



Increase service availability by deploying IBM TM1 Server on AWS

IBM Planning Analytics with Watson on AWS



IBM Planning Analytics with Watson on AWS: Increase service availability by deploying IBM TM1 Server on AWS

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

IBM Planning Analytics with Watson on AWS **i**

IBM Planning Analytics with Watson on AWS Diagram 1

Further reading 2


Diagram history 2

IBM Planning Analytics with Watson on AWS

Publication date: May 16, 2023 ([Diagram history](#))

Increase service availability by deploying IBM TM1 Server on AWS. IBM Planning Analytics is an integrated planning solution that uses artificial intelligence (AI) to automate planning, budgeting, and forecasting activities, and build intelligent workflows.

IBM Planning Analytics with Watson on AWS Diagram



1. Corporate users access IBM Planning Analytics through a private, secure network connection from their corporate office to AWS.
2. Planning Analytics Workspace runs on Red Hat OpenShift Service on AWS, which is a web-based interface that provides full modeling, reporting, and administrative capabilities.
3. IBM TM1 Web runs on **Amazon Elastic Compute Cloud console** instances across multiple Availability Zones for high availability. Users access TM1 server cubes and worksheets using TM1 Web.

4. TM1 Web and Planning Analytics Workspaces access the TM1 server using a DNS record in an **Amazon Route 53** Private Hosted Zone. In a failover scenario, users are redirected to the Stand-by TM1 server.
5. IBM TM1 data tier includes the TM1 admin server and the TM1 server. Both are deployed to the same **Amazon EC2** instance in active-passive failover architecture. The TM1 admin server keeps track of active TM1 servers, while TM1 server manages access to data directory for clients.
6. **Amazon Elastic Block Store** holds TM1 server binaries, data directory, transaction and application logs, and dumps. The data directory contains TM1 cubes, dimensions, and system information that are loaded into memory when the server starts. Cube value changes are stored in memory and in the transaction log. Data is saved to disk when the TM1 server shuts down, or when requested to be stopped by an administrator or a batch process.
7. Backup jobs create snapshots of the TM1 binaries, data and transaction logs volumes. These are used to restore the volumes in the Stand-by TM1 Server **Amazon EC2** instance.
8. TM1 server will save its cube data, from memory to disk, in a failover situation. An **Amazon FSx for Windows File Server** file system is used to replicate the TM1 data across Availability Zones. This data is used to rebuild the TM1 cubes, in memory, when the standby server starts.
9. Logs and metrics of the IBM TM1 server are pushed to **Amazon CloudWatch**, where alerts are defined to detect failures and send push notifications using **Amazon Simple Notification Service**.
- 10 Site reliability engineers receive **Amazon SNS** notifications and start failover of the IBM TM1 server.

Further reading

For additional information, refer to

- [AWS Architecture Icons](#)
- [AWS Architecture Center](#)
- [AWS Well-Architected](#)

Diagram history

To be notified about updates to this reference architecture diagram, subscribe to the RSS feed.

Change	Description	Date
Initial publication	Reference architecture diagram first published.	May 16, 2023

Note

To subscribe to RSS updates, you must have an RSS plugin enabled for the browser you are using.