Relativity aiR for Data Breach Response Pro Study Guide

RelativityOne 2026

Last modified: July 2, 2026

Contents

Exam Overview

This document is designed to be a study guide for the Relativity aiR for Data Breach Response Pro exam. Using this study guide, you can take a deeper dive into the key points of aiR for Data Breach Response features and workflows, and prepare with exam questions and additional resources.

Exam description

The Relativity aiR for Data Breach Response Pro exam certifies knowledge of key features, functionalities, and workflows of the aiR for Data Breach Response product. This exam is designed for eDiscovery counsel, lit support, cyber response, cybersecurity, and sales professionals who want to validate their knowledge.

Study guide format

This study guide is divided into three parts:

- **Exam Details**: This section contains information about the exam structure and format and instructions on how to prepare for the exam.
- **Study Resources:** This section is divided based on exam topics, with one section for each exam topic. Each section contains key points and references to help further your knowledge of aiR for Data Breach Response. Understanding the concepts and references mentioned in the Study Resources is crucial to both your work in aiR for Data Breach Response and successfully passing the exam.
- **Exam Items:** This section contains all exam items. Use the study resources and Study Links table to identify correct answers.

Study guide and exam updates

The Relativity aiR for Data Breach Response Pro exam currently tests on RelativityOne. Review the Exam Update Schedule for upcoming exam changes. Plan accordingly when studying for the exam.

To view changes made to the exam from one version to the next, review the Change Log.

Exam Details

Exam Structure

The following table outlines the various topics tested on the exam.

About aiR for Data Breach Response
Understand value prop and benefits
Prepare for a project
Features and Functionality
Use the Viewer
Use PI and Entity Search
Set up and select detectors
Run Data Analysis
Review reports
Workflows and Use Cases
Identify review populations
Configure tags
Use detectors
Identify and understand use cases

Exam Format

The exam contains a 50-question quiz covering the topics listed above. All questions are multiple choice. You will have 35 minutes to complete the quiz. You must achieve a score of 80% or higher to pass.

Exam Preparation

There are no prerequisites for this certification; however, in order to prepare, we recommend following:

- ***Read and understand this Study Guide.*** This guide contains the necessary materials to prepare for the exam, including topics, key points, resources, and exam items.
- ***Review the*** RelativityOne documentation***.*** This is a free and useful resource for the exam.
- ***Familiarize yourself with the on-demand and live training content.*** We offer a variety of training resources to enhance your knowledge, which can be found on Relativity Learning.

Study Resources

About aiR for Data Breach Response

Key points

Overview

- aiR for Data Breach Response is an AI-powered solution that helps teams identify personal information (PI) and protected health information (PHI), link that data to impacted individuals, assess breach impact, and build notification lists when responding to a data breach.
- aiR for Data Breach Response is purpose-built for breach response workflows, helping teams move faster, reduce manual review, and lower costs while maintaining transparency, human oversight, and defensibility in RelativityOne.
- aiR for Data Breach Response uses a combination of generative AI and pre-built AI and machine learning (ML) models to identify, extract, and link predetermined PI/PHI types.
- Typical aiR for Data Breach Response customers handle a significant volume of data and handle multiple data breaches per quarter.
- Both structured data and unstructured data can be reviewed by aiR for Data Breach Response.
  - In **structured data** such as databases and spreadsheets, aiR for Data Breach Response uses AI to analyze table boundaries and headers, identifies PI within tables, and links names and PI that appear in the same table row.
  - In **unstructured data** such as emails, aiR for Data Breach Response uses generative AI and document context to identify PI types, link PI to names and entities, and consolidate names and entities across documents.

Value proposition

- **Deliver Trusted, Accurate Outcomes:** Minimize missed or duplicate data through intelligent detection, linking, and normalization, all within a secure cloud environment.
- **Ensure Defensibility:** Support transparent, explainable workflows, human review, and audit-ready outputs that help teams make clear, defensible decisions.
- **Increase Efficiency:** Accelerate breach response by automatically identifying and linking PI/PHI, generating accurate notification lists, and reducing manual work.
  - A Relativity customer using aiR for Data Breach Response found reviewers were 78% faster when analyzing relevant documents.
- **Gain Actionable Insights:** Use dashboards, search, and reports to assess breach impact, monitor progress, and guide next steps throughout the workflow.
- **Keep Data Secure in RelativityOne:** Data never leaves RelativityOne's security boundary while teams identify sensitive data and develop notification-ready outputs.

Features and Functionality

Key points

Reviewing documents in the Viewer

- Both unstructured and structured documents can be reviewed in RelativityOne's viewer. For structured documents, you can review how AI identified table boundaries, column headers, the PI that sits within those tables, and the links created. For unstructured data, you can validate the predictions highlighted by aiR. PI detections or records identified by aiR are indicated by a sparkle icon in the **PI Detection panel.**
- In the native viewer, you can:
  - Create and edit spreadsheet tables
  - Capture, edit, or remove PI
  - Add annotations

Setting up and selecting detectors

- Detectors are used to identify PI. From the Settings tab, you can view all detectors and tags, add new detectors and tags, and turn detector tags on and off.
  - A **PI detector** uses a combination of AI models, regular expressions (RegExes) and keywords that detect a string of text and classify it as a form of PI.
  - An **Out of the Box (OOTB) detector** is a pre-built detector created for aiR for Data Breach Response.
  - A **Custom detector** is created by a user to solve the needs of a given project outside of the OOTB offerings.

Data Analysis

- **Data Analysis** is the combination of machine stages that identify PI, link PI to identified entities, normalize data, and prepare data for review and reporting. The stages are:
  - **Run Blocklisting-** The **Blocklist** is a list of excluded PI values that were not relevant for review or were falsely tagged. Prior detections matching blocklisted terms will be marked as Blocklisted and have their links broken.
  - **Run Unstructured Detectors-** automatically identifies and links PI across unstructured documents, normalizes entities across documents, and generates a consolidated notification list,
  - **Run Structured Detectors-** identifies PI by running all enabled PI detectors on unlocked structured documents and automatically links all names and PI.
  - **Run Normalization-** standardizes names and PI into consolidated entities and generates an updated entity report.
  - **Compile Insights-** calculates and consolidates PI and entity statistics for reporting.

Project Dashboard

- The **Project Dashboard** is available to project leads and provides a centralized view of project status, key metrics, and overall progress in the data breach response workflow. The Project Dashboard populates once PI Detection begins during Data Analysis

PI and Entity Search

- With **PI and Entity Search** you can search documents with PI and Data Breach Entities that have been identified by aiR for Data Breach Response. These searches can be saved and used in other Relativity features such as **Review Center**.

- You can search for:
  - Documents that have any PI identified.
  - Documents that have Data Breach Entities identified.
  - Documents containing specific PI types.

Quality control

- aiR for Data Breach Response has built-in QC tools to ensure that there is a minimization of documents with over and under predictions.
  - The **Spreadsheet QC tool** allows a Project Manager to view and edit PI assignments across all spreadsheets.
  - The **Blocklisting tab** shows all terms that have been identified as PI. From here, you can add these terms to the **Blocklist**. You can also view terms that have been added to the Blocklist and remove them if necessary.

Reporting

- You can generate the following reports throughout an aiR for Data Breach Response project lifecycle.
  - **Document Repor**t: provides a breakdown of each document within the data set and the PI found within that document.
  - **Entity Report**: provides a breakdown of each entity found within the data set. An **Entity** is a unique individual identified by their name and their PI, such as SSN, date of birth, and full address. Entities are composed of one or more records. A **Record** is a pairing of a raw name and PI data from the dataset.

Workflows and Use Cases

Key points

Workflow

The steps below illustrate the high level workflow steps of aiR for Data Breach Response.

1. Upload and process documents into your **RelativityOne environment**.
2. Configure the **aiR for Data Breach Response application** into your environment.
3. Start a new **Data Breach ingestion job** and select a saved search of documents for review.
4. In **Settings**, select the personal information detectors to use and add any custom detectors.
5. Run **Data Analysis** to immediately start getting insights on identified, extracted, and linked PI and review documents as PI detection completes.
6. Access the **Project Dashboard** to view PI and PHI identified across the document set.
7. Review the **Entity Analysis tab** to resolve any conflicts and finalize the **notification report** containing impacted individuals and their associated PI.

Use Case

- aiR for Data Breach Response is purpose-built to help teams identify and link impacted PI/PHI, understand which individuals are affected, and generate accurate notification lists so they can respond faster with defensible outcomes after a data breach.

Reference materials

Training Course

- Relativity aiR for Data Breach Response Foundations: Set Up to Entity Reporting

Documentation

- aiR for Data Breach Response
- aiR for Data Breach Response Documentation

Reference Materials

Use the following materials to help you prepare for the exam. We recommend using the Status column to mark activities as Complete as you study.

Resource	Status
Relativity Documentation	
Data Breach Response	
• Installation and configuration	
• Document ingestion	
• Personal Information detectors	
• Data Analysis	
• Project Dashboard	
• PI and Entity Search	
• Quality control	
– Detector QC	
– Spreadsheet QC tool	
– Blocklisting	
• Review documents	
– Unstructured document review	
– Structured document review	
• Entity Analysis	
– Entities	
– Conflicts	
• Review Center: Creating a Review Center queue	
Relativity Learning Courses	
• Relativity Data Breach Response Introduction	
• Relativity Workflow Essentials: Data Breach	
• Relativity aiR for Data Breach Response Foundations: Set Up to Entity Reporting	

Exam Items

Exam Items

This section includes all exam items on the Relativity aiR for Data Breach Response Pro exam. You will receive the 50 questions below on exam day. Use the Study Links table to locate and identify the correct answer to each question. In order to maintain the integrity of the exam, Relativity will not share or confirm answers to any of the questions.

Note: Exam questions are subject to change without notice and in Relativity's sole discretion. All question changes will be reflected in the items below, as well as in the Change Log.

As an additional study tool, the aiR for Data Breach Response Pro Practice Quiz is available in Relativity Learning. You can use it to review the items below.

About aiR for Data Breach Response

1. What is aiR for Data Breach Response?
   a. An AI-powered solution used to reduce the time, cost, and risk to produce an entity notification list.
   b. An AI solution used to reduce the time, cost, and risk to produce a set of redacted documents.
   c. An AI solution used to remove the need for human review in order to generate an entity notification list.

2. After installing aiR for Data Breach Response, how can you provide RelativityOne users access to the application?
   a. Add them to a group that has the correct permissions set.
   b. It is automatically available for all users after installing the application.
   c. Contact Relativity to have their accounts created.

3. What types of technologies does aiR for Data Breach Response use to identify, extract, and link PI and PHI?
   a. User-created AI and machine learning (ML) models to find pre-determined PI/PHI types.
   b. User-created AI and machine learning (ML) models to find unknown PI/PHI types.
   c. A combination of generative AI and pre-built AI and machine learning (ML) models to find pre-determined PI/PHI types.

4. What is a benefit of using aiR for Data Breach Response?
   a. It provides endpoint threat detection capabilities to prevent malicious activity and threat actors.
   b. Data never leaves the RelativityOne security boundary.
   c. You can determine the relationship between impacted individuals and fraudulent activity.

5. What is required to ingest documents into aiR for Data Breach Response?
   a. A list of terms and PI types
   b. A saved search
   c. An entity notification list

6. What regular expression syntax does aiR for Data Breach Response accept when creating a custom detector?
   a. Python
   b. TR1
   c. Java

7. What is the difference between how aiR for Data Breach Response identifies PI in structured data versus unstructured data?
   a. Regular expressions and contextual clues are used in structured data while table boundaries and headers are analyzed in unstructured data.
   b. Table boundaries and headers are analyzed to extract and link in structured data while aiR powers the extraction and linking in unstructured data.
   c. The two data types are treated the same by aiR for Data Breach Response during analysis.

8. How does aiR for Data Breach Response expedite manual entity normalization and deduplication efforts while still providing transparency into the process, leading to faster and cleaner generation of entity reports?
   a. Through automatic normalization and deduplication of personal information.
   b. Through automatic linking of personal information across all document formats.
   c. Through automatic normalization and deduplication of entities and entity review workflows.

9. When does a document in aiR for Data Breach Response begin to be counted toward billing for the application?
    a. Once the Compile Insights stage of Data Analysis is complete.
    b. Once the document has been ingested.
    c. Once the document has been run through the Run Unstructured Detectors stage of Data Analysis.

10. Do you need to prompt aiR for Data Breach Response?
    a. Yes, you must manually craft and refine prompts for each PI type to ensure proper aiR identification.
    b. No, aiR for Data Breach Response is pre-prompted and tested across synthetic data to ensure the highest accuracy.
    c. aiR for Data Breach Response is pre-prompted, but you can choose to use custom prompts instead for various PI types.

Features and Functionality

11. You've added a new regular expression to your custom Employee ID detector. What must you do before rerunning PI Detection to ensure that your new regular expression will be included?
    a. Save and build the Employee ID detector.
    b. Validate the regular expression.
    c. Add it to the existing list of out of the box detectors.

12. How can you identify documents that have failed ingestion into aiR for Data Breach Response?
    a. By reviewing the Failed Ingestion Report.
    b. By filtering for failures in the Document table from the Project Dashboard.
    c. By filtering for failures in the Results column from the Job Details page.

13. What should you do after resolving conflicts generated by the automatic entity normalization process?
    a. Run an Entity Report job from the Reports tab and export the report as an excel document; you do **not** need to rerun Data Analysis.
    b. Rerun the normalization and compile insights stages of Data Analysis to record your changes and generate an updated version of the Entity Report.
    c. Rerun the compile insights stage of Data Analysis to record your changes and generate an updated version of the Entity Report.

14. What is the purpose of configuring your entity normalization and deduplication preferences?
    a. To ensure your coding layout is configured with the fields you want your review team to utilize.
    b. To ensure every entity name requiring redaction is unique.
    c. To tell the algorithm how to normalize and deduplicate the entities found within the data set.

15. What information about entities is searchable with PI and Entity Search?
    a. Jurisdiction the entity resides in.
    b. Entity count on the document.
    c. Other names an entity is known by in the dataset.

16. What is the purpose of resolving entity conflicts?
    a. To make final merge decisions where entity normalization has identified possible similarities but **cannot** make a final decision.
    b. To manually normalize and deduplicate all entities identified by a human review team earlier in the project.
    c. To manually resolve all tertiary conflicts while relying on automatic normalization processes to resolve primary and secondary conflicts.

17. What syntax can you use for a PI and Entities Search to find documents with a breached entity?
    a. CONTAINS DATA BREACH ENTITY
    b. CONTAINS BREACH
    c. CONTAINS PI

18. What type of Review Center queue should be used for aiR for Data Breach Response?
    a. Data breach review queue
    b. Coverage review queue
    c. Saved search review queue

19. Where can you review and potentially merge two entities that have the same name but differing social security numbers once entity normalization completes?
    a. The Clusters subtab within the Entity Analysis tab
    b. The Entities subtab within the Entity Analysis tab
    c. The Conflicts subtab within the Entity Analysis tab

20. Which chart on the Project Dashboard shows the total identified PI across the document set?
    a. Document Types chart
    b. Entity Conflicts to Review chart
    c. PI Types chart

21. You have an existing workspace that has multiple completed aiR for Data Breach Response jobs. How can GenAI Mode be enabled for this workspace?
    a. Navigate to Workspace Settings and toggle on GenAI Mode in your workspace.
    b. The application will need to be redownloaded, and GenAI Mode will automatically be activated for every future job.
    c. This is **not** possible. A new workspace will need to be created that has never run aiR for Data Breach Response before.

22. What is the purpose of the final Entity Report in a data breach?
    a. It demonstrates how many documents in the dataset contain personal information and which documents are likely to contain entities who need notified.
    b. It gives users a list of entities found in the dataset, which can be used as search terms to build a review queue.
    c. It provides a list of affected individuals found in the dataset and which of their personal information was breached.

23. What does the Project Dashboard tell you about a data breach?
    a. It provides a holistic view of PI and entity data across the document set as it updates throughout the lifecycle of a project.
    b. It displays which documents in the dataset may be privileged and should be reviewed outside of aiR for Data Breach Response.
    c. It provides an understanding of which reviewers' work will need a second review to ensure the final Entity Report's accuracy.

24. Where will links between names and personal information appear in the review interface?
    a. Records panel
    b. PI Detection panel
    c. Coding layout

25. Two records were merged due to an identical email address and date of birth. What would the Merge Reason field in the Entity Report show as for this merged record?
    a. Primary PI Match
    b. Secondary PI Match
    c. Tertiary PI Match

26. Where can you review all of the records associated with an entity?
    a. The Entity Analysis tab
    b. The Conflicts subtab
    c. The Records panel in the Document Viewer

27. A reviewer has verified table boundaries and column level PI detections in a spreadsheet. How are names and PI within each table linked?
    a. Names and PI appearing in the same table row will automatically be linked.
    b. The reviewer will link name and PI table columns using the Linked Individuals tab in the coding layout.
    c. The reviewer will link name and PI table columns by selecting each column and clicking the links icon that appears in the document viewer.

28. For spreadsheets, how do you annotate personal information located in a cell outside of a table?
    a. Use the layout to record the personal information in a text box.
    b. Create a new table to capture the single cell.
    c. Right click on the cell and use the Personal Information > Add feature in the context menu to record the personal information.

29. How can a user link personal information PI/PHI that was **not** automatically linked to an individual or organization?
    a. Manually capture it in the layout using a long text field.
    b. Click Add Link from the records panel or multi-select the proper name/organization and PI/PHI and click Link.
    c. Click Add Link from the records panel and manually type in the PI/PHI.

30. How can you tell if an entity in the Records panel was generated automatically using aiR?
    a. It will be listed under the Identified by aiR subheading.
    b. It will be highlighted in red before being accepted by a reviewer.
    c. It will have a sparkle icon next to it.

31. What must all unstructured documents contain to ensure accurate processing by aiR?
    a. Only the native file.
    b. Properly OCR-processed text.
    c. An imaged version of the document.

32. Which feature enables configuring which fields appear in the aiR for Data Breach Response dashboard list view and entity report tabs as well as filtering and sorting by those fields?
    a. Custom layouts
    b. Custom dashboard widgets
    c. Custom column configuration

Workflows and Use Cases

33. When reviewing a document, what option needs to be selected to prevent changes to PI detections when Data Analysis is run?
    a. Preserve Detections
    b. Lock Annotations
    c. Confirm Codings

34. During which Data Analysis stage does Generative AI run?
    a. Unstructured Detection
    b. Structured Detection
    c. Normalization

35. Why should you use the blocklisting tool?
    a. When some instances of a string of text are personal information and sometimes they are not.
    b. To find missed personal information.
    c. To remove false positives from your detector results.

36. What does the blocklisting tab show?
    a. All terms that have been identified as PI in unstructured documents.
    b. All terms that have been identified as PI in structured documents.
    c. All terms that have been identified as PI in both unstructured and structured documents.

37. What method is most effective to verify the accuracy of personal information (PI) predictions in spreadsheets?
    a. Navigate to the Spreadsheet QC tool and confirm the header values in the dataset are accurately classified as a PI type or no PI.
    b. Navigate to the blocklist tool and locate any header values that should be considered no PI in the dataset. Designate any remaining header values as containing PI within the blocklist tool.
    c. Review a sample set of 50 spreadsheets and keep a record of any header values that you will instruct reviewers to update during their review.

38. What is the best way to test the accuracy of a custom regular expression?
    a. Run Data Analysis on the entire dataset to see if the expression works as intended.
    b. Use the built-in testing feature to see if a string of text is captured by a regular expression.
    c. Use a third-party regular expression validator.

39. What is the difference between a global and local keyword when building detectors?
    a. Global keywords look for a phrase and translate it across Relativity's database of 100+ non-English languages. Local keywords only look for a phrase in the English language.
    b. Global keywords look for a regular expression on a document. Local keywords look for a phrase located next to a regular expression.
    c. Global keywords look for a phrase in the entire document, regardless of its proximity to a regular expression. Local keywords look for a phrase within a designated character distance of a regular expression.

40. What is an intended use case for aiR for Data Breach Response?
    a. A cyber incident has occurred and impacted individuals need to be notified.
    b. A set of documents must have all PI redacted for upcoming litigation.
    c. A set of documents are deemed to contain privileged information and need to be logged for litigation purposes.

41. How can you identify documents to QC in aiR for Data Breach Response?
    a. Re-run Data Analysis for every 100 documents reviewed and QC the results from the dashboard to ensure consistency.
    b. Leverage the dashboard to identify the document groups to QC or leverage the custom search provider 'PI and entity search' in the document list.
    c. QC documents while the Run Structured Detectors and Run Unstructured Detectors stages of Data Analysis are running to ensure PI detection consistency.

42. You have a new project where the client wants to quickly know the percentage of documents that contain social security numbers (SSN) or dates of birth (DOB). The data is already ingested into aiR for Data Breach Response, but Data Analysis has **not** yet run. The client has informed you that they are **not** interested in metrics around any other types of PI.

    What are the next steps after documents are ingested into aiR for Data Breach Response to accomplish this request?

    a. Create a custom detector called "SSN or DOB", run Data Analysis, download the Document Report, and filter down to DocIDs containing "SSN or DOB" to calculate percentage.

    b. Toggle off all out of the box detectors except SSN and DOB, run Data Analysis, go to the Project Dashboard, and filter down to DocIDs containing personal information to calculate percentage.

    c. Build regular expressions for SSN and DOB, run Data Analysis, download the Entity Report, and filter down to entities containing SSN or DOB to calculate percentage.

43. Why is it important to manually annotate unstructured documents where text **cannot** be highlighted in the native viewer?

    a. To capture any missing personal information either using extracted text or manual entry and link it to the person it belongs to.

    b. To remove these documents from the project.

    c. To flag the document as needing further review and add any entities from these documents manually in the spreadsheet.

44. Using PI and Entity search, what search criteria would only yield all documents containing both Full Name and Phone Number PI types?

    a. CONTAINS PI TYPE 'Full Name' AND CONTAINS PI TYPE 'Phone Number'

    b. 'PI Count' >=2 AND PI TYPE = 'Full Name', 'Phone Number'

    c. CONTAINS PI TYPE 'Full Name' OR CONTAINS PI TYPE 'Phone Number'

45. What are the two user roles recommended to each have their own permission groups for aiR for Data Breach Response?

    a. Lead and annotator

    b. First-level reviewer and quality-control reviewer

    c. Reviewer and detector

46. Which of the following would you expect to find on the Excluded Cell Values tab?

    a. An SSN in a structured document that consists of only 8 digits instead of 9.

    b. One specific SSN that you created as a blocklisting exception for that PI Type.

    c. SSNs found in unstructured documents that need human review to determine if they match the PI Type.

47. After filtering for all unlocked documents on the Project Dashboard, how can you create a Review Center Queue containing these documents?

    a. Navigate to the document list and create a saved search using the Locked Annotations field.

    b. Create a list using the Save as List mass action and then create a saved search from the list.

    c. You **cannot** query unlocked documents to use for a saved search.

48. Which of the following is true regarding records and entities in aiR for Data Breach Response?

    a. If the same person is found in five documents, five entities and records are created for that person, which will be merged into one entity after normalization.

    b. If the same person is found in five documents, five entities are created for that person, which will be merged into one record after normalization.

    c. If the same person is found in five documents, five records are created for that person, which will be merged into one entity after normalization.

49. If a document contains poor OCR or handwriting, how would you capture personal information in the Viewer?

    a. Click Draw Annotation in the PI Detection panel, draw a box over the personal information, manually enter the value, and then select the PI type.

    b. Highlight the extracted text to capture it as displayed, manually enter the value, and then select the PI type.

    c. Flag the document, add the personal information in comments, rerun OCR, and re-ingest the document.

50. What is the best way to ensure a company-specific employee ID is identified as PI throughout your dataset?

    a. Click Add Detector, fill out the name and category fields, and add regexes and keywords to create a Custom Detector.

    b. Run a PI and Entity Search of the custom regex string that would identify the ID and then use the Identify as PI mass action.

    c. Clone an out of the box (OOTB) detector, update the deduplication identifier to Custom, and add global and local keywords to create a new detector.

Content Distribution Disclaimers

Relativity is aware that a number of apps are available to the public which purport to be study guides for various Relativity exams. These apps are published by third parties that are not affiliated with Relativity, and Relativity does not monitor their content and cannot ensure its accuracy. As such, you should only use Relativity-approved study materials. Your use of any other third party study resources is at your own risk. For approved Relativity study resources please visit the Relativity Community or contact certification@relativity.com if you have any questions.

Relativity may, in its sole discretion, provide certain materials, including but not limited to workbooks, presentations, outlines, webinars, practice examinations, on demand videos and tutorials, and study plans (collectively, the "Materials"), as part of its training, certification, and implementation services. The Materials are unique to each training session and may only be used by individuals who have attended the training session in which they were distributed ("Attendees"). Attendees may request access to Materials that were not distributed during a training session, and Relativity may grant such access in Relativity's sole discretion. Attendees shall not redistribute any Materials to any third party or to any person who has not attended the training session in which the Materials were distributed without Relativity's prior written authorization.

Change Log

1. **March 11, 2026** - Product and exam updated to aiR for Data Breach Response. Content on pages 7-10 updates to include product updates on aiR-powered PI detection in unstructured documents and blocklisting. Questions updated (3, 5, 7, 24, 26, 29, 41) and new questions added (10, 21, 30-32, 34, 36, 45-50). Exam now contains 50 multiple choice questions.
2. **July 2, 2026** - Minor update to Quality Control section on page 9 to reflect product updates.