



Amazon DocumentDB Observability

Performance Insights + DML Auditing

Douglas Bonser

Sr. Specialist Solution Architect
Amazon DocumentDB

Agenda

- Amazon DocumentDB Observability and Monitoring
- Demo
 - Performance Insights
 - Data Manipulation Language (DML) Auditing
- Resources and Next Steps

Amazon DocumentDB

SCALABLE, HIGHLY-DURABLE, AND FULLY MANAGED DATABASE SERVICE FOR OPERATING MONGODB API-BASED WORKLOADS



Fast and scalable

- Scale compute in minutes
- Storage and I/O autoscaling
- Storage scales to 64 TiB
- Scale out to 15 replicas for millions of reads
- Globally distributed



Fully managed

- Built-in high availability
- Backups enabled by default
- Durable by default
- Security best practices by default
- Automatic patching
- Monitoring and alerting




MongoDB compatible

- Applications, drivers, and tools can be used with little or no change
- Supports hundreds of APIs, operators, and stages
- Continually working backward from customers to deliver the capabilities they need

Observability and Monitoring

- Amazon CloudWatch
- Profiler
- Event Subscriptions
- Performance Insights
- Auditing

Amazon CloudWatch




Gather and analyze operational metrics



Instance, cluster, storage, & backup metrics

Monitor in Amazon DocumentDB console



Create metric and composite alarms

Profiler



Identify slow queries



Discover opportunities for new indexes

Detect optimization for existing indexes



Improve query performance

Profiler CloudWatch logs example

```
{
  "op": "query",
  "ts": 1621471508420,
  "ns": "world.people",
  "command": {
    "find": "people",
    "filter": {
      "firstName": "afae"
    },
    ...
  },
  "nreturned": 0,
  "milis": 1861,
  "planSummary": "COLLSCAN",
  "execstats": {
    "stage": "COLLSCAN",
    "nReturned": 0,
    "executionTimeMilliststimate": "1860.690",
  },
  ...
}
```

Timestamp

Profiler CloudWatch logs example

```
{
  "op": "query",
  "ts": 1621471508420,
  "ns": "world.people",
  "command": {
    "find": "people",
    "filter": {
      "firstName": "afae"
    },
    ...
  },
  "nreturned": 0,
  "milis": 1861,
  "planSummary": "COLLSCAN",
  "execstats": {
    "stage": "COLLSCAN",
    "nReturned": 0,
    "executionTimeMilliststimate": "1860.690",
  },
  ...
}
```

Timestamp

Namespace (database.collection)

Profiler CloudWatch logs example

```
{
  "op": "query",
  "ts": 1621471508420,
  "ns": "world.people",
  "command": {
    "find": "people",
    "filter": {
      "firstName": "afae"
    }
  },
  ...
},
"nreturned": 0,
"milis": 1861,
"planSummary": "COLLSCAN",
"execstats": {
  "stage": "COLLSCAN",
  "nReturned": 0,
  "executionTimeMilliststimate": "1860.690",
},
...
}
```

Timestamp

Namespace (database.collection)

Command

`find({"firstName": "afae"})`

Profiler CloudWatch logs example

```
{
  "op": "query",
  "ts": 1621471508420,
  "ns": "world.people",
  "command": {
    "find": "people",
    "filter": {
      "firstName": "afae"
    }
  },
  ...
},
"nreturned": 0,
"milis": 1861,
"planSummary": "COLLSCAN",
"execstats": {
  "stage": "COLLSCAN",
  "nReturned": 0,
  "executionTimeMilliststimate": "1860.690",
},
...
}
```

Timestamp

Namespace (database.collection)

Command

`find({"firstName": "afae"})`

Execution stage

Profiler CloudWatch logs example

```
{
  "op": "query",
  "ts": 1621471508420,
  "ns": "world.people",
  "command": {
    "find": "people",
    "filter": {
      "firstName": "afae"
    }
  },
  ...
},
"nreturned": 0,
"milis": 1861,
"planSummary": "COLLSCAN",
"execstats": {
  "stage": "COLLSCAN",
  "nReturned": 0,
  "executionTimeMilliststimate": "1860.690",
},
...
}
```

Timestamp

Namespace (database.collection)

Command

`find({"firstName": "afae"})`

Execution stage

Execution time

Event Subscriptions




Cluster, instance, cluster snapshot, and parameter group categories



Subscription is at the category level

Subscribe to events on all or specific resources



Subscribe to all or specific events on resource(s)

Event Subscriptions

Source	Type	Message
sample-data	db-instance	DB instance shutdown
sample-data3	db-instance	DB instance shutdown
db-parameter-group	db-parameter-group	Updated parameter profiler to enabled with apply method immediate
sample-data3	db-instance	DB instance restarted
sample-data2	db-instance	DB instance restarted
sample-data3	db-instance	DB instance shutdown
sample-data2	db-instance	DB instance shutdown
demo-clone-manual-snapshot	db-cluster-snapshot	Creating manual cluster snapshot
getting-started-with-documentdb	db-cluster	Completed failover to DB instance: getting-started-with-documentdb

Event Subscriptions

Source	Type	Message
sample-data	db-instance	DB instance shutdown
sample-data3	db-instance	DB instance shutdown
db-parameter-group	db-parameter-group	Updated parameter profiler to enabled with apply method immediate
sample-data3	db-instance	DB instance restarted
sample-data2	db-instance	DB instance restarted
sample-data3	db-instance	DB instance shutdown
sample-data2	db-instance	DB instance shutdown
demo-clone-manual-snapshot	db-cluster-snapshot	Creating manual cluster snapshot
getting-started-with-documentdb	db-cluster	Completed failover to DB instance: getting-started-with-documentdb

Event Subscriptions

Source	Type	Message
sample-data	db-instance	DB instance shutdown
sample-data3	db-instance	DB instance shutdown
db-parameter-group	db-parameter-group	Updated parameter profiler to enabled with apply method immediate
sample-data3	db-instance	DB instance restarted
sample-data2	db-instance	DB instance restarted
sample-data3	db-instance	DB instance shutdown
sample-data2	db-instance	DB instance shutdown
demo-clone-manual-snapshot	db-cluster-snapshot	Creating manual cluster snapshot
getting-started-with-documentdb	db-cluster	Completed failover to DB instance: getting-started-with-documentdb

Event Subscriptions

Source	Type	Message
sample-data	db-instance	DB instance shutdown
sample-data3	db-instance	DB instance shutdown
db-parameter-group	db-parameter-group	Updated parameter profiler to enabled with apply method immediate
sample-data3	db-instance	DB instance restarted
sample-data2	db-instance	DB instance restarted
sample-data3	db-instance	DB instance shutdown
sample-data2	db-instance	DB instance shutdown
demo-clone-manual-snapshot	db-cluster-snapshot	Creating manual cluster snapshot
getting-started-with-documentdb	db-cluster	Completed failover to DB instance: getting-started-with-documentdb

Performance Insights



Visualize database performance




Determine what is causing database load

Determine when there is database load



Alert on database load

Database load



Auditing



Audit operations performed in a cluster



Log cluster operations

Log data operations



Alarm on specific operations

Supported events

- Data Definition Language (DDL)
 - Database, collection, index, and change stream events
 - Connection
 - Authorization
 - Role Management
 - User Management
- Data Manipulation Language (DML)
 - Create, read, update, and delete (CRUD) operations

DML auditing CloudWatch logs example

```
{
  "atype": "authCheck",
  "ts": 1663619344841,
  "timestamp_utc": "2022-09-19 20:29:04.841",
  "remote_ip": "127.0.0.1:40788",
  "users": [ {
    "user": "labUser",
    "db": "demo"
  } ],
  "param": {
    "command": "find",
    "ns": "demo.dml",
    "args": {
      "batchSize": 101,
      "filter": { "version": 1 },
      "find": "dml",
      "limit": 1,
      "projection": { "message": 1 },
      ...
    },
  },
  "result": 0
}
```

Timestamp

DML auditing CloudWatch logs example

```
{
  "atype": "authCheck",
  "ts": 1663619344841,
  "timestamp_utc": "2022-09-19 20:29:04.841",
  "remote_ip": "127.0.0.1:40788",
  "users": [ {
    "user": "labUser",
    "db": "demo"
  } ],
  "param": {
    "command": "find",
    "ns": "demo.dml",
    "args": {
      "batchSize": 101,
      "filter": { "version": 1 },
      "find": "dml",
      "limit": 1,
      "projection": { "message": 1 },
      ...
    },
  },
  "result": 0
}
```

Timestamp

Source

DML auditing CloudWatch logs example

```
{
  "atype": "authCheck",
  "ts": 1663619344841,
  "timestamp_utc": "2022-09-19 20:29:04.841",
  "remote_ip": "127.0.0.1:40788",
  "users": [ {
    "user": "labUser",
    "db": "demo"
  } ],
  "param": {
    "command": "find",
    "ns": "demo.dml",
    "args": {
      "batchSize": 101,
      "filter": { "version": 1 },
      "find": "dml",
      "limit": 1,
      "projection": { "message": 1 },
      ...
    },
  },
  "result": 0
}
```

Timestamp

Source

Database user

DML auditing CloudWatch logs example

```
{
  "atype": "authCheck",
  "ts": 1663619344841,
  "timestamp_utc": "2022-09-19 20:29:04.841",
  "remote_ip": "127.0.0.1:40788",
  "users": [ {
    "user": "labUser",
    "db": "demo"
  } ],
  "param": {
    "command": "find",
    "ns": "demo.dml",
    "args": {
      "batchSize": 101,
      "filter": { "version": 1 },
      "find": "dml",
      "limit": 1,
      "projection": { "message": 1 },
      ...
    },
  },
  "result": 0
}
```

The diagram consists of four orange lines on the right side of the JSON log entry. The top line connects the label 'Timestamp' to the 'timestamp_utc' field. The second line connects 'Source' to the 'remote_ip' field. The third line connects 'Database user' to the 'user' field within the 'users' array. The bottom line connects 'Namespace (database)' to the 'ns' field within the 'param' object.

DML auditing CloudWatch logs example

```
{
  "atype": "authCheck",
  "ts": 1663619344841,
  "timestamp_utc": "2022-09-19 20:29:04.841",
  "remote_ip": "127.0.0.1:40788",
  "users": [ {
    "user": "labUser",
    "db": "demo"
  } ],
  "param": {
    "command": "find",
    "ns": "demo.dml",
    "args": {
      "batchSize": 101,
      "filter": { "version": 1 },
      "find": "dml",
      "limit": 1,
      "projection": { "message": 1 },
      ...
    },
  },
  "result": 0
}
```

Timestamp

Source

Database user

Namespace (db)

DML command



Demo

Getting Started

Amazon DocumentDB Free Trial



What's included?

750 hours – t3.medium instance
(2 vCPUs, 4GB RAM)

5 GB storage

30 million I/Os

5 GB backup

Thank You



Q & A