AI agents inside customer service

Who answers, who sends, and what an agent may change

IN BRIEF

01 Customer service is the department where an agent's output reaches a customer. An answer about a fare. A reply about a refund. The company is held to it. That is why the question to ask first is not where an agent runs. It is who answers, who sends, and what the agent may change.

02 The public material answers that question differently. One vendor documents two modes side by side: one "sends the email automatically", the other "creates a draft for a customer service representative to review and send". One supplier charges by the outcome and says its agent "processes payments and refunds". Another vendor's page says of its drafts: "Nothing sends until they approve it." Hosting separates them less: at least two of the supplier pages we read offer deployment on the customer's own infrastructure.

03 The record explains why the first question matters. A tribunal in Canada held an airline to what the chatbot on its website told a passenger. A software company's support bot told users about a policy that did not exist, and the company's co-founder had to correct it in public. One energy supplier has published the opposite practice: every AI-written email is labelled, and complaints and vulnerable customers stay with staff.

04 This paper is about the service department's own work, from first answers and email to complaints, refunds, calls and quality. It sets out what an agent can read, how it reaches service systems, what it may do without a person, and what the rules already say. It ends with seven questions to put to any supplier, and with a place to start.

05 The position throughout is simple. An agent prepares. A named person decides and sends. No agent issues a refund, a credit or compensation. None changes an account, an order or a contract, closes a ticket or a complaint, or sends a reply about a customer's own case. Two agents answer general questions directly, inside a fence: published and approved help articles only. They say they are AI, and an adviser takes over on request. No agent scores an adviser or reads mood. The agent runs on infrastructure the company controls.

IN THIS PAPER

01 Where customer service is today 2
02 The work nobody has packaged 3
03 Three doors, one gate 4
04 What an agent may do 6
05 What the rules already ask 7
06 A worked example: one home appliance retailer's service team 9
07 Seven questions before an agent talks to a customer 10
08 Where to start 10

PART 01

Where customer service is today

$812.02

in damages, interest and fees, which a Canadian tribunal ordered an airline to pay after its website chatbot suggested a passenger could apply for a bereavement fare after travelling

$0.99

per outcome is the price one supplier publishes for its agent; one outcome it charges for is a resolution, defined as "No further help is requested after Fin's last answer"

The suppliers have arrived, and most sell the answer itself. Vendor pages now describe agents for chat, email and voice, for triage, drafting and quality review.

They disagree on who sends, and on what the agent may change. This is the clearest difference in the published material.

- Microsoft's documentation for its case management agent sets two modes side by side. One "sends the email automatically". The other "creates a draft for a customer service representative to review and send". A custom agent can also ask for the case to be closed [1]. In the page's worked example the agent processes a refund once the customer confirms the amount. We found no approval step in it.
- Intercom says of its agent: "It updates accounts, processes payments and refunds, troubleshoots technical issues and much more." Its price is "$0.99 per outcome". One of the outcomes it charges for is a resolution, defined as: "No further help is requested after Fin's last answer" [2]. By that wording, a customer who goes quiet counts.
- One published workflow template ends in a step named "Send Auto-Reply (Gmail)", which "Delivers the response directly to the customer's inbox" [3].
- A second template's default is: "Close tickets when AI classification indicates resolved and add an internal closure note". A third does the opposite: "Post drafts as private/internal notes by default for human review" [3].
- Relevance AI's reply drafter places a cited draft in front of the adviser. "Nothing sends until they approve it" [4].

These are different answers, and a buyer should know which one it is buying. Three things are in play: whether the customer hears from the agent or an adviser, whether a person reads a reply before it goes, and whether the agent can move money, change an account or close the matter.

Few pages say what the customer is told. On the supplier pages we opened, we found no statement that the customer is told the agent is AI. Hand-over is described as a feature: Intercom's agent "hands off to your human team effortlessly, maintaining full customer context" [2]. Voiceflow's documentation is plainer. "A tool that creates a support ticket does not, by itself, connect the customer to a live human conversation" [4].

Hosting separates them less than buyers expect. Rasa's documentation begins "Deploy Rasa Pro on Kubernetes or OpenShift with the Helm chart", on a cluster the customer provides, with "A valid Rasa License" [5]. IBM's page for its agent product says it "Supports multiple clouds (IBM Cloud, AWS and more) or on-premise deployment" [5]. Talkdesk's page speaks of going "from pilot to production on any contact center, third-party system, or on-premises" [5]. Other pages we read do not say. So where it runs is still a question to ask. It is not the question that tells suppliers apart.

A tribunal has held a company to its chatbot's answer. In Moffatt v. Air Canada, decided on 14 February 2024 by British Columbia's Civil Resolution Tribunal, a passenger had asked the chatbot on the airline's website about bereavement fares. The chatbot suggested he could apply for the fare after travelling. The airline's policy did not allow that, and the airline argued it could not be held liable for what its chatbot said. The tribunal called that "a remarkable submission" and wrote: "It should be obvious to Air Canada that it is responsible for all the information on its website. It makes no difference whether the information comes from a static page or a chatbot." It found the airline "did not take reasonable care to ensure its chatbot was accurate" and ordered it to pay $812.02 in damages, interest and fees [6]. The sum is small. The reasoning is the point.

PART 01 · CONTINUED

A support bot has invented a policy. In April 2025 The Register reported that the email support bot of Cursor, a coding tool, told users about a limit on use that the company did not have. Its co-founder wrote in public: “Hey! We have no such policy”, and called the message “an incorrect response from a front-line AI support bot”. He added: “Any AI responses used for email support are now clearly labeled as such” [7]. The label came after the reply had gone out.

A public chatbot has given a wrong answer about the law. In March 2024 The Markup reported that New York City’s chatbot for business owners said: “There are no regulations in New York City that require businesses to accept cash as a form of payment.” The city had passed a law requiring exactly that in 2020 [8].

One supplier has published a narrower practice. Octopus Energy said in July 2026 that in its trial “every AI-written email is clearly labelled and customers can ask to speak to a human advisor at any stage”. Its assistant “never handles vulnerable customers, sensitive cases or complex complaints” [9]. The release does not say whether a person reads each email first.

One company has said what it learned. In May 2025 the chief executive of Klarna, a payments company, told a newswire: “I just think it’s so critical that you are clear to your customer that there will always be a human if you want” [10].

We found no named, dated, independently measured result for a customer service agent, and this paper quotes none as proof.

PART 02

The work nobody has packaged

We listed the jobs a service department does and looked for each among the templates and prebuilt agents we opened. They are thick around the reply and the ticket. We did not find them where the work is assembling a file for a named person.

WHERE TEMPLATES OR PREBUILT AGENTS EXIST	WHERE WE FOUND NONE
A reply written and sent to the customer’s inbox	**The documents of a complaint put in date order for a review or an ombudsman**
A ticket closed on the model’s own classification	**Help articles that disagree with a current policy, or with each other**
A draft reply posted as an internal note	**The cases where the company promised an update by a date**
A drafted reply with the article behind each statement	**What a customer has said about their circumstances, and what the procedure offers**
A refund processed once the customer confirms the amount	**The questions customers ask that no help article answers**

PART 02 · CONTINUED

Two findings matter more than the list. First, in the templates that send, the reply reaches the customer before anyone in the service team has read it [3]. Whether a person reads it is a setting, and whoever chooses it has decided what the company will be held to.

Second, the careful templates show the better pattern. A draft lands in front of the adviser with its sources, and nothing goes until a person sends it [3][4].

The gaps sit where the reading and assembling is slow and the decision already belongs to a named person. That makes them the safer place to begin.

PART 03

Three doors, one gate

An agent is only as useful as what it can reach. In a service department that means three kinds of thing.

Documents.

Help articles, procedures, refund and returns policies, product terms, reply templates. These are indexed once as knowledge sources, with each reader's permissions kept, so an agent answers with the passage it relied on and says when the documents are silent.

Systems of record.

Tickets, cases, orders, bills, contact history. Most are reached through an interface the company already runs. The company gives the platform the description of that interface, chooses which operations become tools, and decides which of them need approval.

Conversations and figures.

Call transcripts, chats, survey comments, waiting times, backlog. These arrive as files, through the same interfaces, or from a reporting database that is read only.

All three pass one gate: every tool call is checked against the company's policy, which can allow it, record it, hold it for a person, or block it.

Where the company's interface checks the caller, the agent acts as the adviser using it and cannot claim another identity.

SIX WAYS INTO A SERVICE DEPARTMENT'S SYSTEMS

	ROUTE	WHAT IT REACHES	WHAT TO KNOW
01	**A file the user uploads**	A customer's letter, a call transcript, a claim form, a survey export	Nothing is connected. The safest start
02	**A knowledge source**	Help articles, procedures, policies, product terms, reply templates	Permissions are kept per reader
03	**An imported description of an interface**	Tickets, cases, orders, bills, contact history	You choose the operations. Changes wait for approval
04	**A reporting database**	Contact reasons, waiting times, backlog, survey results	Read only
05	**A server from the tool catalogue, run in your cluster**	Databases, web search, document conversion	It runs inside; check what each tool does
06	**A channel the customer uses**	The web chat or the voice line of the two answering agents	Published articles only. It says it is AI. An adviser on request

71

public interface descriptions read, from 16 publishers; 68 import

2,488

of the 4,774 operations in the 68 that import are sent as changes

What is published. We read 71 public interface descriptions from 16 publishers and ran each through the platform's importer. 68 import, 1 is refused, and 2 are over the size limit. The 68 carry 4,774 operations, of which 2,488 are sent as changes, and the method alone cannot tell which only read [11].

What that exercise taught us. More than half of the operations are sent as changes. In these files a change is often a message sent, a ticket closed, a call placed or a refund raised.

What is not published. The 16 publishers are help desks, messaging and voice platforms, contact centre platforms, a commerce platform, and incident and chat tools. Several of the largest service suites publish none a file address reaches. For those, the description comes from the company's own system, or the agent reads from the reporting database.

Voice. A voice agent needs speech recognition and speech synthesis. Where the company serves both itself, the caller's audio is processed on its own infrastructure. Where it uses a hosted speech service, the audio goes to that service.

Recordings and card numbers. The company sets what is recorded and for how long. The safe design keeps card numbers out of transcripts wherever it can.

So choose the operations at import, start with the ones that only read, and add nothing that reaches a customer or moves money until the approval step has been seen working.

PART 04

What an agent may do

Two rules run through every agent in this paper. It prepares, and a named person decides and sends. And a customer can always reach a person.

01	**It reads** what the adviser using it could read, and no more.
02	**It drafts** the reply, the note, the response, the report, with each statement tied to its source.
03	**It lists** what is missing, what is out of date, what does not fit, and what falls due.
04	**It stops** before anything that changes a record or reaches a customer about their own case. A tool that changes something waits for approval.
05	**It does one job.** An agent is given the tools for one duty. The agent that checks a refund request holds no tool that pays one.
06	**It scores no adviser and reads no mood.** It gives no score and no rank to a member of staff, and it infers no emotion from a voice or a text, the customer's or the adviser's.
07	**Customers are told.** The two agents that answer a customer say they are AI at the start, and an adviser takes over on request, with the conversation so far.

That gives 35 agents in nine areas: first answers; email, chat and messages; tickets and cases; complaints; refunds, credits and claims; the knowledge base; calls; quality and training; reporting and customer feedback. The full list is a separate document. Four lines are drawn tighter than the rest.

The two that answer.

The Self-Service Answer Agent and the Voice Answer Agent answer a customer directly. The fence is narrow. They answer general questions only, from help articles the company has published and approved. They say they are AI, give the article, and say so when the articles do not cover the question. They touch no account and promise nothing. They hand over to an adviser when the customer asks, and whenever the question is about the customer's own account, order or complaint. The tribunal's reasoning is why: an approved article is an answer the company already stands behind.

Replies and complaints.

An agent may draft a reply from the ticket, the contact history and the help articles. It may check a contact against the company's definition of a complaint, set out what happened, and draft the response. It sends nothing. An adviser edits and sends every reply. A complaints handler decides the outcome and any redress, and no agent closes a complaint.

Refunds, credits and claims.

An agent may check a refund request, a return or a warranty claim against the policy and the order, and say which rule applies and what is missing. It issues no refund, credit or compensation and changes no order. An adviser or a supervisor approves or refuses.

Advisers.

An agent may check a sample of conversations against the quality checklist and list, by topic and by process step, where the procedure was not followed. It gives no score and names no adviser. The quality lead reviews the findings and decides any coaching. Escalation works from rules the company writes, such as a case's age or a repeat contact, not from a reading of anyone's mood.

PART 05

What the rules already ask

This part needs a consumer lawyer's read, and a privacy officer's, before it is relied on. It reports what the texts say; it is not advice. It does not say how any agent would be classified under any law. That is for counsel, system by system. Nothing in it says that a platform makes a company compliant.

Telling the customer

A person is told they are dealing with an AI system. Article 50(1) of the AI Act: "Providers shall ensure that AI systems intended to interact directly with natural persons are designed and developed in such a way that the natural persons concerned are informed that they are interacting with an AI system", unless that is obvious to a reasonably well-informed person in the circumstances [12]. Article 50(5) adds that the information is given "in a clear and distinguishable manner at the latest at the time of the first interaction or exposure" [12].

That duty already applies. The Act's general date of application is 2 August 2026 [12]. That is not the date for high-risk systems. Regulation (EU) 2026/1744 moved only those: the high-risk duties "shall apply from: (i) 2 December 2027 as regards AI systems classified as high-risk pursuant to Article 6(2) and Annex III; and (ii) 2 August 2028 as regards AI systems classified as high-risk pursuant to Article 6(1) and Annex I" [13].

Reaching a person

In Spain, a machine cannot be the only channel. Law 10/2025 on customer service bars the use of answering machines or similar means as the only way to serve customers. Where a company uses them, or conversational bots, the customer must be able to ask for a human operator at any point from the start, and the operator identifies themselves when the conversation begins. The law covers companies providing listed basic services, among them water, gas and electricity, passenger transport, post, telecoms and financial services, and large companies. It entered into force on 28 December 2025 and gives companies twelve months to adapt [14]. We read the Spanish text; the English here is ours.

In financial services sold at a distance, a person on request. Directive (EU) 2023/2673 requires that, where the trader uses online tools, "the consumer shall have a right to request and to obtain human intervention at the pre-contractual stage, and in justified cases after the distance contract has been concluded". Its recitals name "chatbots, roboadvice, interactive tools or similar means". Member States apply the measures from 19 June 2026 [15].

Decisions about a customer

A person decides. Under the GDPR, a person "shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her", with stated exceptions. Where an exception applies, the person keeps "at least the right to obtain human intervention on the part of the controller" [16]. Whether a refused refund is such a decision is for counsel. No agent in this paper refuses or grants one.

PART 05 · CONTINUED

Advisers

Evaluating workers is on the high-risk list. Annex III of the AI Act lists, under employment, "AI systems intended to be used to make decisions affecting terms of work-related relationships, the promotion or termination of work-related contractual relationships, to allocate tasks based on individual behaviour or personal traits or characteristics or to monitor and evaluate the performance and behaviour of persons in such relationships" [12]. Its date is 2 December 2027 [13].

Reading emotion at work is prohibited. Article 5(1)(f) bans "the use of AI systems to infer emotions of a natural person in the areas of workplace and education institutions", except for medical or safety reasons [12]. By its words that article is about the workplace and education. This paper does not say it covers a customer's mood. The agents here read no mood on either side of the conversation, as a design choice.

The United Kingdom and the United States, in brief

The United Kingdom's data protection law, as amended in 2025, says "a decision is based solely on automated processing if there is no meaningful human involvement in the taking of the decision" [17]. The Financial Conduct Authority's guidance on its Consumer Duty asks firms to "ensure that customers do not face unreasonable barriers" [17].

In the United States we opened one text. The Federal Communications Commission ruled in February 2024 that the restrictions on calls using an "artificial or prerecorded voice" "encompass current AI technologies that generate human voices", so such calls need the called party's prior consent [18]. No agent in this paper places a call to a customer. We opened no state law, and say nothing about them here.

MARKET	WHAT WE READ	WHAT IT ASKS
European Union	AI Act, Article 50(1) and (5)	People are told they are dealing with an AI system, at the first interaction [12]
European Union	GDPR, Article 22	No decision "based solely on automated processing" with a significant effect [16]
European Union	Directive 2023/2673	Human intervention on request in distance financial services, from 19 June 2026 [15]
Spain	Law 10/2025	A machine is never the only channel; a human operator on request [14]
United Kingdom	UK GDPR, Article 22A	"meaningful human involvement in the taking of the decision" [17]

PART 05 · CONTINUED

What the rules have in common

01	**The company answers for what its agent says,** as it does for its website.
02	**The customer is told** when the other side is an AI system.
03	**A person can be reached** on request.
04	**Advisers are protected too.** Evaluating workers is a listed use, and reading their emotion is barred.

PART 06

A worked example: one home appliance retailer's service team

Take a made-up retailer: home appliances sold in two countries and a service team of 60. Nothing here is a customer's story, and every figure is illustrative; it is section 4 put to work.

01

A customer asks in the web chat how long the warranty on a washing machine lasts. The agent says it is AI, answers from the published warranty article and gives the link. She then asks about her own repair. The agent hands the chat to an adviser, with the conversation so far.

02

An email arrives about a late delivery. The agent drafts the reply from the ticket, the order record and the delivery article, with sources. The adviser changes one sentence and sends it.

03

A customer asks for a refund on an oven bought 40 days ago. The agent checks it against the refund policy and the order, says which rule applies, and lists what is missing. A supervisor decides.

04

A complaint comes in about three missed repair visits. The agent sets out what happened in date order and drafts the response against the complaints procedure. The complaints handler decides the outcome and the redress, and sends it.

05

After a call, the agent drafts the note from the transcript: the reason for the call, what was agreed, what is open. The adviser corrects the callback date and saves it.

Every one of these reads articles or records the company already has, changes no record by itself, sends nothing about a customer's own case without a person, and leaves a record of who asked and what was shown.

PART 07

Seven questions before an agent talks to a customer

01 **Who answers the customer?**
Which questions does the agent answer directly, and from which approved content?

02 **Who sends?**
Which replies or complaint responses go out without a person reading them first? Is that a setting, and who holds it?

03 **What may it change?**
Can it issue a refund or a credit, change an account or an order, or close a ticket or a complaint? Does each of those wait for a named person?

04 **Is the customer told, and how do they reach a person?**
Does it say it is AI at the start? Is an adviser one request away, with the conversation passed on?

05 **Where does it run?**
On our infrastructure, in our private cloud, or in the supplier's service? Where is the model? For voice, where do speech recognition and speech synthesis run?

06 **What is counted, scored and recorded?**
What counts as a resolved contact, and does a customer's silence count? Does anything score an adviser or read mood? What does the record hold, and for how long?

07 **Can we stop it?**
One agent, one tool, or all of them, at once, by us.

PART 08

Where to start

Pick work where reading is slow, the decision stays with a person, and nothing reaches a customer on the agent's word.

Reply drafting
The agent drafts from the ticket and the articles, with sources. An adviser sends every reply.

Call summaries
The agent drafts the after-call note. The adviser checks and saves it.

Knowledge gaps
The agent lists the questions no article answers. The knowledge owner decides what to write.

One team, one channel, invented data first. Then beside the manual process, measured against it.

About Katonic AI

Katonic AI builds the platform that organisations use to build, run and govern AI agents on their own infrastructure: in their own data centre, in a private or sovereign cloud, or with no connection to the internet at all. The platform is in production with governments, banks, telecom operators and utilities. Katonic is certified to ISO 27001.

KATONIC.AI/SOLUTIONS/CUSTOMER-SERVICE

SOURCES

Read on 22 November 2026. Figures are the publishers' own and were not audited by us.

01 Microsoft Learn, "Integrate Case Management Agent with custom Copilot Studio agents" (Dynamics 365 Customer Service developer documentation), "Last updated on 2026-11-06" as shown on the page. https://learn.microsoft.com/en-us/dynamics365/customer-service/develop/case-management-agent-integration

02 Intercom, Fin home page and pricing page, undated, read 22 November 2026. https://fin.ai/ ; https://fin.ai/pricing

03 n8n workflow template 8744, "Automate Zendesk Support replies with GPT-4o-mini, Google Sheets & Gmail" ("Last update 9 months ago" on the page; no publication date). https://n8n.io/workflows/8744-automate-zendesk-support-replies-with-gpt-4o-mini-google-sheets-and-gmail/ ; Zapier automation templates "Update and close support ticket with AI reply" and "Create AI-drafted private ticket reply for support agents", undated, read 22 November 2026. https://zapier.com/automations/customer-service-success/conversational-support/ai-customer-support/update-and-close-support-ticket-with-ai-reply ; https://zapier.com/automations/customer-service-success/conversational-support/ai-customer-support/create-ai-drafted-private-ticket-reply-for-support-agents

04 Relevance AI, "Reply Drafter" use case page, undated, read 22 November 2026. https://relevanceai.com/use-cases/reply-drafter ; and Voiceflow documentation, "Human agent hand-off" overview, "Last updated Nov 12, 2026" on the page. https://www.voiceflow.com/docs/documentation/handoff/overview

05 Rasa documentation, "Deploying to Kubernetes", undated, read 22 November 2026. https://rasa.com/docs/pro/deploy/deploy-kubernetes/ ; IBM, watsonx Orchestrate product page, undated, read 22 November 2026. https://www.ibm.com/products/watsonx-orchestrate ; Talkdesk, home page, undated, read 22 November 2026. https://www.talkdesk.com/

06 Moffatt v. Air Canada, 2024 BCCRT 149, Civil Resolution Tribunal of British Columbia, file SC-2023-005609, date issued 14 February 2024, paragraphs 2, 27, 28 and 44. The tribunal's own PDF: https://decisions.civilresolutionbc.ca/crt/crtd/en/525448/1/document.do

07 The Register, Thomas Claburn, report on Cursor's AI support bot inventing a usage policy, 18 April 2025. https://www.theregister.com/2025/04/18/cursor_ai_support_bot_lies/

08 The Markup, Colin Lecher, "NYC's AI Chatbot Tells Businesses to Break the Law", 29 March 2024. https://themarkup.org/news/2024/03/29/nycs-ai-chatbot-tells-businesses-to-break-the-law

09 Octopus Energy, "Octopus Energy's AI trial wins customer approval", press release, 7 July 2026. https://octopus.energy/press/more-news-press-releases/octopus-energy-s-ai-trial-wins-customer-approval/

10 Entrepreneur, Sherin Shibu, report on Klarna's chief executive reversing course on AI in customer service, 9 May 2025, quoting his remarks to Bloomberg. https://www.entrepreneur.com/business-news/klarna-ceo-reverses-course-by-hiring-more-humans-not-ai/491396

SOURCES continued

11 Templates and prebuilt agents of agent and workflow platforms, as opened for sources 1 to 5; and 71 interface descriptions from 16 publishers fetched 22 November 2026 and run through the platform's importer (Katonic research).

12 Regulation (EU) 2024/1689 (AI Act) of 13 June 2024, Article 5(1)(f), Article 50(1) and (5), Article 113, Annex III point 4(b). http://publications.europa.eu/resource/celex/32024R1689

13 Regulation (EU) 2026/1744 of 8 July 2026 (Official Journal 24 July 2026), the replaced Article 113 point (c). http://publications.europa.eu/resource/celex/32026R1744

14 Spain, Ley 10/2025, de 26 de diciembre, por la que se regulan los servicios de atención a la clientela (BOE-A-2025-26698; BOE no. 312 of 27 December 2025), Articles 1, 2 and 8, the transitional provision and final provision 8; consolidated text. https://www.boe.es/eli/es/l/2025/12/26/10/con

15 Directive (EU) 2023/2673 of 22 November 2023, inserting Article 16d into Directive 2011/83/EU (paragraph 3), recitals 15 and 40, and Article 2 (application from 19 June 2026). http://publications.europa.eu/resource/celex/32023L2673

16 Regulation (EU) 2016/679 (GDPR) of 27 April 2016, Article 22(1) and (3). http://publications.europa.eu/resource/celex/32016R0679

17 United Kingdom, Data (Use and Access) Act 2025, section 80, inserting Article 22A into the UK GDPR; commencement note "S. 80 in force at 5.2.2026". https://www.legislation.gov.uk/ukpga/2025/18/section/80 ; and Financial Conduct Authority, FG22/5, "Final non-Handbook Guidance for firms on the Consumer Duty", July 2022, chapter 9. https://www.fca.org.uk/publication/finalised-guidance/fg22-5.pdf

18 United States, Federal Communications Commission, Declaratory Ruling FCC 24-17, CG Docket No. 23-362, adopted 2 February 2024, released 8 February 2024, paragraph 2. https://docs.fcc.gov/public/attachments/FCC-24-17A1.pdf