

SAS® for Containers Running on Domino Data Lab



Accelerate Your SAS® Analytics and Save on CapEx – Take It to the AWS Cloud.

SAS for Containers on Domino streamlines SAS users' ability to run data science workloads in the AWS cloud. Using SAS for Containers within Domino, users such as data scientists, statisticians and business analysts can create SAS programs, develop SAS models, and publish applications. Users can interact with a familiar SAS environment via SAS Studio in a web browser or use Python in Jupyter Notebooks – an open-source, web browser notebook-style interface.

Domino provides an open, end-to-end data science platform for users to develop, validate, deliver, and monitor models using their tool of choice. With SAS for Containers on Domino, users can readily launch multiple SAS environments for model building and deploy SAS simultaneously on top of Domino's elastic compute infrastructure in the public cloud. They do not need to have any prior experience with the container technologies. Every detail of each SAS experiment is tracked: data¹, code, SAS version, environment, discussions, parameters, and results. Users can build on each other's work, enhancing collaboration and uncovering breakthroughs.

SAS for Containers on Domino not only eases software management and deployment for IT, but also accelerates model development and delivery for data science teams.

Benefits

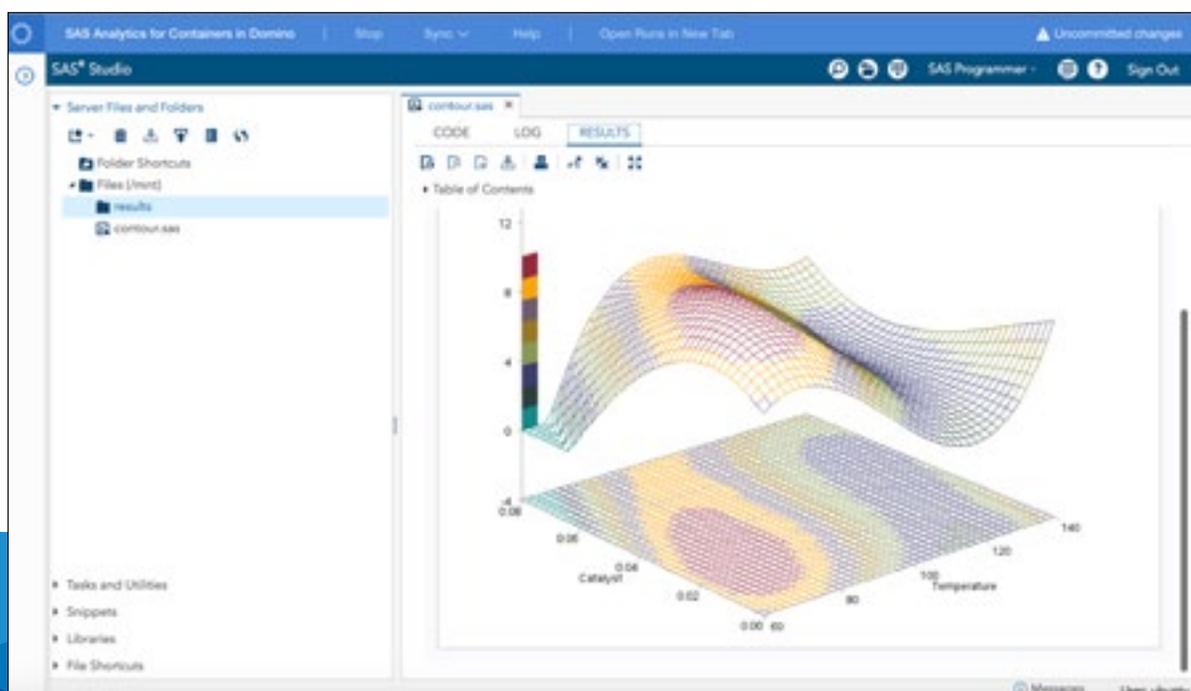
- Scale SAS model deployment. SAS models can be deployed² on top of Domino's elastic compute infrastructure, so that users can scale horizontally to support hundreds of models running simultaneously. Different versions can be deployed simultaneously on different instances for A/B testing, in a self-service environment.
- Accelerate SAS research with elastic compute and parallel experimentation. Domino's elastic compute grid can be hosted on a cloud platform such as AWS, which can reduce on-premises hardware costs. Data scientists can readily run SAS research workloads with their preferred compute resources in isolation with a single click, without having to wait for other workloads to finish. Additionally, data science teams can spin up parallel experimentation for faster iterations of model development and validation.



¹ Imported data, queries executed on external databases and query results are revisioned in Domino.

² Requires SASPy and Flask.

- Reproduce any experiment with one click and create a knowledge flywheel with collaboration. Running SAS for Containers on Domino provides a solid foundation for tracking and reproducibility. Domino tracks all versions of SAS projects and captures each experiment's state, including data³, code, SAS version, environment, discussions, parameters, and results. These projects can be shared amongst data science teams for collaboration and as the foundation for future projects. Capturing these artifacts helps users achieve model provenance and governance and provides auditability in their workflows. This capability is crucial for highly regulated industries such as health care, life sciences, finance, insurance, and others.
- Take advantage of the existing data lake/hub. Domino connects seamlessly to existing data sources such as Apache Hadoop. Put SAS to work on data stored in Hadoop and leverage the ability to execute Apache Spark within a Domino project.



To learn more about SAS for Containers, visit sas.com/containers.

To learn more about Domino Data Lab, visit dominodatalab.com.

Licensing

Domino and SAS have attractive licensing options available for joint customers. To learn more, visit sas.com/domino_questions.

³ See footnote 1.