

Backing up and Recovering DynamoDB Tables



Stefan Roman

DEVOPS ENGINEER

www.katapult.cloud




On-demand Backup and Restore



Backups for long term data archival



No performance degradation




Fast backup process



No additional cost except storage



On-Demand Backup and Restore



On-Demand Backup and Restore




What gets backed up and restored?

- Secondary indices
- Streams
- Provisioned read and write capacity



On-Demand Backup and Restore




What does not get backup and restored?

- Autoscaling policies
- Access policies
- CloudWatch metrics and alarms
- Stream settings
- TTL settings



Demo



Create DynamoDB table

Create items and create backup

Restore the backup

Observe configuration of the new table



Point-in-Time Recovery



Protects from accidental deletion



No scheduled on-demand backups needed



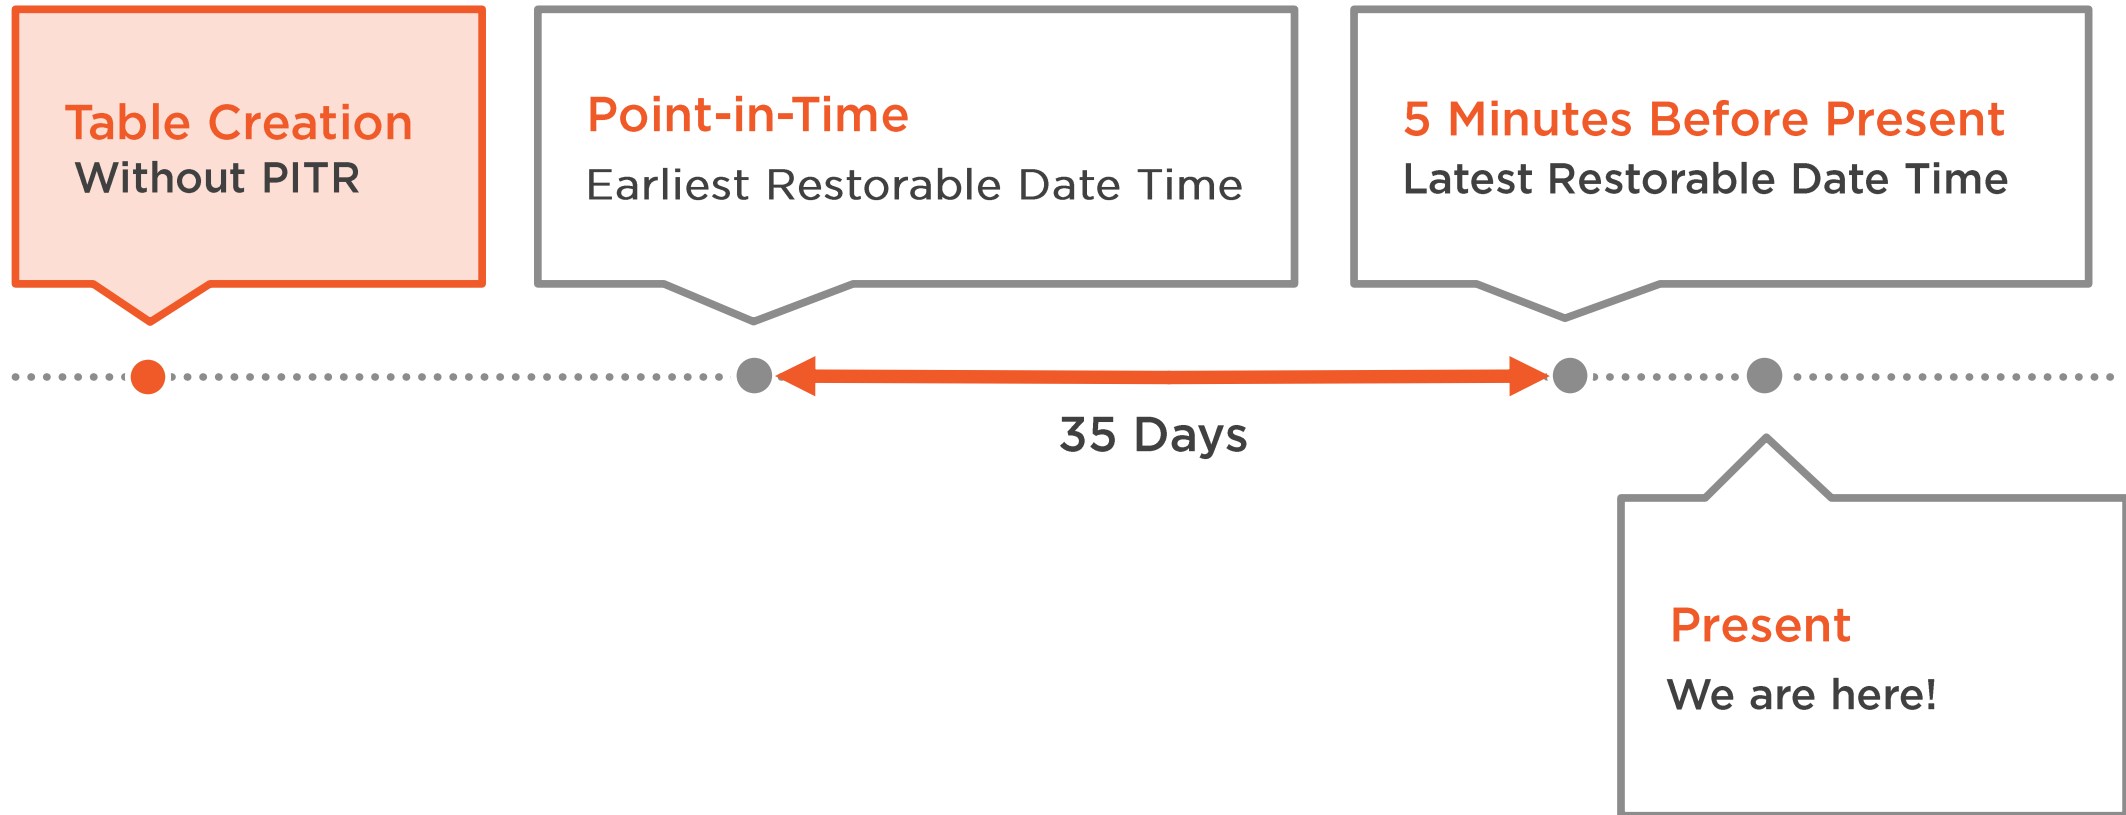
Available up to 35 days



Restorable to 5 minutes from current time



Point-in-Time Recovery



Point-in-Time Recovery




What is restored?

- Secondary indices
- Provisioned read and write capacity
- Encryption settings



Point-in-Time Recovery



What is not restored?

- Auto scaling policies
- Access policies
- CloudWatch metrics and alarms
- Stream settings
- TTL settings
- Point-in-time recovery settings



Point-in-Time Recovery



On-Demand Backup

Designed for long term data archival




Point-in-Time Recovery

Designed as a solution for disaster recovery



Demo



Enable Point-in-Time recovery

Delete some items

Create a table using Point-in-Time recovery



Summary



Backing up and recovering

- On-demand
- Point-in-Time

Recovering removes some settings

Use cases for each solution

Backing up does not degrade performance

Observed item restoration

