payment initiation provider fills in the form but never signs it; only the account holder authorises, with strong customer authentication, in their own bank's app.
5. Two legal instruments underpin the ecosystem: the Payment Services Regulations 2017, which require access from every account provider, and the CMA Order 2017, which produced the specification and bound only the CMA9.
6. Everybody outside the CMA9 built what they liked, so coverage is a matrix of institutions rather than a national fact.
7. Card fees are percentages and open banking fees are flat in pence, which is why the saving grows with the size of the payment rather than staying constant.
8. The money arrives in seconds at twenty past nine on a Wednesday evening, because Faster Payments runs every hour of every day including Christmas morning.

9. It also arrives for keeps: a push payment has no chargeback, no section 75 protection and no reversal mechanism.
10. Regulation 76 refunds unauthorised transactions, and a payment approved with strong customer authentication is authorised, so it is unavailable for a sofa that never arrived.
11. The authorised push payment regime covers deception by a criminal rather than a dispute with a merchant.
12. So an open banking payment is a bank transfer with a good user interface: better than a card for paying your own energy bill, materially worse for buying a mattress from a company you have never heard of.
13. Consent is at least two permissions, one contractual with the third party and one an authorisation event at the bank, and they drift apart in ordinary use.
14. Since March 2022 the ninety-day obligation has been a reconfirmation of consent obtained by the account information provider, which moved the friction rather than removing it.
15. Status is not settlement: the provider sees the bank's view of the payment order, statuses lag, and a checkout that waits for a terminal status is a checkout that hangs.
16. A variable recurring payment is one authorisation plus control parameters, a maximum per payment, periodic limits, a named payee and a validity window, enforced by the bank rather than by the third party.
17. Sweeping, moving your own money between your own accounts, was mandated on the CMA9 with non-discriminatory access, which is precisely why it is free at the point of use.
18. Commercial VRP needed a commercial model the CMA Order never provided, which arrived with the UK Payments Initiative scheme launched on 2 June 2026 and a regulatory forbearance on its centralised access fee running to July 2027.
19. Whether regulation 79's refund right attaches to a commercial VRP is genuinely unsettled, and the Direct Debit Guarantee, an immediate refund on request with no time limit, is the standard against which VRP will be measured.
20. Everything good about open banking payments and everything difficult about them come from the same fact wearing two faces: they run on Faster Payments.

Sources: Open Banking Limited, on the Open Banking Standard v4.0 and v4.0.1, the Read-Write API Profile and VRP API Profile v3.1.10, the Payment Initiation API Profile v4.0, the standards proposition on Variable Recurring Payments, Directory and certificate guidance on OBWAC and OBSeal, the API performance statistics for June 2026, the ecosystem highlights published on openbanking.org.uk, and the year-in-review analysis for 2025; the Competition and Markets Authority, Retail Banking Market Investigation Order 2017, the roadmap completion confirmations of January 2023 and September 2024, the letters on VRP for sweeping of 27 July 2021, 15 November 2021, March 2022 and 3 December 2025; the Payment Services Regulations 2017 (SI 2017/752), regulations 68 to 71, 74 to 80, 95 and 100, as published on legislation.gov.uk; the Financial Conduct Authority, PS21/19 on changes to the SCA-RTS, the strong customer authentication pages, FS25/4 on the design of the Future Entity, the letter to trade associations of 19 December 2025, the open banking progress statement of 2025, the joint statement with the PSR of 20 January 2026 on open banking pricing models, and the statement on the launch of the UK Payments Initiative scheme; the Payment Systems Regulator, PS24/7 on the maximum level of reimbursement for Faster Payments APP scams; Pay.UK, on how Faster Payments work and on Confirmation of Payee; UK Payments Initiative, launch announcement of 2 June 2026; Directive (EU) 2015/2366 and Commission Delegated Regulation (EU) 2018/389; the Data (Use and Access) Act 2025; the Council of the European Union, final compromise texts for PSD3 and the Payment Services Regulation of 23 April 2026; the Consumer Credit Act 1974, sections 75 and 75A, and the Financial Ombudsman Service guidance on section 75 and chargeback.

ABOUT THE AUTHOR

Shikhar Singh

Founder & Chairman

KedByte Technologies Private Limited

Shikhar Singh is the Founder & Chairman of KedByte Technologies Private Limited.

He wrote this book from the position of someone building payment software rather than someone commenting on it, and the difference shows in what it chooses to explain. The chapters that go deepest are the ones where he has watched capable engineers lose days: the gap between authorisation and settlement, the several identifiers all called "the transaction number", and reconciliation, which is unglamorous, unavoidable, and where more young financial technology companies fail than for any more interesting reason.

His conviction is that payments has no genuinely hard idea in it, only a very long chain of simple ones, most of which are documented for people who already understand them. The barrier is not difficulty. It is that the industry writes for itself. That is a solvable problem, and this book is an attempt at solving it.

Under his direction, KedByte Technologies Private Limited treats technical education as infrastructure rather than as a product feature: knowledge that should be transferable, verifiable, and free of the vocabulary barriers that keep capable people out of the field.



ABOUT THE PUBLISHER

KedByte Technologies Private Limited

KedByte Technologies Private Limited publishes technical work under the KedByte Press imprint.

The editorial standard applied to this volume is simple to state and difficult to meet. Every explanation must work twice: once for a reader who has never opened a terminal, and once for an engineer who needs the exact figure. Every claim that can be checked must be checked. Every figure that will go stale must carry a date. Where the field genuinely disagrees, both positions are given rather than one being quietly chosen.

Where this book uses real evidence, it uses it unaltered. The network trace in Part F, the failed pushes in Part G and the error messages quoted throughout are reproduced exactly as they were observed. Diagnosis is taught against ambiguous real data, because that is the only kind there is.



KEDBYTE

TECHNOLOGIES PRIVATE LIMITED

COLOPHON

This first edition of *How Money Moves* runs to five volumes and fifty-four chapters, and was set entirely from plain text.

The body text is set in TeX Gyre Pagella, a digital revival of Hermann Zapf's Palatino, chosen for long-form reading. Headings, tables and the KedByte identity are set in Poppins. Code, command output and diagrams are set in DejaVu Sans Mono, and every diagram in this book is drawn in plain ASCII so that it survives being copied anywhere.

The book was composed with pandoc and typeset with XeLaTeX on A4 stock. The single accent colour used on rules, chapter numbers and section headings is KedByte navy.

Shikhar Singh

Founder & Chairman, KedByte Technologies Private Limited



First Edition • KedByte Press