Service Name

Amazon Timestream

Service Overview

Amazon Timestream is a fully managed, serverless time series database service designed for storing and analyzing trillions of time-series data points daily. It offers two engine options: Timestream for LiveAnalytics for real-time analytics with fast SQL queries, and Timestream for InfluxDB for open-source compatibility. The service automatically manages data lifecycle by keeping recent data in memory for fast access and moving historical data to cost-optimized magnetic storage. It provides built-in time series analytics functions, simplified data access across storage tiers, and seamless integration with AWS services for data collection, visualization, and machine learning applications.

Key Use Cases

- IoT sensor data storage and analysis: Collecting and analyzing massive volumes of sensor data from connected devices with real-time monitoring capabilities
- Industrial telemetry management: Managing operational data from manufacturing equipment, power systems, and infrastructure monitoring with predictive maintenance insights
- Application performance monitoring: Tracking application metrics, user interactions, and system performance data for real-time insights and anomaly detection
- Infrastructure monitoring: Collecting and analyzing system metrics, logs, and operational data from servers, networks, and cloud resources for observability

Features

- **Serverless Auto-scaling**: Fully managed service with automatic capacity adjustment based on workload demands
- **Dual Storage Architecture**: Memory store for high-throughput writes and fast queries, magnetic store for cost-effective long-term storage
- **Built-in Time Series Functions**: SQL with time series-specific extensions for analytics, aggregations, and data processing
- **Timestream Compute Units (TCUs)**: Serverless compute capacity with 4 vCPUs and 16 GB memory per unit for predictable query performance
- **Multi-AZ Replication**: Data automatically replicated across three Availability Zones for high availability and durability
- **Data Lifecycle Management**: User-configurable policies for automatic data tiering between memory and magnetic storage

- **Always Encrypted**: Data encrypted at rest and in transit with AWS KMS integration for customer-managed keys
- **Multi-measure Records**: Store multiple measures in single table rows to reduce storage costs and improve query performance

Value Proposition

Amazon Timestream provides significant advantages over traditional relational databases and self-managed solutions for time-series workloads. It delivers up to 10x cost reduction compared to relational databases while handling trillions of events per day with millisecond query response times:

- The serverless architecture eliminates infrastructure management overhead, while automatic data lifecycle management optimizes costs by intelligently tiering data between fast memory storage and cost-effective magnetic storage
- Built-in time series functions and SQL compatibility accelerate development, and seamless AWS service integrations enable comprehensive analytics pipelines
- The service provides enterprise-grade security, compliance certifications, and 99.999999999% data durability with Multi-AZ replication

Service Integration

Amazon Timestream integrates deeply with core AWS services for comprehensive time-series solutions. Data ingestion flows through AWS IoT Core for device data, Amazon Kinesis for streaming analytics, Amazon MSK for Kafka workloads, and AWS Lambda for serverless processing:

- EventBridge Pipes enables real-time data streaming from Kinesis with transformation capabilities
- Analytics integrations include Amazon QuickSight for visualization, Amazon SageMaker for machine learning, Amazon Athena for ad-hoc queries via UNLOAD to S3, and Amazon Managed Grafana for operational dashboards
- Security and management leverage AWS KMS for encryption, IAM for access control, VPC for network isolation, AWS Backup for table backups, and CloudWatch for monitoring with custom metrics and alarms

Onboarding and Offboarding

Customers can quickly start with Amazon Timestream using the AWS Management Console tutorial that creates a database with sample IoT and DevOps datasets, or explore the fully functional sample applications available in multiple programming languages (Java, Python, Node.js, Go, .NET). The getting started process involves creating a database and table through the console, configuring retention policies for memory and magnetic stores, and connecting via AWS SDKs, CLI, or APIs:

- Prerequisites include AWS account setup, IAM permissions configuration, and optionally VPC and security group setup for network isolation
- Sample applications on GitHub demonstrate database creation, data ingestion, and query execution patterns to accelerate development and best practices adoption

Getting Started Guide:
https://docs.aws.amazon.com/timestream/latest/developerguide/getting-started.html

Data Removal

Data removal during offboarding involves deleting individual tables, entire databases, or complete service resources. Tables can be deleted through the console, CLI, or APIs with permanent data removal across all storage tiers. Database deletion removes all contained tables and associated data. For Timestream for InfluxDB clusters, deletion permanently removes all cluster nodes and data stored in InfluxDB Object Storage on Amazon S3 with no recovery options. Critical data should be exported using InfluxDB tools before deletion. The deletion process takes 5-10 minutes for clusters, and billing stops immediately upon completion.

Backup/Restore and Disaster Recovery

Amazon Timestream provides backup capabilities through AWS Backup integration for centralized data protection. Backups are created at table granularity with the first backup being full and subsequent backups incremental. Cross-region and cross-account backup copying enables disaster recovery scenarios. Restore operations can recreate tables in different regions or accounts with configurable retention settings. Multi-AZ data replication provides built-in resilience against zone failures. However, Timestream for InfluxDB lacks native snapshot capabilities, requiring data export for backup purposes.

Backup Capabilities

Timestream for LiveAnalytics integrates with AWS Backup for comprehensive table backup management. Users can create on-demand backups or scheduled backups through backup plans with configurable retention periods and lifecycle policies:

- Incremental backups optimize storage costs while maintaining complete recovery capabilities
- AWS Backup provides centralized policy management, cross-region replication, and automated backup scheduling
- Timestream for InfluxDB does not have native backup capabilities and requires manual data export using InfluxDB tools for data protection

Disaster Recovery

Amazon Timestream supports disaster recovery through multi-region backup copying and restoration capabilities via AWS Backup. Cross-region backup replication enables recovery in different AWS regions during regional outages:

- The service's Multi-AZ architecture provides automatic failover within regions for high availability
- Application Recovery Controller can orchestrate traffic routing during failures
- However, Timestream for InfluxDB has limited disaster recovery options due to lack of native backup features, requiring custom data export and replication strategies

Recovery Objectives

Amazon Timestream helps customers achieve recovery objectives through automated Multi-AZ replication for sub-minute RPO within regions and cross-region backup capabilities for disaster recovery. The serverless architecture enables rapid scaling during recovery scenarios. AWS Backup's incremental backup strategy minimizes recovery time while backup automation reduces RPO. For RTO objectives, the service's cellular architecture and auto-scaling capabilities support fast recovery and capacity restoration. Customers can implement backup and restore strategies with 2-4 hour RTO/RPO for most workloads through proper backup planning and cross-region replication.

Service Constraints

Service Quotas: https://docs.aws.amazon.com/general/latest/gr/timestream.html

FAQ: https://aws.amazon.com/timestream/faqs/

Technical Requirements

Service Documentation: https://docs.aws.amazon.com/timestream/latest/developerguide/what-is-timestream.html

User Guide: https://docs.aws.amazon.com/timestream/latest/developerguide/getting-started.html

API Reference: https://docs.aws.amazon.com/timestream/latest/developerguide/API_Reference.html

Pricing Overview

Please see the AWS UK G Cloud 15 Pricing Document accompanying this service in the Digital Marketplace.

Public Pricing: https://aws.amazon.com/timestream/pricing/

Cost Optimization

Amazon Timestream offers unique cost optimization through intelligent data lifecycle management with automatic tiering between memory and magnetic storage based on data age and access patterns. The Timestream Compute Units (TCU) model provides predictable query costs with per-second billing and configurable MaxQueryTCU limits for budget control:

- Multi-measure records reduce storage costs by storing multiple metrics in single rows
- Batch writes optimize ingestion costs, while query optimization through time predicates, measure name filtering, and spatial pruning reduces compute consumption
- The adaptive query engine automatically optimizes performance and costs, and scheduled queries enable automated data processing to minimize interactive query expenses

Savings Considerations

Primary cost savings come from Timestream's serverless architecture eliminating infrastructure management overhead and providing 10x cost reduction versus relational databases for time-series workloads. Automatic data lifecycle management significantly reduces storage costs by moving older data to cost-optimized magnetic storage while maintaining query access:

- The service's compression and optimized storage format minimize storage footprint
- Pay-per-use pricing for writes, storage, and queries ensures customers only pay for actual consumption without pre-provisioning
- Batch writing reduces per-record costs, while query optimization techniques like limiting scanned data and using time filters minimize compute charges
- AWS Backup's incremental backup strategy reduces backup storage costs compared to full backup alternatives