
FREE DOWNLOAD · INCIDENT RESPONSE

The Incident Response Handbook

What the first days of a breach actually demand, compiled from the documented guides at vendorquals.com: choosing responders under fire, the first 24 hours, ransomware response and negotiation, command and the board briefing, every notification deadline on one page, and the day after. Evidence-first, no marketing.

VendorQuals · vendorquals.com · October 2026

THE INCIDENT RESPONSE HANDBOOK

Contents

What the first days of a breach actually demand, in the order the incident demands it. Chapters 1 through 4 are the response itself: the first hour, the responder choice, ransomware, and negotiation. Chapters 5 through 7 are the structures around it: command, the notification clocks, and the words. Chapter 8 is the day after. Chapter 9 is what the incident should change. Every rule here traces to the same sources: claims processes, privilege doctrine, notification statutes, and NIST SP 800-61.

1. Hour zero: the moves that matter first	3
2. Choosing a responder in hours, not weeks	5
3. Ransomware: the first twenty-four hours	7
4. The ransom question: negotiation is its own skill	9
5. Command: one named person	11
6. The clocks: who must be told, and when	13
7. Words: the board, the customers, the press	15
8. The day after: the post-incident review	17
9. Turning panic into practice	19
Where to go next: the free downloads and the guides behind this handbook	20

CHAPTER ONE

Hour zero: the moves that matter first

Contain without destroying, preserve before you remediate, and call in the right order. Everything else in this handbook depends on this hour being done right.

The first hour decides what the investigation, the insurance claim, and the regulator report can ever say. Machines can be isolated and evidence preserved, or evidence can be destroyed by well-meaning cleanup. Every rule in this chapter exists because someone, under pressure, did the destructive thing with good intentions.

Log the discovery time to the minute, first. Every notification clock in chapter 6 keys off discovery or awareness, and a discovery time reconstructed later from memory is contested in a way a timestamped one is not. Start the incident log at the same moment: every action, decision, and time goes in it, in one place, from the beginning.

Contain without destroying

Isolate affected systems at the network level. Disconnect segments, disable compromised accounts, block command-and-control addresses. If your platform allows it, snapshot encrypted machines before doing anything else. Do not power off encrypted systems if you can avoid it: volatile data in memory is evidence, and powering off destroys it. Do not reimage, restore, or clean anything before forensics. The sentence 'we reimaged it' damages the investigation, the insurance claim, and the regulator report at the same time, and none of the three can be undone.

The call order

Three calls, in this order, each faster than a vendor search:

Order	Call	Why
1	Your carrier's breach hotline	Panel rules decide who gets paid and how much. Panel firms can usually start under the claim immediately. Engaging off-panel without pre-approval is how buyers pay twice.
2	Counsel	Response runs under privilege when counsel directs it, and on most panels the breach lawyer formally engages the responders. Evidence gathered without privilege can end up in the regulator's hands in full.
3	A responder	Your retainer firm, or a panel firm through the claim. Chapter 2 is how to choose one in hours. The clocks in chapter 6 run whether or not a responder is on site.

The incident log

One document, one owner, continuous entries: what was observed, what was decided, by whom, at what time. The log is the spine of everything that follows: the board briefing in chapter 7, the notification filings, the claim, and the review in chapter 8. If it is not in the log, it did not happen.

THE TWO UNRECOVERABLE MISTAKES

Remediation before forensics, and off-panel engagement without the carrier's pre-approval. Both are made by competent people under pressure, both are common, and neither can be undone after the fact. Every other mistake in this handbook is recoverable.

CHAPTER TWO

Choosing a responder in hours, not weeks

Mid-incident you run three gates, not six checks: mobilization with named people, environment fit, and documentation discipline. Everything else waits for the day after.

The full evaluation of an incident response firm has six checks, and that evaluation is for peacetime. Mid-incident you run three gates, in order, and you can run them in a single day of calls. If someone is already running incident command (chapter 5), that person runs the gates: whoever is choosing responders should not also be fighting fires.

Gate 1: Can they mobilize, with named people

- A named engagement lead. A real person who will be on your incident, not a salesperson and not 'our team.'
- Hours to first action, in writing, today. Hours to acknowledge, hours to remote start, hours to on-site if you need boots. 'Immediately' without numbers is enthusiasm, not an SLA.
- The concurrency answer. 'What happens if you are working three other incidents when I invoke?' The honest answers involve surge bench, priority terms, or a retainer queue. Silence on this question is a red flag at 2 a.m. and a disqualifier by morning.
- If you hold a retainer, invoke it now. Pre-agreed response terms are the thing it exists for.

Gate 2: Do they know your environment

A responder who has never worked in your environment type spends the first billable day learning it. The failure modes are not subtle:

- OT/ICS: containment means production trade-offs. A firm without ICS experience may isolate a segment that stops a production line, or miss the engineering workstation that matters.
- Cloud and cloud-native: the evidence lives in API logs, audit trails, and cluster state that evaporate. Standard endpoint forensics on a Kubernetes incident is archaeology in the wrong ruin.
- Regulated sectors: in medtech and automotive environments, the response has to anticipate the regulator's questions, not just the attacker's.

One question exposes fit: 'What incidents like ours, in environments like ours, have your actual responders worked, and which can you describe without an NDA?' Firms with documented engagements answer fast.

Gate 3: Will the work survive the claim and the regulator

The audience for the response is larger than you: the carrier's adjusters, possibly a regulator, possibly a court. Two checks, asked fast:

- Documentation discipline. 'Show us a sanitized report structure now.' You are listening for a timeline with evidence citations, scope separated from speculation, and prose a non-technical executive can read.

- Chain of custody and privilege. Confirm they work under counsel's direction and preserve evidence to a defensible standard. Firms that shrug at this mid-incident will shrug at it in the deposition too.

The first-call script

Five questions, every firm, same order. Twenty minutes across three calls makes the choice comparable when you are least able to compare:

1. Who is the named lead on our incident, and which responders are with them?
2. How many hours from a signed engagement to your first action, and what does first action mean?
3. What other incidents are you working right now, and what happens to our priority if they escalate?
4. What incidents in environments like ours can you describe without an NDA?
5. What do you need from us in the first four hours, and what will you absolutely not do without our approval?

The last question is the tell. Strong firms answer with evidence handling, isolation approvals, and communication protocol. Weak firms answer with 'we'll take it from here,' which is exactly wrong: it is your incident, your evidence, and your regulator.

What to skip on purpose

- The full RFP. Structured procurement is for peacetime. Emergency engagements need three gates, not a scoring matrix.
- Reference calls. No one you can reach at 3 a.m. is a useful reference. Named responders and documented engagements substitute.
- Rate negotiation to the third decimal. Get the rate structure in writing, flag disputes, negotiate the delta after containment. The exception: minimums and out-of-scope definitions must be stated before work starts, because those are where emergency invoices surprise buyers.
- Security questionnaires. Vendor risk reviews are a procurement artifact, not a response tool.

Red flags at 2 a.m.

- 'Just restore from backup and we'll check it after.' Remediating before forensics destroys the record of what happened and what left.
- No named humans by the end of the first call. Teams that cannot produce a lead before engagement will not produce one during.
- SLA by adjective. A commitment with an asterisk about concurrent incidents is not a commitment.
- Panel discouragement. Any firm that talks you out of calling your carrier's hotline is volunteering you for the bill.
- The retainer ambush. Your emergency is not the moment to buy a retainer as a condition of mobilization. Buy it later, deliberately (chapter 9).

CHAPTER THREE

Ransomware: the first twenty-four hours

Contain the spread, preserve the evidence, find out what you actually lost. Only then decide what the attack costs you, in money, data, or downtime.

Everything about ransomware response is sequencing. The payment question cannot be answered until scope, backup state, and exposure are established, and all three take hours to establish honestly.

Hours zero to four: stop the spread

1. Isolate at the network level. Disconnect affected segments, disable compromised accounts, block command-and-control addresses. Snapshot encrypted machines before anything else if the platform allows it.
2. Do not power off, reimage, or restore yet. The evidence requirements of your claim and any regulatory defense are decided in these hours.
3. Identify the variant. Ransomware families carry known behaviors, known flaws, and occasionally known decryptors. Identification changes recovery options immediately.
4. Check your offline, immutable backups first. The most important question in hour two: what exists that the attacker could not reach?

Hours four to twelve: know what you lost

- Scope the encryption and the exfiltration separately. Modern ransomware is a two-part event: systems locked, data stolen. The stolen data, not the locked systems, usually drives notification obligations and the true exposure.
- Preserve the ransom note and any attacker infrastructure. They carry version and timeline data for the investigation and the negotiation record.
- Engage the response path: carrier hotline, counsel, responder (chapters 1 and 2).
- Brief leadership on facts, not fears. The board gets knowns and unknowns (chapter 7), never a cost promise.

Hours twelve to twenty-four: the decision inputs

The payment decision is made on evidence, and the inputs are all establishable within a day:

1. Recoverability. What do working backups cover, and what downtime gap remains?
2. Exposure. What was stolen, whose data is it, and what do notification laws require you to tell, and when?
3. Negotiation posture. Attackers price to perceived ability to pay and urgency. The first demand is an opening position, not a finding (chapter 4).
4. Sanctions screening. Payments to sanctioned actors or wallets are illegal regardless of the incident. This check is mandatory before any payment conversation advances.

Only after these inputs exist does the business make the payment, partial recovery, or full recovery call, with counsel in the room and the decision documented.

What not to do

- Do not pay in the first hours. Early payment is the most expensive version of every ransom: unestablished scope, unchecked backups, maximum fear at maximum urgency.
- Do not let IT clean up before forensics. Restoration before imaging destroys the record of both the attack and the exfiltration.
- Do not email the attacker from your own mailbox. It identifies decision-makers, signals urgency, and locks in a bad anchor.
- Do not promise the board a number. The cost estimate of hour four is never the cost of hour twenty-four.

CHAPTER FOUR

The ransom question: negotiation is its own skill

Talking with a ransomware crew is closer to hostage negotiation than to IT work. Most IR firms do not have a real negotiation practice. Find out before the ransom note arrives.

Most incident response firms can help contain a ransomware incident on the technical side. Far fewer have a real, trained negotiation practice. The work takes knowing how attacker groups behave, payment mechanics including crypto handling and OFAC sanctions checks, pressure management, and when the right move is not to pay at all. Treating negotiation as a standard add-on to any IR engagement is a common and costly buyer mistake, usually discovered after the ransom note arrives.

What a negotiation firm actually does

- Opens and runs the talk with the attacker. A named negotiator, not your own staff, which lowers the odds of a nervous employee handing the attacker more power.
- Checks the attacker's claims. Was data truly taken, or just claimed? Does the unlock tool actually work? Both are verified before payment talk gets serious.
- Runs sanctions checks. Many ransomware groups sit on OFAC lists. Paying a listed group carries legal risk regardless of the excuse, and a firm with a real compliance arm checks this before payment is on the table.
- Haggles the amount and terms if payment becomes the path. First demands are often open to talks, sometimes by a lot.
- Works with police and lawyers throughout. Payment calls bring regulatory and disclosure duties beyond the technical side.

The choice to pay is not a technical call

Paying or not paying is a business, legal, and often regulatory decision. A negotiation firm's job includes helping you make it with true facts: verified claims, sanctions status, real recovery-without-payment timelines. A firm whose pitch is all 'we get your data back fast,' with no word on verification or sanctions, is skipping the steps that protect you from paying for nothing.

Questions to ask before you need this

- Does your retainer include ransomware negotiation itself, or is it a separate service with separate terms?
- How many ransomware negotiations has your team run in the past 12 months, and with which attacker groups?
- What is your sanctions check process, and how fast can you run it under time pressure?
- Have you worked with the specific ransomware family we might face, or is your method generic across strains?
- If payment is a bad idea, what is your process for saying so, and does advising against payment cut against any fee you earn from making one?

Red flags

1. No named negotiator. The person doing the technical work is also doing the talking. Those are different skills.
2. No mention of sanctions checks. This is not optional homework.
3. Pressure to pay fast without proof. Proof of data theft and unlock-tool testing are the steps that protect you from paying for nothing.
4. No stated experience with your specific attacker or strain. Different groups have different records of honoring deals; a firm without that knowledge is negotiating blind.

CHAPTER FIVE

Command: one named person

One person holds overall responsibility for the on-scene incident. Not a committee. Urgent calls need a named owner, which is the entire point of the Incident Command System.

Who has overall responsibility for managing the on-scene incident? One named person: the incident commander. The principle comes from FEMA's Incident Command System, the emergency management method taught in its ICS-100 course, and cyber teams adopted the role on purpose. An incident has one commander at any moment. When a shift ends, command passes to a named relief person with a written handoff. A committee never runs the show in the moment.

What command means

ICS teaches the incident commander's duties in order. First, team safety. Second, goals. Third, the plan and moves to reach them. Fourth, the people and tools to do the work. Safety comes first on purpose: a response that hurts people has already failed. Command is also not technical skill. The commander picks what to do and when; experts pick how inside their own fields. Mixing the two is how incidents stall.

Why one person

Incidents move faster than committees. Someone must own the hard calls in real time: commit people, escalate, stand down, accept risk. A group cannot hold that power in the moment. Committees advise between incidents; commanders decide during them. The same idea shows up in NIST SP 800-61 and in every serious incident response plan: the plan names a commander, spells out what they may decide without asking anyone, such as cutting off a network segment or calling the retainer firm, and names their backups. The plan gives the power ahead of time. The commander uses it.

The ICS idea	In a cyber incident
One incident commander	One named response lead with real power during a declared incident
Written transfer of command	Shift handovers with a written status note. Long incidents fail at bad handoffs
Team safety first	In cyber: legal risk, privilege, evidence handling, and not breaking the business while fixing it
Command staff vs workers	The commander steers. Technical experts do containment and forensics
Pre-approved calls	The plan's pre-approvals, like isolation power or counsel calls, are the tools the commander commands with

Who should not command

Not the person deep in the technical work. A forensic lead buried in a disk image cannot also weigh business trade-offs. The commander does not need to be a security expert; they need command skill: running the decision loop, trusting experts, deciding under pressure. In cyber that usually means security fluency, but the core skill is deciding, not disk forensics.

Outside firms and insurers

A retainer firm joins operations under the commander's steering. Your people know your systems; the firm brings deep response skill. The plan should state how they hand off work before the incident, not during it. The carrier's panel rules (chapter 1) say which firms need no approval first. The commander steers both.

CHAPTER SIX

The clocks: who must be told, and when

The clocks stack: one incident can be on five calendars at once. They start at discovery or awareness, not at the end of your investigation.

The clocks run whether or not a responder is on site, and they run in parallel. The table below is the complete set a US-headquartered organization typically faces. Print it, date it, and verify current rules with counsel before relying on any of them.

Regime	Deadline	Clock starts	Goes to
GDPR (EU/UK residents)	72 hours to the authority	Awareness of the breach	Lead supervisory authority; individuals without undue delay if high risk
HIPAA (protected health info)	60 days to individuals	Discovery of the breach	Individuals; contemporaneous log; HHS annually; media if 500+ in a state
US state breach laws	Most: without unreasonable delay. 30 days in CO, FL, ME, ND, WA; TX 60 days at 250+ residents	Discovery (statute-dependent)	Affected residents; state AGs above thresholds (commonly 250 or 500)
SEC Item 1.05 (listed companies)	4 business days (Form 8-K)	Determination the incident is material	SEC filing, via disclosure counsel
Banking (interagency rule)	36 hours	Determination a notification incident occurred	Primary federal regulator (OCC, FDIC, Federal Reserve)
NYDFS Part 500	72 hours	Qualifying cybersecurity event	NY Department of Financial Services superintendent
FTC Safeguards Rule (non-bank finance)	30 days	Discovery; unencrypted info of 500+ consumers	The FTC directly
SEC Regulation S-P (brokers, advisers)	30 days to individuals	Discovery of unauthorized access to sensitive info	Affected individuals, per the rule's content requirements
State insurance (NAIC model law)	72 hours (varies by state)	Determination a cybersecurity event occurred	State insurance commissioner
Contractual notices	Per contract	Per contract terms	Customers, partners, counterparties (missed notices pay real damages)

Three things people get wrong

- GDPR starts at awareness, and awareness means having a reasonable degree of certainty that a security incident has occurred that led to personal data being compromised. That standard can be met before the forensic investigation ends.
- The SEC's Item 1.05 clock starts at a determination that the incident is material, which is a disclosure counsel decision, not a security team decision.
- HIPAA runs a contemporaneous log: every breach goes in it as it is discovered, with notification to individuals within 60 days, and the log itself is examinable.

CONTRACTUAL CLOCKS

Customer and partner contracts carry their own notification deadlines, often 24 to 72 hours, and missed contractual notices pay real damages. Pull the notification matrix out of your contracts during peacetime, not mid-incident.

CHAPTER SEVEN

Words: the board, the customers, the press

An incident runs on two tracks. The technical one decides the technical cost. The words decide the reputational one.

One track is technical: contain, investigate, recover. The other is words: what the board hears, what customers are told, what the press is answered with, and what regulators read. The second track decides the reputational cost as surely as the first decides the technical one.

The board briefing

The board's first briefing should arrive within hours of discovery, and it has a strict shape:

1. What is known. Discovery time, systems affected, response engaged (who, under what terms), containment status. Facts with timestamps from the incident log.
2. What is not yet known. Scope of data affected, whether exfiltration occurred, attacker identity. Naming unknowns is not weakness; it is how the board calibrates confidence in everything else you say.
3. What happens next. The response plan for the next 24 hours, the decision points coming (notification triggers, engagement terms, ransom decisions if applicable), and what the board will be asked to decide versus merely know.

What the board needs from the first briefing is the ability to govern, not technical mastery: who commands the incident (chapter 5), what money is authorized, what legal exposure counsel is managing. Give them a cadence, hourly for the first day, then a rhythm the incident sets, and never let a board member learn about the incident from the press.

One voice, in order of obligation

- Regulators first where clocks demand it (chapter 6). Notification content is a legal document; counsel signs off.
- Affected customers honestly: what happened, what data, what you are doing, what you will do for them, how to reach you.
- Staff before press. Employees who learn details from journalists become a leak surface.
- Press through the one voice. 'We are investigating and will update' is a complete statement.

The discipline rules

Every statement is written, reviewed by counsel, and logged with its time. Nothing goes out that the technical team has not confirmed. No executive posts individually. What you say in writing can be discovered years later, which is why crisis communications run alongside breach counsel, not around them.

What not to say, at any stage

- No scope promises before confirmation. 'Only three systems' and 'no data affected' are the two sentences that age worst in incident response.
- No blame narratives. Blaming the attacker, a vendor, or an employee in public creates legal exposure and burns goodwill you will need later.
- No unauthorized confirmations. Press 'confirming' a detail is not a reason to confirm it.
- No silence as strategy. Saying nothing while the story forms elsewhere is itself a communication, usually read as concealment.

THE PANEL FACT MOST BUYERS MISS

Carrier-published breach panels are not only forensics and law. VendorQuals panel research documents 13 crisis communications relationships across the published panels, read October 2026: FleishmanHillard on three carriers' panels, Edelman and FTI Consulting on two each, plus single-panel firms. If you carry cyber insurance, professional crisis communications may be a claim-covered call. Call the hotline before hiring.

CHAPTER EIGHT

The day after: the post-incident review

The review examines the response, not just the intrusion: detection, escalation, decisions, communications, and the plan's promises versus what happened.

The incident is contained, systems are restored, notifications are out. The post-incident review is where the organization actually learns, and NIST SP 800-61 treats it as a phase of the response, not a courtesy. Organizations that skip it, or run it as a blame exercise, leave the plan exactly as wrong as it was.

When, and who

Hold it within two weeks of closure: long enough that the fires are out, short enough that memory and logs are fresh. Everyone who made or escalated a decision attends: security, IT operations, legal, communications, and the business owner of the affected system. The executive sponsor attends or explicitly receives the output. The leadership question matters: internal facilitators carry internal politics into the room, external facilitators arrive without a stake in who looks good, and when breach counsel directs the review the work product can sit under privilege, which changes how honestly people speak. Decide who leads, and under whose direction, before the invitations go out.

The structure that works

Run it in this order and timebox the whole thing to two or three hours. Longer reviews rehash; they do not reveal.

1. Timeline reconstruction. One shared timeline from first detection to closure. Disagreements about the timeline are findings, not noise: if two people remember the escalation differently, the escalation path is ambiguous.
2. Detection analysis. Time from initial compromise to detection, and detection to escalation, compared against targets. If detection took three weeks, the finding belongs to monitoring, not the response team.
3. Decision audit. Revisit the hard calls: containment scope, shutdown decisions, ransom posture, notification timing. The question is never 'would you decide differently now.' It is whether the plan pre-authorized the call, and whether the decider had the authority and information the plan promised.
4. Communications audit. Which statements held, which aged badly, which audiences heard late, how much rework the facts required.
5. Control and plan changes. The concrete list: detections to deploy, access to revoke, contracts to renegotiate, plan sections to rewrite, drills to schedule.

The blame question

Examine decisions, not people. Every action is assessed against what the actor knew, the authority they had, and what the plan said. A responder who improvised outside the plan is not the finding; the plan's failure to authorize or forbid the action is the finding. The exceptions are process violations: bypassed evidence handling, deleted data against counsel's instruction. Those get named, because a review that

cannot name them protects the wrong thing.

The deliverables

- The action list. Every finding with an owner and a due date, tracked like the project it is. The only deliverable that changes anything.
- The board summary. One page: what happened in business terms, what it cost, what changes now.
- The plan update. The response plan gets its revisions and a new version date. The review is also the natural moment to revisit retainer terms and panel access while the gaps are fresh.

The next drill should test at least one change the review produced. That closes the loop: incident, review, change, drill, and a plan that is measurably harder to break than it was in the last cycle.

CHAPTER NINE

Turning panic into practice

The emergency taught you things no RFP could. Schedule the real evaluation, then make the next 2 a.m. a checklist instead of a decision.

Once containment is stable and responders are working, schedule the real work: converting the emergency into preparation. The incident taught you things no request for proposal could: whether the firm's bench was real, whether its documentation discipline held under pressure, whether its terms matched what you were told on the first call.

Run the full six-check evaluation against the firm that just responded, and against two alternatives, then decide the retainer question with evidence instead of adrenaline. A retainer buys pre-agreed response terms and queue priority; whether that premium is worth it is a risk decision you can now make with incident-grade information.

The three peacetime artifacts

1. The plan. A written incident response plan with roles, severity tiers, a contact tree, and the notification matrix from chapter 6 pre-filled. The fill-in template this handbook ships alongside is a starting point.
2. The drill. A tabletop exercise that rehearses the communications track (chapter 7) as seriously as the technical one, so the next first hour runs on muscle memory.
3. The retainer. Pre-agreed terms with a firm that passed the gates in chapter 2, in writing, before the next 2 a.m.

And if this incident ended without a responder engaged, because it was small or the panel handled it, you just received the cheapest drill money can buy. Run it as one. The difference between organizations that survive breaches and organizations that are shaped by them is rarely the first incident. It is what the first incident changed.

APPENDIX

Where to go next

This handbook condenses guides published at vendorquals.com, where each topic carries its full reasoning and its sources. Everything below is free, with no signup.

Download	What it is	Where
Breach Notification Clocks cheat sheet	The table in chapter 6, alone on one printable page	vendorquals.com/resources
The First 24 Hours cheat sheet	Chapters 1 and 3 condensed to one printable page	vendorquals.com/resources
Incident Response Plan template	The fill-in plan from chapter 9, as a DOCX	vendorquals.com/resources
The Breach Cost Calculator	HIPAA and GDPR penalty exposure, calculated per record and per tier	vendorquals.com/tools/breach-cost-calculator

The guides behind this handbook

- You're breached right now: choosing a responder. vendorquals.com/blog/breached-right-now-choosing-a-responder
- How to choose an incident response firm (the full six-check evaluation). vendorquals.com/blog/how-to-choose-an-incident-response-firm
- Ransomware response: the first 24 hours. vendorquals.com/blog/ransomware-response-first-24-hours
- Ransomware negotiation. vendorquals.com/blog/ransomware-negotiation-guide
- The incident commander, explained. vendorquals.com/blog/incident-commander
- State breach notification laws (the matrix method). vendorquals.com/blog/state-breach-notification-laws
- Crisis communications during a breach. vendorquals.com/blog/crisis-communications-during-a-breach
- The post-incident review. vendorquals.com/blog/post-incident-review-guide

VendorQuals

VendorQuals is the evidence-graded incident response firm directory: firms documented from registry-checked, panel-documented, and research-sourced evidence, filterable by specialization, carrier panel status, and response time. Browse it at vendorquals.com/firms. No firm paid to be included anywhere in this handbook. Have a correction or addition? vendorquals.com/corrections

This handbook is operational guidance, not legal advice. Notification rules change; verify current requirements with counsel before relying on them. Compiled October 2026.