PAYMENTS · DATA · AI · COMPLIANCE

Instant Payments (RTP and FedNow) Delivery and Support Runbook

From requirements to 24/7 L1, L2 and L3 support: RTP and FedNow basics, ISO 20022, exit criteria, sign-off, live pilot, real-time fraud controls, liquidity, alerts, disaster recovery and audit.

SP Sankar · Payment Product Consultant

sp@paymentconsultant.top · paymentconsultant.top · linkedin.com/in/spsankar

Version 2026 · Always confirm current rules with The Clearing House, the Federal Reserve and your own bank policies.

Contents

1. Purpose and how to use this runbook
2. Planning vs delivery alignment
3. Instant payments basics for the team
4. Rules and regulations
5. The nine kinds of work, step by step
6. Exit criteria and sign-off: E2E and UAT
7. Proving it in production
8. Controls that run across every payment
9. Security, certificates and certification
10. Production support: L1, L2 and L3
11. Operational runbooks for common incidents
12. Disaster recovery
13. Internal and external audit
14. Measures to watch
15. Go-live checklist
16. Sources

Purpose and how to use this runbook

This runbook takes an instant payments change, for example joining RTP or FedNow, adding a new use case such as payroll or Request for Payment, or upgrading the platform, from idea to live production and support. It is written for product managers, delivery leads, engineers, testers, operations and compliance teams.

Instant payments change two things for every team:

- **They never close.** Payments run 24 hours a day, 7 days a week, every day of the year. There is no overnight batch window, no weekend and no holiday to hide a release or a fix.
- **They are final in seconds.** Money is settled and available to the receiver almost at once. A wrong or fraudulent payment can't be pulled back; you can only ask the receiving bank to return it.

Golden rules for Agile instant payments delivery

- Work in sprints, but never skip the hard gates: certification with the network, E2E and UAT exit criteria, change approval and a live pilot.
- Design for zero downtime: rolling or blue-green releases, no "maintenance window" for customers.
- Fraud checks must finish in milliseconds and inside the network's time limit; test speed as hard as accuracy.
- Staff support for 24/7 from day one: on-call rota, weekend coverage and holiday cover.
- Watch liquidity: your settlement balance must cover payments at 3 a.m. on a Sunday too.

Planning vs delivery alignment

Instant payments projects carry fixed outside dates: network certification slots, RTP and FedNow release dates and your own bank's change freezes. Because the system never closes, plan release windows for the lowest-traffic hours and agree them with the network.

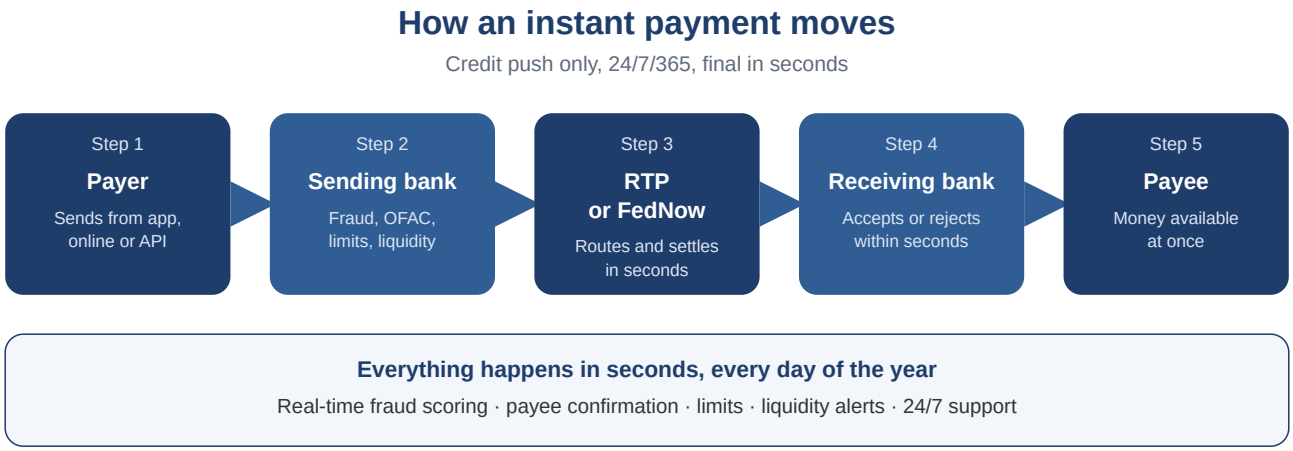
Level	What it holds	Who owns it	Reviewed
Product roadmap	Use cases (receive, send, Request for Payment), network release dates	Product manager	Quarterly
Release or PI plan	Features committed for the quarter, certification slots, risks	Product manager, delivery lead	Every planning increment
Sprint plan	User stories for 2 weeks	Product owner and team	Every sprint

Level	What it holds	Who owns it	Reviewed
Release calendar	Low-traffic release windows, freezes, network test windows	Release manager	Weekly

- **Receive first, then send:** most institutions go live receiving payments first, then add sending once fraud controls and liquidity are proven.
- **Plan to real capacity** with 15 to 20% buffer for 24/7 support work.
- **Track outside dependencies:** the network, your processor or service provider, core banking and fraud vendors.
- **Measure it:** say/do ratio, feature predictability, scope change and milestone status, reported weekly.

Instant payments basics for the team

The US has two instant payment networks run by different operators. Most large banks connect to both.



How an instant payment moves

> **How to read this picture:** Read left to right: each box is one step of a single instant payment. The whole journey takes seconds, day or night. The band at the bottom lists what must work in real time for this to be safe, because once the receiving bank accepts, the payment is final.

Topic	RTP network	FedNow Service
Operator	The Clearing House (owned by large banks)	Federal Reserve Banks
Launched	2017	July 2023
Hours	24/7/365	24/7/365

Topic	RTP network	FedNow Service
Transaction limit	$10 million (since February 9, 2025)	$10 million network limit (since November 12, 2025); each bank can set lower limits
2025 activity	More than $1.3 trillion in payments, up about 428% on 2024	About 8.4 million payments worth $853.4 billion, average $101,435
Participants	About 1,135 banks and credit unions (January 2026)	More than 1,500 financial institutions
Settlement	Prefunded joint account held at the Federal Reserve	Federal Reserve master accounts
Message format	ISO 20022	ISO 20022

Sources: The Clearing House press releases (January 2025, February 2025, December 2025); Federal Reserve Financial Services announcements and 2025 statistics.

How a payment works: both networks are credit push only. The payer's bank sends the money; nobody can pull money from your account. The receiving bank must accept or reject within seconds, and once accepted the money is final.

ISO 20022 messages your team must know

Message	Use
pacs.008	Credit transfer (the payment)
pacs.002	Payment status: accepted or rejected, with reason
pain.013 / pain.014	Request for Payment and its response
camt.056	Request to return a payment (recall)
camt.029	Response to a return request
pacs.004	Return of funds
camt.035 / admi messages	Acknowledgements and system messages

Exact message versions differ by network; build to each network's published specifications.

Rules and regulations

Rule	What it means for delivery
RTP Operating Rules	Rules for RTP participants, including timing, returns and Request for Payment
Regulation J, Subpart C and Operating Circular 8	Rules for payments sent through FedNow
UCC Article 4A	Governs the funds transfer for these credit push payments, as applied by the network rules
Regulation E	Consumer protection for consumer accounts: error resolution and unauthorized transfers
OFAC sanctions	Screen payments in real time before accepting or sending
BSA/AML	Monitor for suspicious activity, including money mule accounts, and file SARs
UDAAP and complaints	Clear disclosure that instant payments are final and can't be cancelled

The nine kinds of work, step by step

These are not waterfall phases. In sprints they overlap; the gates are what matter.

Phase 1: Requirements

- Use cases: receive only, send, Request for Payment, business payouts, payroll, account-to-account transfers.
- Limits by customer type and channel, below the network limit.
- Fraud rules for sending: new payee, first payment, unusual amount or time, mule account signals.
- Customer experience: clear "this is final" warning, payee name shown before sending, instant confirmation.
- 24/7 posting: core banking must post and show the balance in real time.

Exit: requirements signed by product, operations, fraud, compliance and treasury (liquidity).

Phase 2: Design

- Real-time flow: request, validation, fraud and sanctions checks, liquidity check, send, status, posting and notification.

- Time budget: every step from receive to accept or reject must finish well inside the network's time limit.
- Liquidity: how the settlement account is funded and monitored 24/7, including weekends when Fedwire is closed.
- Idempotency and duplicate checks; a unique end-to-end ID per payment.
- High availability: active-active or fast failover; no single point of failure.

Exit: architecture, security and resilience reviews passed.

Phase 3: Development

- Build to the RTP and FedNow ISO 20022 specifications and connection guides.
- Real-time core posting (or a shadow ledger) with clear rules for when the customer can use the money.
- Code review, static analysis and dependency scanning on every change.

Phase 4: Unit testing

- Message building and parsing for each message; validation and reason codes.
- Limits, fraud rule triggers, timeouts and retries.

Phase 5: Performance testing

- Response time at peak (for example payday and tax refund days) with the full fraud and sanctions stack switched on.
- Sustained load over hours, not minutes; memory leaks show up on day 3, not hour 1.
- Failover under load: no lost or duplicated payments.

Exit: 99th percentile response time inside the agreed budget at 2 times expected peak.

Phase 6: End-to-end (E2E) testing

- Send and receive in the network test environment, including rejects, timeouts, returns, return requests and Request for Payment.
- Late responses: what happens if your accept arrives after the network timeout.
- Liquidity alerts, posting, statements and notifications.

Phase 7: User acceptance testing (UAT)

- Operations runs real cases: investigation, return request, fraud claim, customer complaint.
- Customer journeys in every channel, including weekend and night tests.

Phase 8: Release management

- Zero-downtime release plan; agree low-traffic windows; rollback that never loses an in-flight payment.

- Network notified if required for planned changes.

Phase 9: Deployment

- Deploy, smoke test, then a live pilot (see "Proving it in production").
- Hypercare: extra on-call cover for the first weeks, including weekends.

Exit criteria and sign-off: E2E and UAT

E2E exit criteria

- [ ] All network certification test cases passed and accepted by the network
- [ ] 100% of planned test cases run; at least 95% passed; no open Severity 1 or 2 defects
- [ ] Response times inside the agreed budget at peak load
- [ ] Every message type sent and received, including rejects, returns and return requests
- [ ] Real-time posting and reconciliation balanced
- [ ] Fraud and sanctions scenarios behave as designed

UAT exit criteria

- [ ] Business scenarios signed by the business owner
- [ ] 24/7 operations procedures tested, including a weekend shift
- [ ] Customer warnings, confirmations and notifications checked

Sign-off matrix

Who signs	What they confirm
Product owner	Features meet acceptance criteria
QA lead	Exit criteria met, evidence stored
Operations head	24/7 team trained and staffed
Treasury	Liquidity funding and alerts ready
Fraud and compliance	Fraud, OFAC and Reg E requirements met
Information security	Security testing passed
Change manager	Release approved

Proving it in production

1. Receive a small live payment (for example $1) from an account you control at another institution; confirm it posts in seconds.
2. Send a small payment to an account you control; confirm the accepted status and the credit at the other side.
3. Send a return request and complete a return.
4. If in scope, send and pay a Request for Payment.
5. Repeat on the second network if you use both.
6. Check liquidity balances and reconcile all pilot payments.
7. Run the pilot at night or on a weekend too, to prove 24/7 support.

Exit: all pilot payments settled and reconciled; sign-off by product, operations, treasury and compliance.

Controls that run across every payment

Control	What good looks like
Real-time fraud scoring	Every outgoing payment scored before sending; risky ones held or blocked
Payee confirmation	Show the payee name or confirm the account before the customer sends
Limits	Per payment, per day and per customer, below the network limit
Mule account detection	Watch incoming payments for accounts that receive and quickly send money out
OFAC screening	Real time, inside the response time budget
Liquidity monitoring	Alerts when the settlement balance falls below a safe level
Customer warnings	Clear message that the payment is final and can't be cancelled
Audit trail	Every request, decision and status logged with time

Authorized push payment scams are the main fraud risk: a criminal tricks the customer into sending money themselves. Strong warnings, payee confirmation and real-time scoring are the defence. The

Clearing House reported that in April 2025 only 123 fraud cases were reported out of 35 million RTP payments, but each case is final, so prevention matters.

Security, certificates and certification

- **Access:** least privilege, quarterly reviews, PAM for privileged access, NIST SP 800-63B for passwords and multi-factor authentication.
- **Network connection:** FedLine for FedNow; the RTP connection for RTP. Track every certificate and key expiry; alert 60, 30 and 7 days before.
- **Certificate renewal with no downtime:** renew in a planned change with both old and new certificates valid, because there is no closed window.
- **Network certification:** both networks require certification testing before go-live and after major changes. Keep evidence.
- **Vendors, licences and open source:** track versions, support dates and licences; scan dependencies on every build.

Production support: L1, L2 and L3

Level	Who	What they do	Target
L1	24/7 operations and service desk	Monitor dashboards, answer customers, first triage	Acknowledge within 5 to 15 minutes, day or night
L2	Application support (on call)	Diagnose, restart, fail over, manage incidents	Respond within 15 to 30 minutes
L3	Engineering and vendor (on call)	Code fixes, hotfixes, root cause	Engaged within 1 hour for Sev 1

Severity levels

Severity	Example	Response
Sev 1	Can't receive or send, network connection down, rising timeouts	Bridge within 15 minutes, any time of day
Sev 2	Slow responses, one channel down, liquidity near threshold	Within 30 minutes
Sev 3	Single payment issue	Next business day

Alerts and notifications

- Real-time alerts to Slack or Teams and paging for: connection loss, reject or timeout rate above normal, response time above budget, liquidity below threshold, certificate expiry, posting failures.
- Email updates to stakeholders for Sev 1 and Sev 2 every 30 minutes.
- Tell the network operator about any outage that affects your ability to receive.

Escalation: L1 to L2 at 15 minutes; L2 to L3 and the incident manager at 30 minutes; senior leaders for any Sev 1.

Hotfixes: emergency change with a named approver, peer review, test in a lower environment, zero-downtime deploy and a review within 5 business days.

Operational runbooks for common incidents

1. Rising timeouts or rejects

1. Check response time at each step (fraud, sanctions, core posting).
2. If one component is slow, fail over or switch to its backup.
3. If you can't recover quickly, consider signing off from receiving (as the network allows) rather than timing out on customers.
4. Reconcile any payment with an unclear status before retrying.

2. Liquidity running low (for example on a weekend)

1. Alert treasury on call.
2. Move funds into the settlement account as the network allows.
3. If needed, slow or pause sending until funded.

3. Customer reports a scam payment

1. Log the claim and send a return request to the receiving bank at once.
2. Report to fraud, file a SAR if needed, and follow Reg E for consumer accounts.
3. Check whether the receiving account looks like a mule account across other payments.

4. Duplicate or wrong payment sent

1. Confirm the status and end-to-end ID.
2. Send a return request; do not send a second payment to "fix" it.
3. Fix the cause before the next release.

Disaster recovery

Item	Example target
Recovery time objective (RTO)	Minutes, not hours: an active-active design or automated failover
Recovery point objective (RPO)	Zero: no accepted payment may be lost

- [ ] Second site or zone connected and certified with each network
- [ ] Automatic failover tested under load
- [ ] Liquidity and treasury procedures work from the backup site
- [ ] DR test at least once a year, with evidence
- [ ] After failover, reconcile every payment in flight

Internal and external audit

Internal audit checklist

- [ ] Instant payments policies and procedures current and approved
- [ ] Network agreements and certifications on file
- [ ] Limits set, reviewed and enforced by the system
- [ ] Real-time fraud and OFAC screening working (sample payments)
- [ ] Return requests and fraud claims handled on time; Reg E met for consumer accounts
- [ ] Liquidity alerts and funding procedures tested
- [ ] Access reviews quarterly; no direct database changes without approved tickets
- [ ] 24/7 rota staffed and incidents reviewed
- [ ] DR test in the last 12 months

External audit and exam checklist

- [ ] Evidence index linking each control to its evidence
- [ ] Payment population lists ready to sample
- [ ] Change records with approvals, certification and pilot evidence
- [ ] Access lists and review sign-offs
- [ ] Incident records and root cause reviews
- [ ] BSA/AML and OFAC policies, training and independent testing
- [ ] Vendor and service provider SOC reports
- [ ] Open findings tracker with owners and dates

Measures to watch

Measure	Why it matters
Response time (median and 99th percentile)	Slow responses become timeouts and failed payments
Reject and timeout rate	Early sign of system or data problems
Availability (24/7)	Instant payments are expected to always work
Fraud rate and scam claims	The main customer harm
Liquidity low points	Risk of not being able to send
Return requests and outcomes	How often money must be chased back

Go-live checklist

- [ ] Network certification complete
- [ ] E2E and UAT exit criteria met and signed
- [ ] Live pilot payments settled and reconciled, including a night or weekend test
- [ ] Fraud scoring, payee confirmation and limits switched on
- [ ] Liquidity funding and alerts live
- [ ] 24/7 on-call, alerts and escalation tested
- [ ] DR failover tested
- [ ] Customer warnings and disclosures live

Sources

- The Clearing House: RTP network press releases on 2024 and 2025 results, the $10 million limit and fraud data
- Federal Reserve Financial Services: FedNow transaction limit increase (November 12, 2025) and 2025 statistics
- Regulation J Subpart C and Operating Circular 8; RTP Operating Rules; UCC Article 4A; Regulation E; OFAC