

RapidMiner Server


Highly scalable, collaborative enterprise data science deployment platform

“RapidMiner allows us to go from an anecdotal approach to a data-supported approach for operational decisions. This enables us to create more meaningful interventions and provide better patient care.”

L.Ellenburg,
Senior Manager of Informatics

Collaboration for data analytics teams

RapidMiner Server is designed to improve the productivity of analytics teams so you can do more data science. You can share knowledge and best practices across your organization, inside a centralized repository, using built-in security controls to limit access as necessary. Optimize with the advanced queuing mechanism: RapidMiner Server can slice out resources and dedicate to teams, use cases or projects.

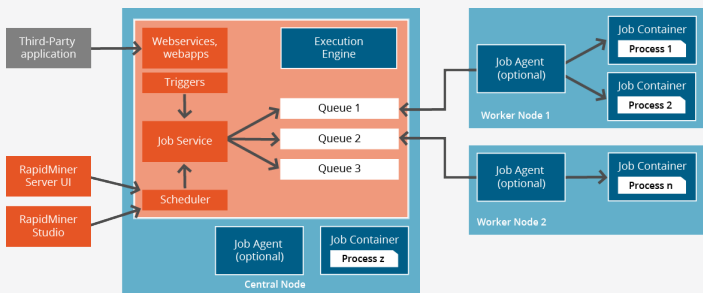
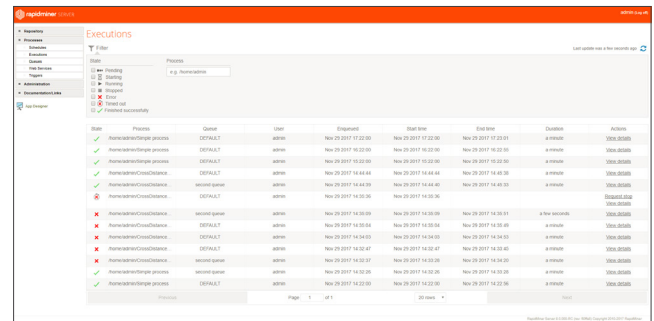


Speed predictive modeling

RapidMiner Server is designed to take full advantage of enterprise-grade compute environments. You can offload large jobs to the server, dramatically increasing modeling speed by taking advantage of high performance hardware.

Deploy models into production

Turn predictions into prescriptions by embedding predictive models into your business processes in just a few clicks. With RapidMiner Server, you can deploy predictive models into applications, databases, BI tools, and more. Models can also be consumed via APIs and web services, allowing for integrations with anything. Monitor model performance, detect model degradation, and react by automatically retraining to keep up with the changes. With just a few clicks, you can set-up scheduled processing to keep retraining models and continuously score data in real-time.



Highly scalable, distributed architecture


Highly scalable micro-services based architecture delivers reliable and lightning fast data science in production. Know that you're doing everything to be in control, reduce risk and downtime.

A distributed architecture means that now multiple Job Agents can be installed on multiple machines for process execution. This allows you to scale your environment as much as you require, easing your worries of future scalability as your data science teams grow and processes require more dedicated resources.

The distributed nature of containerized job executions allows for operational continuity even in case one 'rogue' process was to fail. The failing process is 'sandboxed' and can no longer affect the execution of other processes, creating a highly stable environment.

Real-time scoring

Predict at scale, with very low latency, and deliver actionable intelligence in real-time to the decision maker or machine. This back-end capability is designed for demanding use cases requiring very fast scoring, like predicting how your customers behave, when your industrial parts will break or, calculating the risks associated with an action or a client. Enable real-time online scoring from web portals, phone apps, or desktop applications.



Key features

Computation & scalability

- Computational service: Use the computational power of enterprise servers and free Studio resources for development
- Virtually unlimited scalability: Create RapidMiner Server clusters of any size by adding more Job Agents and machines to the environment

Queues

- Easy creation of queues for resource management: Resources can be split among users or use cases, they can be shared or made exclusive by the administrator depending on the company's structure
- Flexible environment configuration: multiple options when creating queues that can adapt to any environment

Scheduling

- Periodic scheduling of workflow executions within the Server's UI
- Set up actions based on a triggered event
- Remote execution of analysis processes

Integration & operationalization

- One-click deployment using Web Services: any process can be readily made into a published web service. Web Services allow integration with third-party tools like custom web consoles, BI tools and others.
- Web services / processes can deliver XML, JSON, static / dynamic visualizations and binary files among others
- The Real-time scoring agent allows you to predict at scale, with very low latency, and deliver actionable intelligence in real-time

Repository

- Shared repository with security controls. Groups can share or protect their models and processes as needed
- "Copy&paste" processes from Studio to Server or from Server to Server
- Fine-grained permissions for processes, models and data

Collaboration

- Reusable templates and processes: make processes available to the whole team to act as templates or best practices
- LDAP integration or RapidMiner user system
- Process version management

Management & monitoring

- Webapps: custom dashboards for monitoring, management and showcasing of process results
- Logging and auditing of executions
- Monitoring of current schedules and executions

Connections

- Define and protect your database connections for processes to securely access data

Extensions

- Use any of Studio extensions in a RapidMiner Server enterprise environment