One Model Security & One AI Transcription

CHAPTER 1

Intro, Topics Covered, & Learning Outcomes

Hey. My name is Hayley, and I'm on the One AI team here at One Model. Before diving into the fundamentals of machine learning and building models in One AI, I want to make sure you are familiar with how security works. We will examine how One Model security interacts and functions within the One AI machine learning model tool and the storyboards where model insights are shared. This knowledge will help you make informed decisions when creating roles and permissioning One AI application access roles.

This module is specifically designed for site administrators and those managing security within One Model. A basic understanding of application and data access roles is recommended before moving forward.
Model creators may also benefit as they could encounter issues related to incorrect permissions.

I will talk about how security interacts and functions with One AI. However, One Model security is very customizable, and security configurations can vary by client. Please reach out to your customer success team with any site-specific questions or concerns about your security setup.

We will cover security within machine learning model creating, editing, and managing, and security within storyboards with model insights and components.

After watching, you will understand the ethical and strategic considerations of promoting full data access for users involved in building or managing machine learning models in One AI; become familiar with the necessary application access roles for effectively using machine learning in One AI; learn scenarios where data access roles are not fully recognized during the model creation and management process; and understand that storyboards with machine learning insights fully comply with One Model's role based security, so viewers don't need full access.

CHAPTER 2

Security in Model Creation, Editing, & Management

Section 2 - Security in Machine Learning Model Creation, Editing, and Management

We will start by discussing how One Model security interacts with machine learning models within the One AI tab. Model building is intended for users with full data access, usually an admin or people analytics or data science resource with temporary full access.

Model creators need this full access to be able to build robust, transparent models and accurately interpret them. This enables them to determine whether the model should be deployed and shared with stakeholders or if they need further refinement.

Full data access is essential for making informed decisions during model building, evaluation, and troubleshooting.
It ensures the creation of strong, high performing, and relevant models.
This comprehensive understanding helps model creators identify and address potential biases, fairness issues, or ethical concerns in the data resulting in more performant and effective models.

To grant users access to the One AI tab in the main ribbon menu, their role must include the CanAccessOneAIMenu application access role. This tab takes users to the window where they can create, edit, and manage machine learning models. Without this role, users cannot view, create, manage, or share machine learning models. However, they still can view model insights and results on storyboards without this role.

When building machine learning models in the One AI query builder, users can access every metric, dimension, and column within the One Model site regardless of their data access role. The One AI query builder does not follow the role-based security that is applied in the rest of One Model. We will go through some scenarios and examples to clarify this as it can be a bit complex and dense.

CHAPTER 2.1
Security Scenarios in Model Creation, Editing, & Management

For our first scenario, for example, if a user without access to diversity, equity, and inclusion (DEI) data is building a model, they can still add DEI columns like gender or ethnicity as core attributes or filter the model population by DEI dimensions.

In most cases, even if the model is otherwise built correctly, it will run with these unpermitted columns included as core input attributes, but they will be made fully null and automatically dropped.

Even if the unpermitted column is predictive, the model won't use it. While this might not matter for niche, single purpose columns; overall, better models are often created by users with full data access because otherwise we are missing out on predictive columns that are being automatically dropped.

For our next scenario, you must know that when creating or editing a model, the user selects which metric to use to define what they're predicting, such as voluntary terminations, new hire failure, etcetera, and which metric they will define the model population, such as headcount or external hires.

They will see the entire metric list for both, regardless of their data access role. If they select a metric they do not have permission to access, they can still create the model, but running it will result in failure.

If a user with access to the unpermitted metric then tries to run the model created by someone without access, it will still fail because the model always runs under the permissions of its original creator.

In order for that user with access to successfully run the model, they would have to recreate it themselves.

Conversely, if a user with access creates the model and another user without access edits or runs it, the model will still run successfully because again, it operates under the permissions of its original creator.

Therefore, it is advisable to restrict users without full data access from participating in machine learning model building to minimize confusion and mitigate data security risks.

For our next scenario, it's helpful to understand that model creators and editors can create generative attributes during the model building or editing process.

These are new input variables derived from the original model dataset that can be selected as features in the model. To create generative attributes, users need both the CanAccessOneAIMenu and CanConfigureOneAIGenerativeAttributes application access roles.

While optional for model creators and managers, the generative attribute role provides valuable tools for strengthening models by offering a range of input variables, so it's highly recommended.

With this role, users can view and use all non-calculated metrics available in the One Model site to build generative attributes regardless of their data access role. If they build a generative attribute using a metric they don't have access to and include it in the model dataset, the model will fail when run. However, they can still view the data produced by the generative attribute in the data downloads I'm going to discuss in the next section.

Moreover, even if a user lacks access to the metric used to create generative attributes, they can still use models with these attributes if they weren't the original creator. Because as I said before, the model operates under the permissions of the original creator.

Users with the can access One AI menu permission can download the train and test data and the predict data for any models on your site, regardless of who created them. These downloads contain row by row data identified by person ID or another unique identifier that was used to build the model, including values for each core attribute. This may involve generative attributes, salary information, DEI data, performance ratings, and any other data available in One Model.

Even if the columns are automatically dropped and shown as 100% null in the exploratory data analysis report for users without access, they can still see the actual row by row data in these reports.

Users who shouldn't have access to this information shouldn't have access to the One AI menu tab.

Additionally, users can view results from models they didn't create even if they don't have access to every metric and column used in the model.

From the EDA report, users can conduct variable analysis and download correlation information on all columns included in the model regardless of their access. They can also run and deploy models created by other users and analyze the data using Explore or in storyboards.

Therefore, to promote ethical model building and prevent confusion within the tool, users should have full data access if they will be creating or managing models within One AI. Insights and results from machine learning models can still be shared with users who do not have full data access through storyboards.

In the next section, we will discuss how security interacts with storyboards containing One AI machine learning insights and components.

CHAPTER 3

Security & Storyboards with Machine Learning Components

Section 3 - Security and Storyboards with Machine Learning Components

Now let's discuss how security interacts and functions within storyboards created to share machine learning insights and model results with stakeholders.

The good news is that this process is much more straightforward and similar to One Model's role based security throughout the rest of One Model.

As I mentioned earlier, users do not need full data access to view storyboards that display model insights and results. When a model creator is satisfied with their model, they can deploy it, making it available on the front end of One Model to be used in Explore or storyboards.

The first time this is done, a data engineer will add a script that generates the tables and dimensions needed to create the metrics and storyboards to display the model results.

These metrics, dimensions, columns, and storyboards follow standard permissioning within data access roles or right from the storyboard, allowing you to choose which groups have full, partial, or no access.

For instance, if a user lacks access to a particular data point, such as salary range, and there's a tile that segments predictions based on salary range, that tile will show as "Security Restricted" for that user, and no data will be displayed.

Creators should consider their storyboard audience while developing model insight storyboards.

They should be creative with tiles and filters to ensure that viewers see only information appropriate for their role and use case.

CHAPTER 4

Conclusion & Thanks

Understanding the intricacies of security within One AI is important for maintaining data integrity and confidentiality in machine learning projects. Throughout this module, we've covered the key principles of security in One AI, enabling you to make informed decisions in model creation, management, and storyboard development and sharing.

Balancing accessibility and data security is essential for fostering trust and optimizing your machine learning projects.

We recommend that model creators and managers have full data access, while those viewing machine learning storyboards do not need full access.

Remember that security configurations may vary among One Model clients, so consult with your customer success team for site-specific questions or concerns regarding One AI machine learning models. We are here to provide guidance and best practices to ensure a simple, secure experience. Happy modeling!