

The background of the image is a dark navy blue. On the right side, there are large, overlapping, semi-transparent shapes in shades of purple and magenta. Two thin, light blue lines cross the image diagonally from the top right towards the bottom left. The text 'AWS re:Invent' is positioned on the left side in a white, sans-serif font.

AWS re:Invent

DECEMBER 2 – 6, 2024 | LAS VEGAS, NV

OPN402

Gain expert-level knowledge about Powertools for AWS Lambda

Andrea Amorosi

(he/him)

Senior SA Engineer, Powertools
Amazon Web Services



© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

01 Powertools for AWS Lambda

02 Structured logging

03 Event handler

04 Idempotency

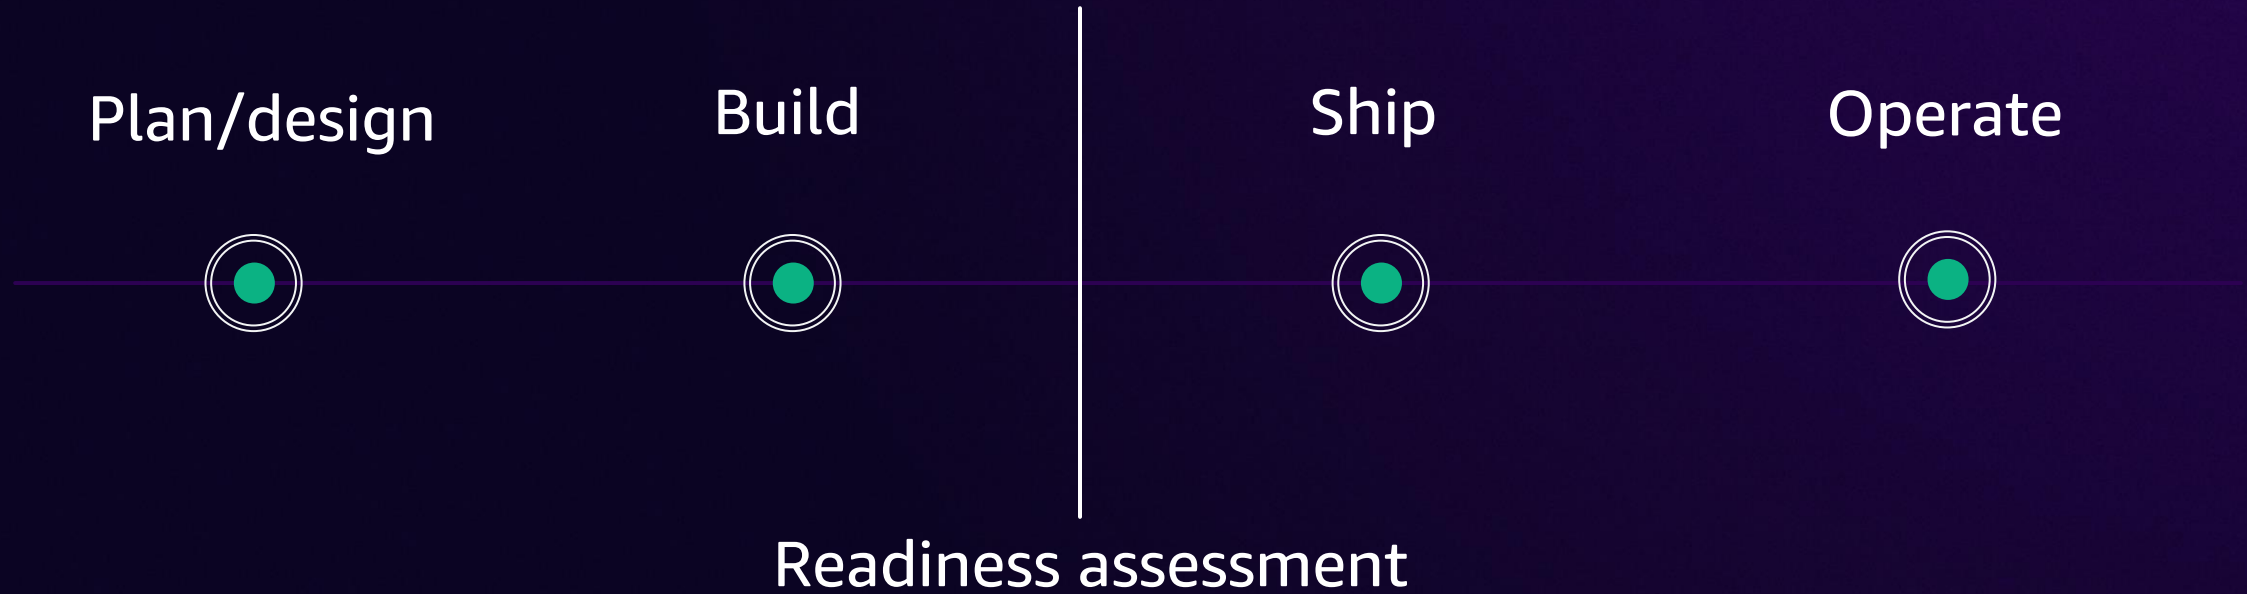
05 Batch processing

06 Roadmap and wrap up

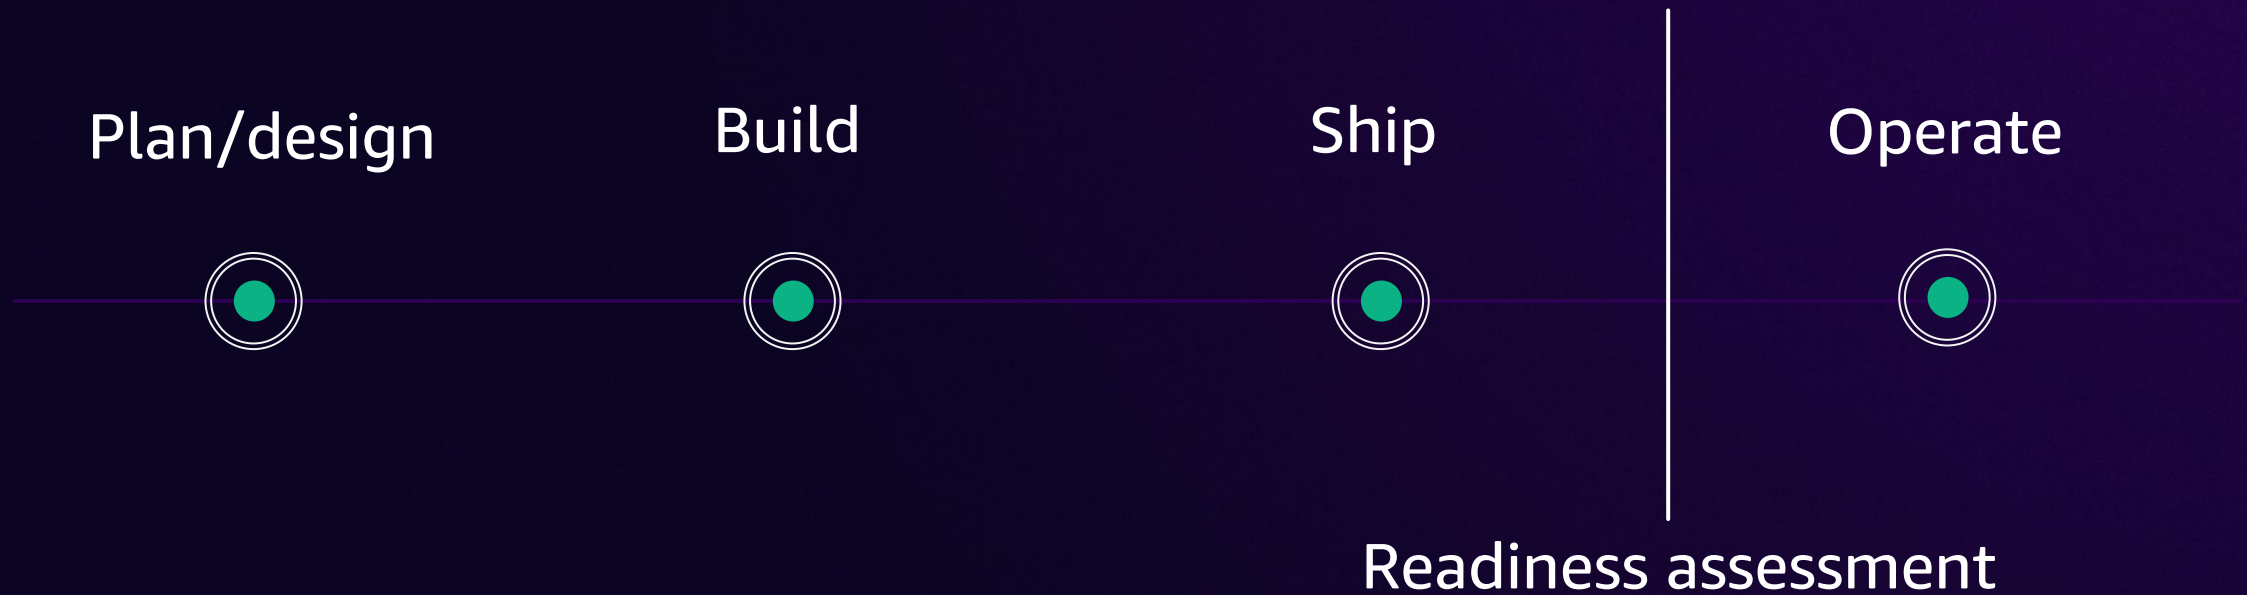
Systems development life cycle



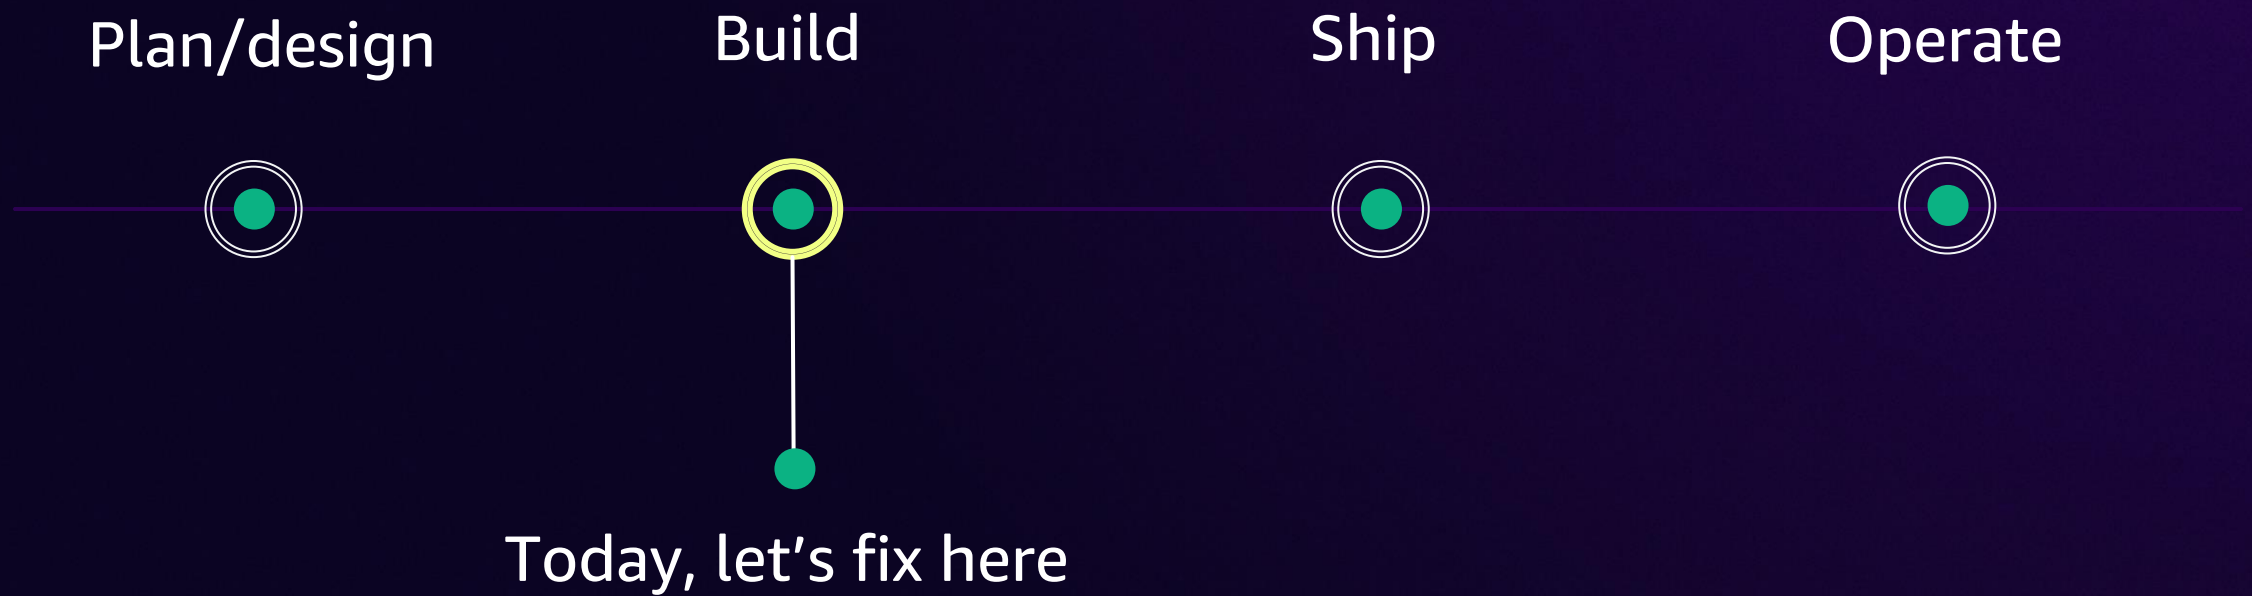
Findings come in late



Findings come in late



Let's fix this



Powertools for AWS Lambda

Logger

Metrics

Tracer

+ Many more

A toolkit to implement
serverless **best practices** and
increase **developer velocity**

Powertools for AWS Lambda: Toolkit

Best practices for **everyone**

Python | TypeScript | Java | .NET

Batch processing

Config management

Observability

Secrets handling

REST/GraphQL API

Idempotency

Input/output validation

BYO middleware

Self-documented schemas

Feature flags

Caching

Streaming

Data extraction

*feature set may vary across languages

Powertools for AWS Lambda

200B+
requests/week

40%
Community developed

Structured logging



Structured logging: Recap

Raw

Semi-structured

Canonical

Structured

Structured logging: Without Powertools

Raw

```
1 import logging
2
3 logger = logging.getLogger()
4 logger.setLevel(logging.INFO)
5
6 logger.info("Hello world")
```

[INFO] 2024-12-02T22:.. 1c8df7d3... Hello world

Structured logging: Without Powertools

Semi-
structured

```
1 import logging
2
3 logger = logging.getLogger()
4 logger.setLevel(logging.INFO)
5
6 logger.info({"message": "Hello world"})
```

[INFO] 2024-12-02T22:... ... {"message": "Hello world"}

Structured logging: Without Powertools

Canonical

```
1 import logging
2 import os
3
4 from logfmt import Logfmt
5
6 handler = logging.StreamHandler()
7 handler.setFormatter(Logfmt())
8 logging.basicConfig(handlers=[handler], level=os.getenv("LOG_LEVEL", "INFO"))
9
10 logging.info("Aha!", extra={"request_latency": 0.1})
11
```

at=INFO msg=Aha! request_latency=0.1

Structured logging: Without Powertools

Structured

```
1 import logging
2 import os
3
4 from pythonjsonlogger import jsonlogger
5
6 logger = logging.getLogger()
7 structured_handler = logging.StreamHandler()
8 formatter = jsonlogger.JsonFormatter(
9     fmt='%(asctime)s %(levelname)s %(name)s %(message)s'
10 )
11
12 structured_handler.setFormatter(formatter)
13 logger.addHandler(structured_handler)
14 logger.setLevel(os.getenv("LOG_LEVEL", "INFO"))
15
16 logger.info("Aha!")
17
```

Structured logging: Without Powertools

Structured

```
1 import logging
2 import os
3
4 from pythonjsonlogger import jsonlogger
5
6 logger = logging.getLogger()
7 structured_handler = logging.StreamHandler()
8 formatter = jsonlogger.JsonFormatter(
9     fmt='%(asctime)s %(levelname)s %(name)s %(message)s'
10 )
11
12 structured_handler.setFormatter(formatter)
13 logger.addHandler(structured_handler)
14 logger.setLevel(os.getenv("LOG_LEVEL", "INFO"))
15
16 logger.info("Aha!")
17
```

```
1 {
2     "asctime": "2024-12-02...",
3     "levelname": "INFO",
4     "name": "root",
5     "message": "Aha!"
6 }
```

Structured logging: Powertools

Structured



```
1 from aws_lambda_powertools import Logger
2
3 logger = Logger(service="payment")
4
5 logger.info("Aha!")
```


Structured logging: Powertools

Structured

```
1 from aws_lambda_powertools import Logger
2
3 logger = Logger(service="payment")
4
5 logger.info("Aha!")
```

```
1 {
2   "level": "INFO",
3   "location": "main.handler:5",
4   "message": "Aha!",
5   "timestamp": "2024-12-02...",
6   "service": "payment",
7   "sampling_rate": 0.0
8 }
```

Structured logging: Powertools

Structured

```
1 from aws_lambda_powertools import Logger
2 from aws_lambda_powertools.logging import correlation_paths
3
4 logger = Logger(service="payment")
5
6 @logger.inject_lambda_context(correlation_id_path=correlation_paths.API_GATEWAY_REST)
7 def handler(event, context):
8     logger.info("Aha!")
9
```

Structured logging: Powertools

Structured

```
1 from aws_lambda_powertools import Logger
2 from aws_lambda_powertools.logging import correlation_paths
3
4 logger = Logger(service="payment")
5
6 @logger.inject_lambda_context(correlation_id_path=correlation_paths.API_GATEWAY_REST)
7 def handler(event, context):
8     logger.info("Aha!")
9
```

```
1 {
2   "level": "INFO",
3   "location": "main.handler:8",
4   "message": "Aha!",
5   "timestamp": "2024-12-02 11:47:12",
6   "service": "payment",
7   "cold_start": true,
8   "lambda_function_name": "test",
9   "lambda_function_memory_size": 128,
10  "lambda_function_arn": "arn:aws:lambda:eu-west-1:123456789012:function:test",
11  "lambda_request_id": "52fdcf07-2182-154f-164f-5f0f9a621d84",
12  "correlation_id": "5a2f91db-c398-4afe-8d89-089df20c5cdc"
13 }
```

Structured logging: Powertools

BYO formatter

```
1 from aws_lambda_powertools import Logger
2 from aws_lambda_powertools.logging.formatter import LambdaPowertoolsFormatter
3 from aws_lambda_powertools.logging.types import LogRecord
4
5
6 class CustomFormatter(LambdaPowertoolsFormatter):
7     def serialize(self, log: LogRecord) -> str:
8         return self.json_serializer({
9             "level": log["level"],
10            "message": log["message"],
11            "timestamp": log["timestamp"],
12            "correlation_ids": {
13                "api_request_id": log["correlation_id"],
14                "service": log["service"],
15                "fn_request_id": log["lambda_request_id"]
16            },
17            "function_metadata": {
18                "name": log["lambda_function_name"],
19                "memory": log["lambda_function_memory_size"],
20                "arn": log["lambda_function_arn"],
21                "cold_start": log["cold_start"]
22            }
23        })
24
25
26 logger = Logger(service="payment", logger_formatter=CustomFormatter())
27 logger.info("hello")
```

Structured logging: Powertools



```
1 { "level": "INFO", "message": "validated user token", "timestamp": "2021-05-3 11:47:12", "customer_id": "868a2c8ef", ... }  
2 { "level": "INFO", "message": "updated subscription", "timestamp": "2021-05-3 11:47:30", "customer_id": "868a2c8ef", ... }  
3 { "level": "INFO", "message": "email notification ok", "timestamp": "2021-05-3 11:47:30", "customer_id": "868a2c8ef", ... }
```


Structured logging: Powertools

Wide logs

```
1 from aws_lambda_powertools import Logger
2
3 logger = Logger(service="payment")
4
5 @logger.inject_lambda_context()
6 def handler(event, context):
7     customer = validate_user_token(event)
8     logger.append_keys(customer_id=customer.id)
9
10    try:
11        subscription_id = update_subscription(customer_id=customer.id)
12        logger.append_keys(subscription_id=subscription_id)
13    except Exception as exc:
14        logger.append_keys(subscription_update_fault=exc)
15
16    notification_id = send_email_notification(email=customer.email)
17    logger.append_keys(notification_id=notification_id)
18
19    logger.info('subscription process')
```

Structured logging: Powertools

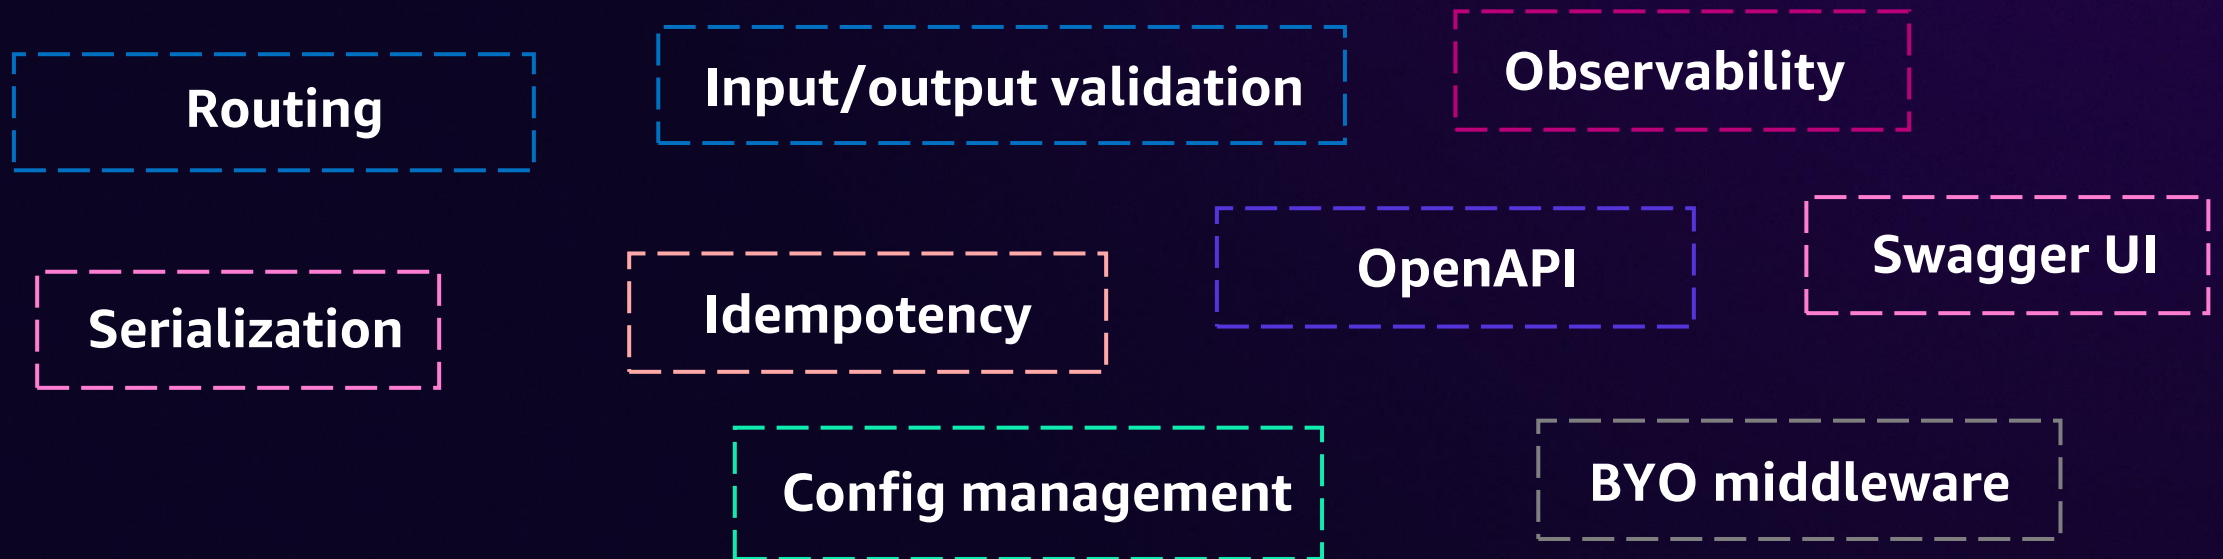

Wide logs

```
1 {
2   "level": "INFO",
3   "message": "subscription process",
4   "timestamp": "2024-12-02 11:47:12",
5   "customer_id": "868a2c8ef",
6   "subscription_id": "4f5ccf70f56c4998b191ba3af7bf3783",
7   "notification_id": "0b478d8ca6fc43388ba6",
8   "correlation_ids": {
9     "service": "payment",
10    "fn_request_id": "addc193-497c-a446-290d-16002cb6b134a"
11  },
12  "function_metadata": {
13    "name": "test",
14    "memory": 128,
15    "arn": "arn:aws:lambda:eu-west-1:123456789012:function:test",
16    "cold_start": false
17  }
18 }
```

Event handler



A common pattern



API design & trade-offs

How many Lambda functions do you need for an API?

API design & trade-offs


Lambda
monolith

?

Micro-functions

Lambda monolith

/orders



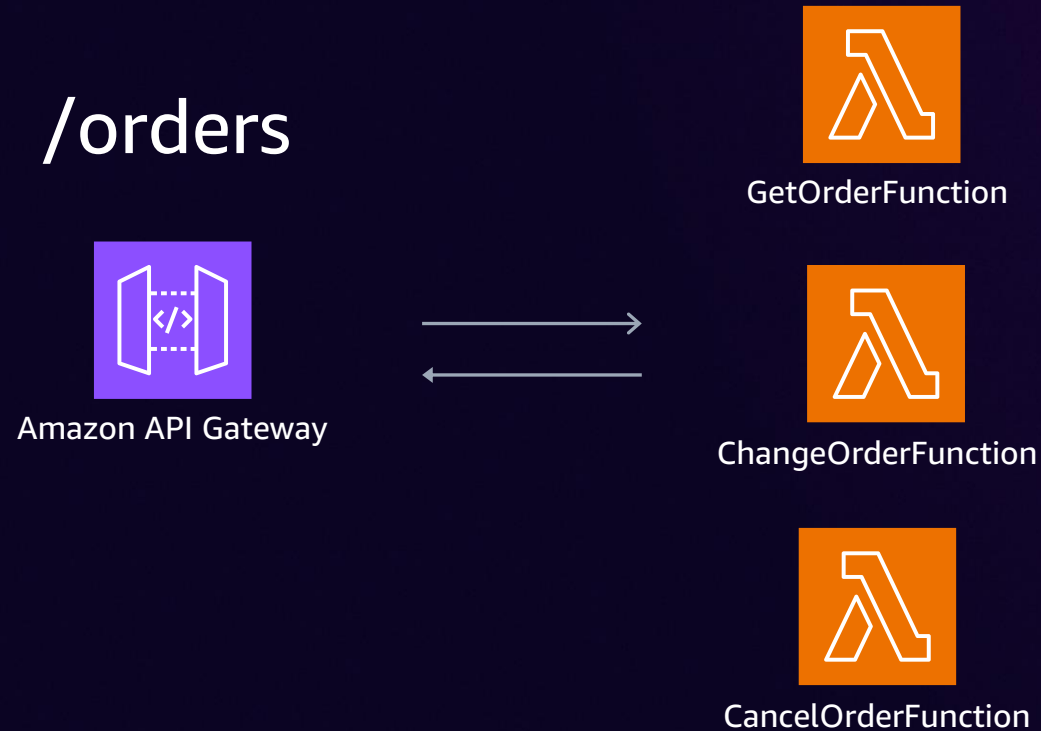
Amazon API Gateway



AWS Lambda

- Simplicity
- Lower cold start chance
- Higher cold start time
- Scaling & quotas
- Broad permissions
- Simpler CI/CD

Micro-functions



- Complexity
- Higher cold start chance
- Lower cold start time
- Independent scaling
- Granular permissions
- Multiple deployments

Event handler

Understand input structure

Create routing

Extract and transform input

Create response structure

Extract path and query parameters

Handle exceptions

```
1 def handler(event, contex):
2     try:
3         http_method = event["httpMethod"]
4         if http_method == "POST":
5             payload = json.loads(event["body"])
6             order_id = create_order(payload)
7             if not order_id:
8                 return {
9                     "statusCode": 500,
10                    "body": json.dumps({"message": "Could not create order"})
11                }
12            return {
13                "statusCode": 201,
14                "body": json.dumps({"order_id": order_id})
15            }
16        if http_method == "GET":
17            order_id = event["pathParameters"]["order_id"]
18            resp = get_order(order_id)
19            return {
20                "statusCode": 200,
21                "body": json.dumps(resp)
22            }
23    except:
24        return {
25            "statusCode": 500,
26            "body": json.dumps({"message": "Internal Server Error"})
27        }
28
```

Routing

```
1 from aws_lambda_powertools.event_handler.api_gateway import APIGatewayRestResolver
2 from aws_lambda_powertools.utilities.typing import LambdaContext
3 import requests
4
5 app = APIGatewayRestResolver()
6
7
8 @app.get("/todos/<todo_id>")
9 def get_todo_by_id(todo_id: int):
10     todo = requests.get(f"https://jsonplaceholder.typicode.com/todos/{todo_id}")
11     todo.raise_for_status()
12
13     return todo.json()
14
15
16 @app.route("/todos", methods=["POST", "PUT"])
17 def create_todo():
18     todo_data: dict = app.current_event.json_body # deserialize json str to dict
19     todo = requests.post("https://jsonplaceholder.typicode.com/todos", data=todo_data)
20     todo.raise_for_status()
21
22     return todo.json()
23
24
25 def lambda_handler(event: dict, context: LambdaContext):
26     return app.resolve(event, context)
27
```

Exception handling

```
1 @app.exception_handler(ValueError)
2 def handle_invalid_limit_qs(ex: ValueError): # receives exception raised
3     metadata = {"path": app.current_event.path, "query_strings": app.current_event.query_string_parameters}
4     logger.error(f"Malformed request: {ex}", extra=metadata)
5
6     return Response(
7         status_code=400,
8         content_type=content_types.TEXT_PLAIN,
9         body="Invalid request parameters.",
10    )
11
12
13 @app.get("/todos")
14 @tracer.capture_method
15 def get_todos():
16     # educational purpose only: we should receive a `ValueError`
17     # if a query string value for `limit` cannot be coerced to int
18     max_results: int = int(app.current_event.get_query_string_value(name="limit", default_value=0))
19
20     todos: requests.Response = requests.get(f"https://jsonplaceholder.typicode.com/todos?limit={max_results}")
21     todos.raise_for_status()
22
23     return {"todos": todos.json() }
```


Validation

```
1 from typing import Optional
2
3 import requests
4 from pydantic import BaseModel, Field
5
6 from aws_lambda_powertools.event_handler import APIGatewayRestResolver
7 from aws_lambda_powertools.utilities.typing import LambdaContext
8
9 app = APIGatewayRestResolver(enable_validation=True)
10
11
12 class Todo(BaseModel):
13     userId: int
14     id_: Optional[int] = Field(alias="id", default=None)
15     title: str
16     completed: bool
17
18
19 @app.get("/todos/<todo_id>")
20 def get_todo_by_id(todo_id: int) -> Todo:
21     todo = requests.get(f"https://jsonplaceholder.typicode.com/todos/{todo_id}")
22     todo.raise_for_status()
23
24     return todo.json()
25
26 def lambda_handler(event: dict, context: LambdaContext) -> dict:
27     return app.resolve(event, context)
```

OpenAPI spec

```
1 import requests
2
3 from aws_lambda_powertools.event_handler import APIGatewayRestResolver
4 from aws_lambda_powertools.utilities.typing import LambdaContext
5
6 app = APIGatewayRestResolver(enable_validation=True)
7
8
9 @app.get(
10     "/todos/<todo_id>",
11     summary="Retrieves a todo item",
12     description="Loads a todo item identified by the `todo_id`",
13     response_description="The todo object",
14     responses={
15         200: {"description": "Todo item found"},
16         404: {
17             "description": "Item not found",
18         },
19     },
20     tags=["Todos"],
21 )
22 def get_todo_title(todo_id: int) -> str:
23     todo = requests.get(f"https://jsonplaceholder.typicode.com/todos/{todo_id}")
24     todo.raise_for_status()
25
26     return todo.json()["title"]
```

And more...

OpenAPI spec

Fine-grained responses

Middleware

Custom serializer

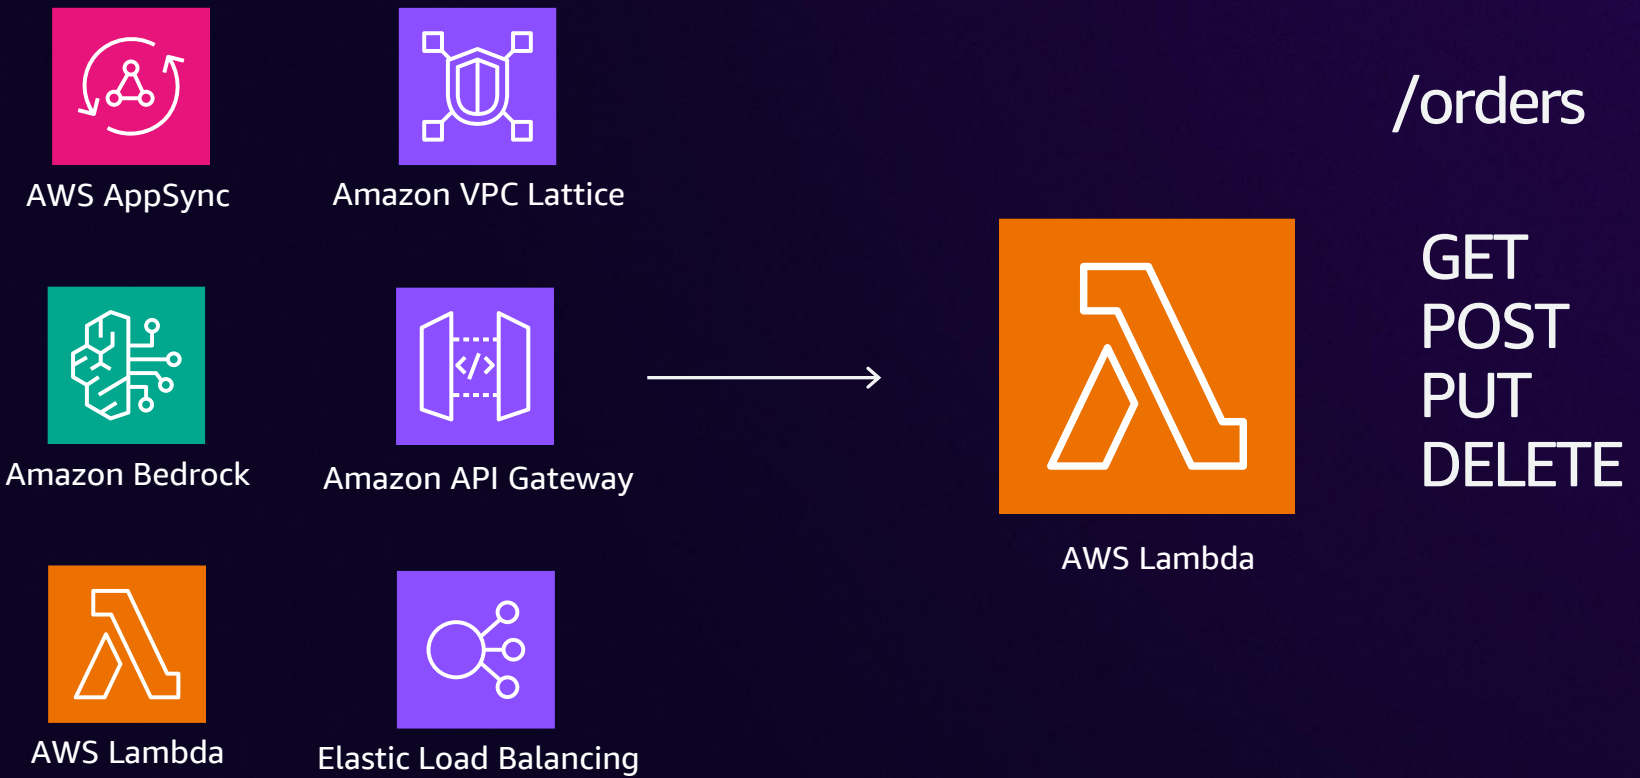
Swagger UI

Binary response

CORS

Compression

Event handler



Event handler

```
1 import requests
2 from requests import Response
3
4 from aws_lambda_powertools.event_handler import APIGatewayRestResolver
5 from aws_lambda_powertools.utilities.typing import LambdaContext
6
7 app = APIGatewayRestResolver()
8
9
10 @app.get("/todos")
11 @tracer.capture_method
12 def get_todos():
13     todos: Response = requests.get("...")
14     todos.raise_for_status()
15
16     return {"todos": todos.json()[0:10]}
17
18
19 def lambda_handler(event: dict, context: LambdaContext) -> dict:
20     return app.resolve(event, context)
```

```
1 import requests
2 from requests import Response
3
4 from aws_lambda_powertools.event_handler import ALBResolver
5 from aws_lambda_powertools.utilities.typing import LambdaContext
6
7 app = ALBResolver()
8
9
10 @app.get("/todos")
11 @tracer.capture_method
12 def get_todos():
13     todos: Response = requests.get("...")
14     todos.raise_for_status()
15
16     return {"todos": todos.json()[0:10]}
17
18
19 def lambda_handler(event: dict, context: LambdaContext) -> dict:
20     return app.resolve(event, context)
```


Idempotency



What is idempotency?

idem same
potent having power

idempotent (*ī-dem pə-tən-sē*)
adjective

Same effect regardless of retries

Safely retry without side effects

Stable outcomes

Consistency

Resilience

Why do I need idempotency?

"The world is **asynchronous**."

Dr. Werner Vogels
VP and CTO at Amazon.com

At-least-once delivery

Lambda retries

SDK retries

Transient failures

Eventual consistency

Auto-scaling operations


Idempotency



Idempotency



Idempotency



Timeout	Concurrency	Consistency	Expiration
Validation	Idempotency key		Serialization

Scenario 1: First invocation



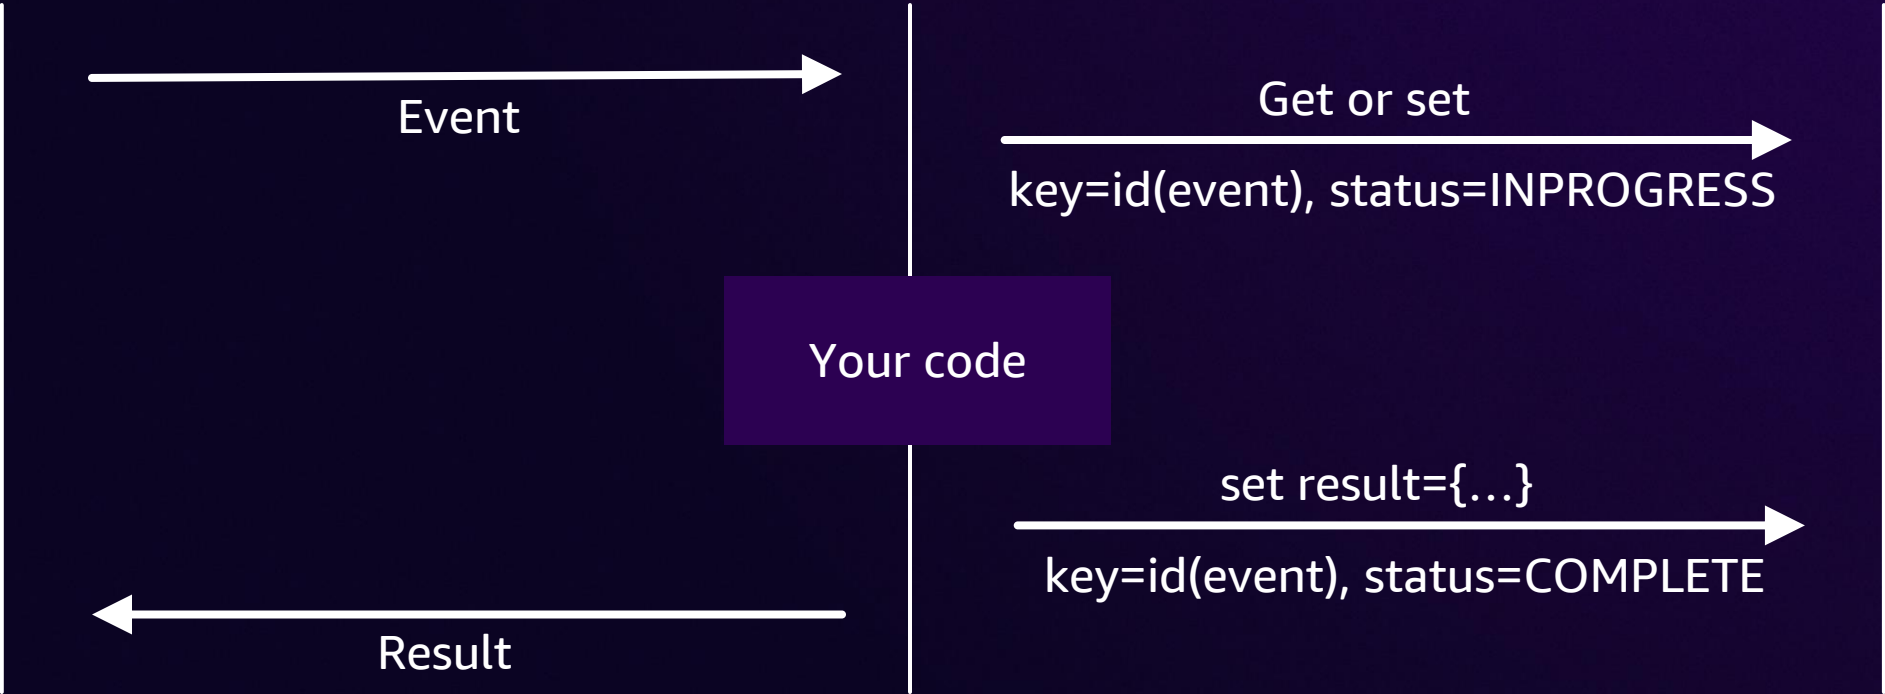
Caller



Your Lambda



Storage



Scenario 2: Second invocation



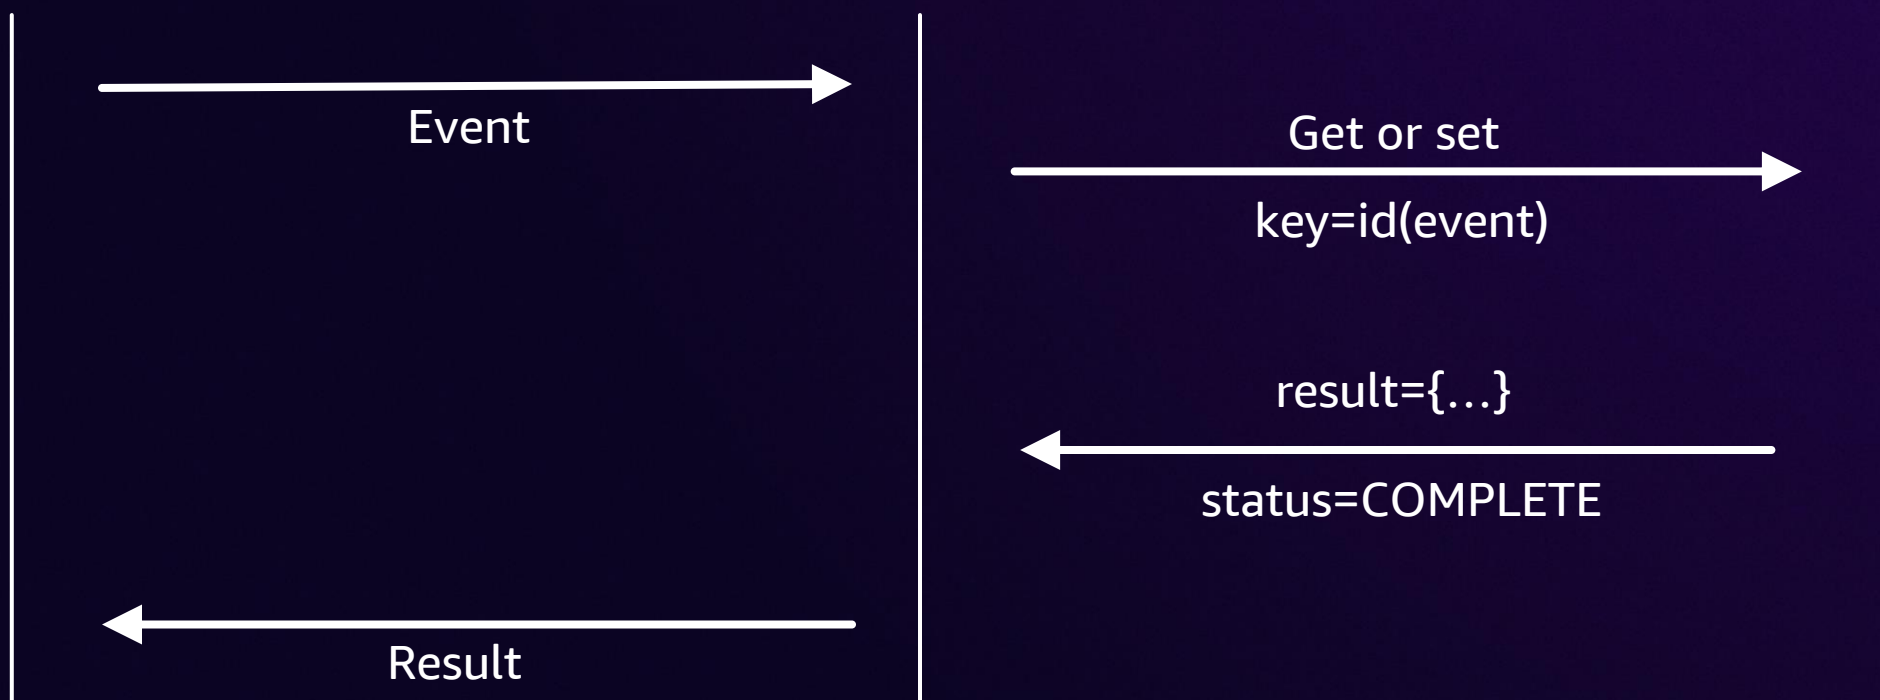
Caller



Your Lambda



Storage



Scenario 3: Concurrent invocations




Caller



Your Lambda



Storage



Scenario 4: Handling exceptions



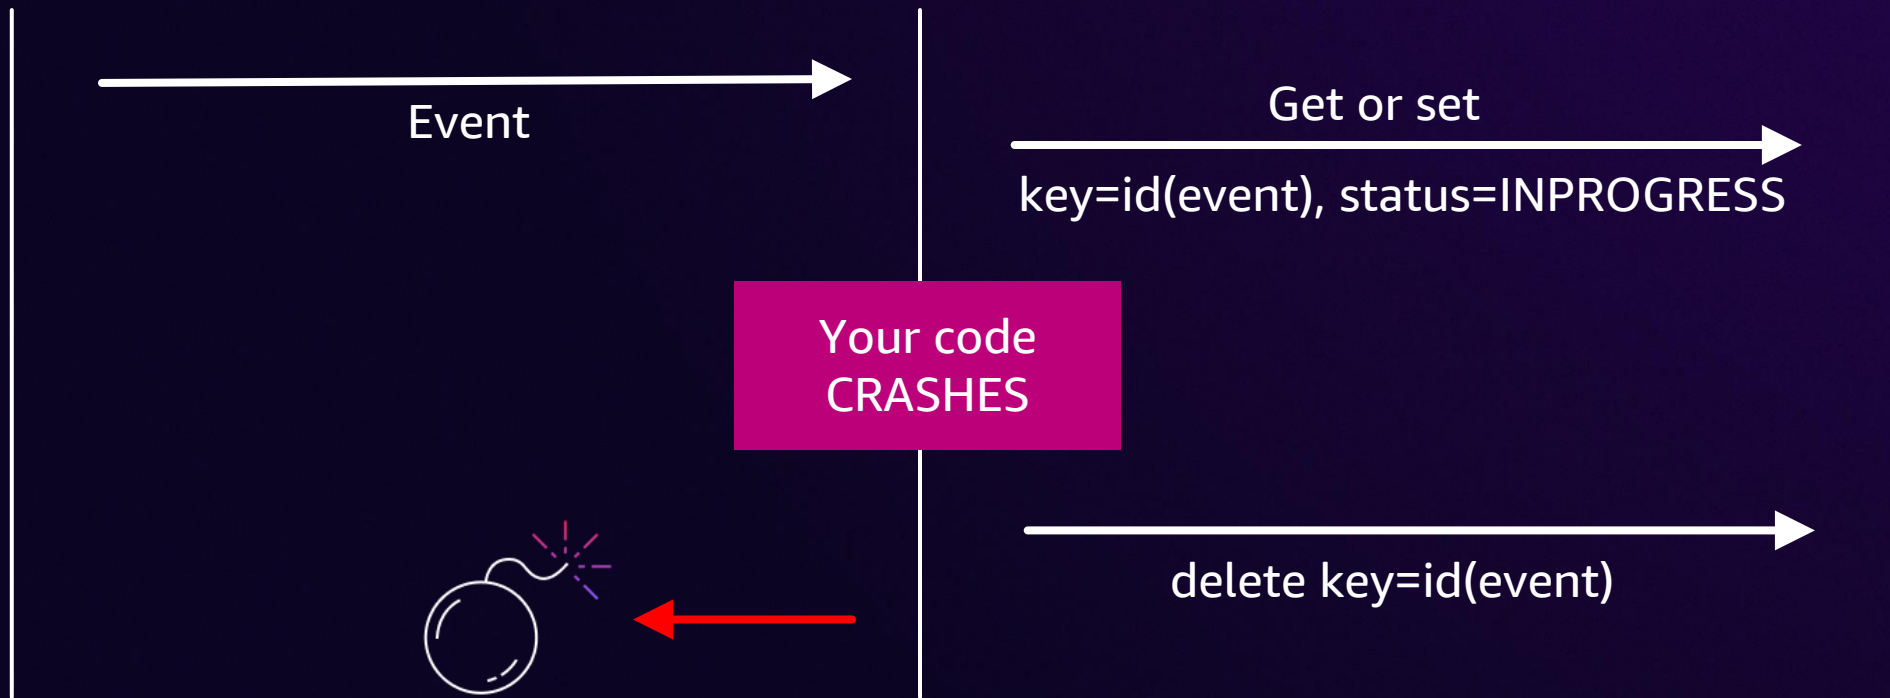
Caller



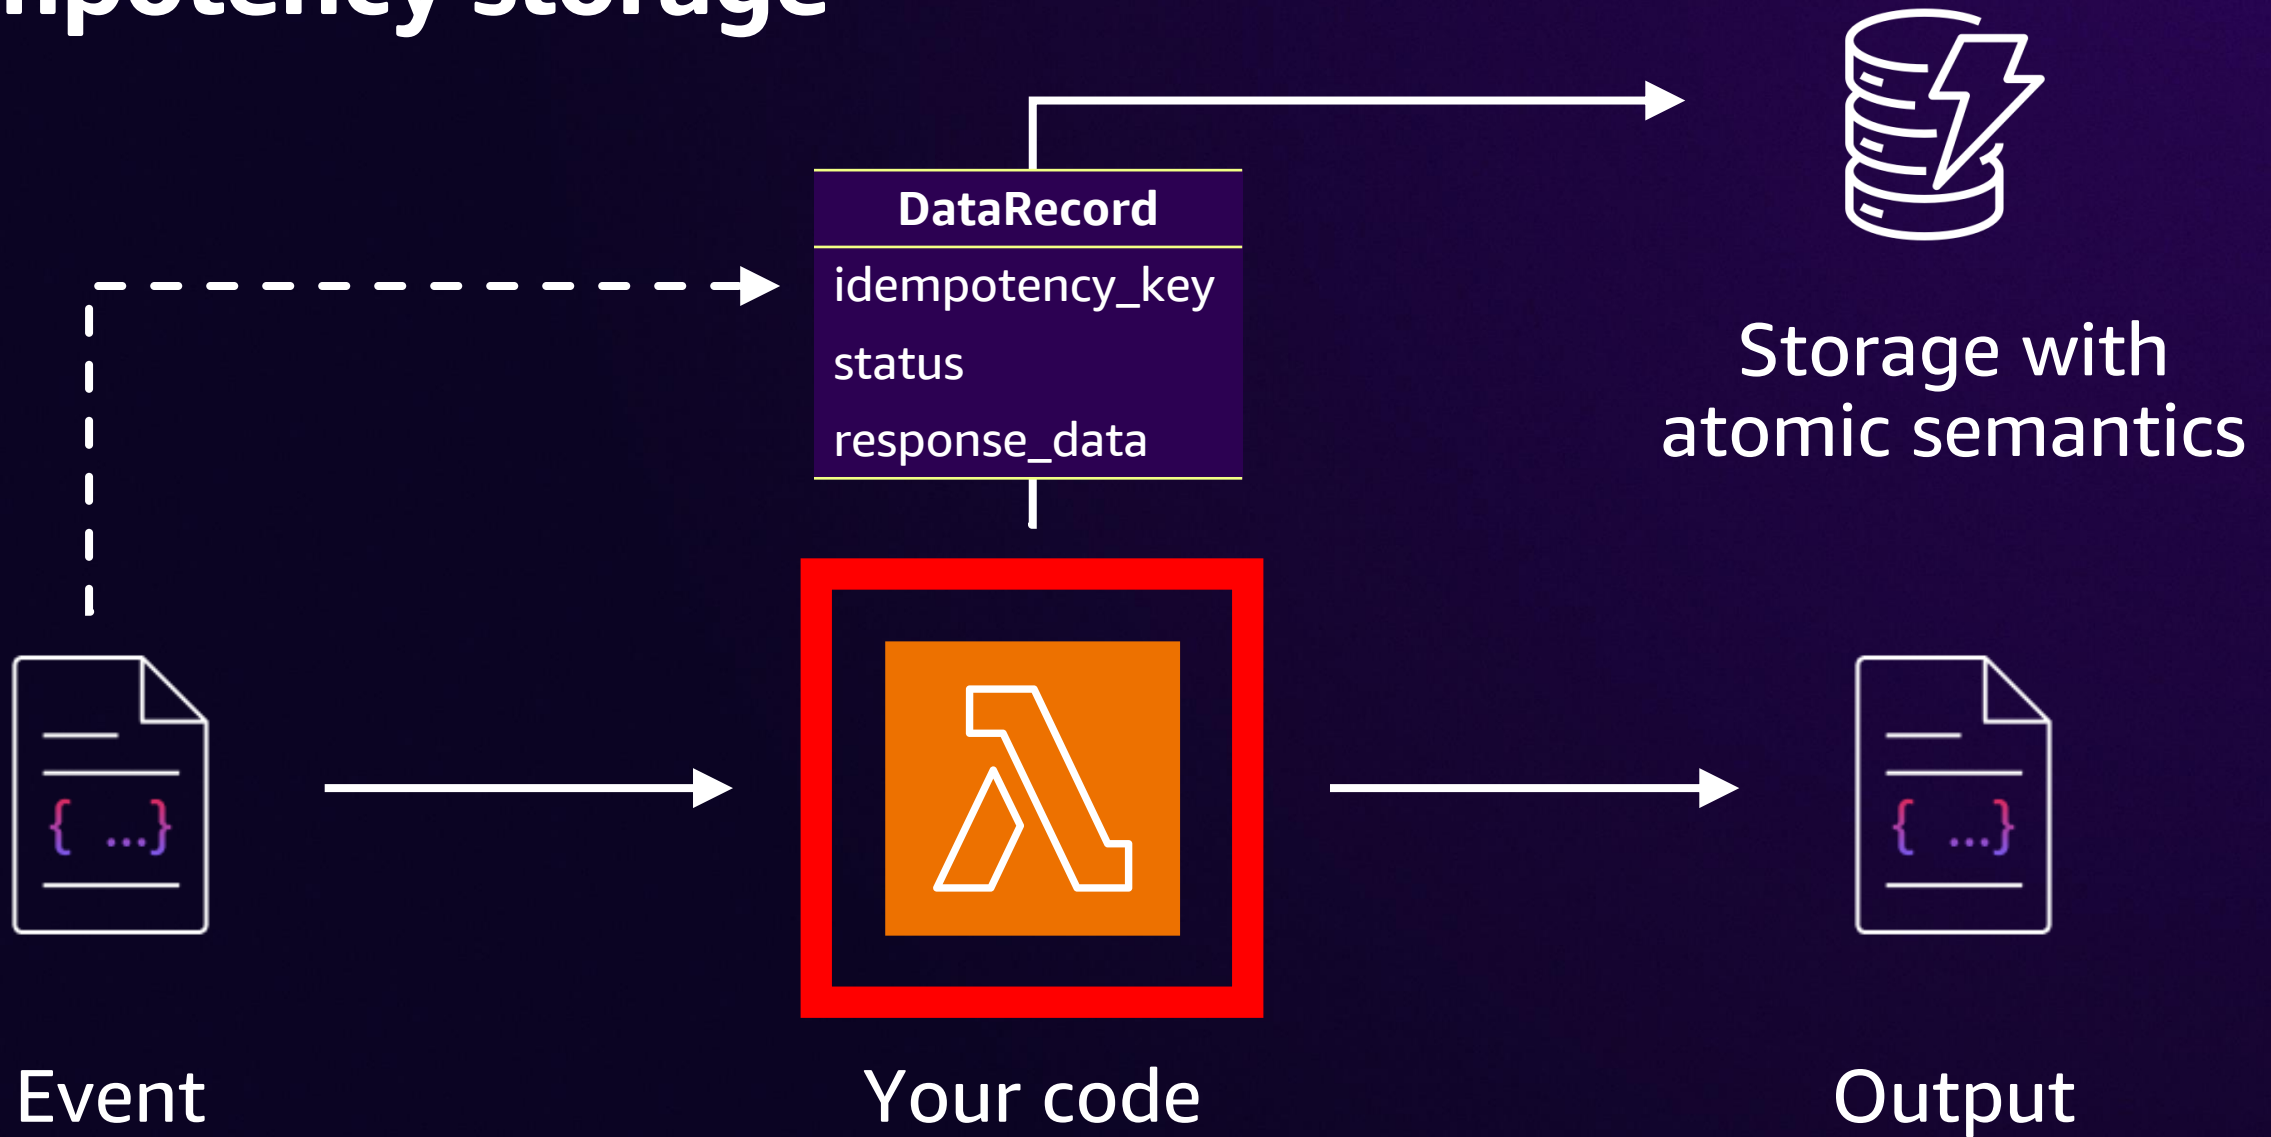
Your Lambda



Storage



Idempotency storage



Idempotency

```
1 from aws_lambda_powertools.utilities import LambdaContext
2 from aws_lambda_powertools.utilities.idempotency import (
3     DynamoDBPersistenceLayer,
4     idempotent,
5 )
6
7 ddbPersistenStore = DynamoDBPersistenceLayer(table_name="IdempotencyTable")
8
9 @idempotent(persistent_store=ddbPersistenStore)
10 def lambda_handler(event: dict, context: LambdaContext):
11     try:
12         payment = create_subscription_payment(event)
13         return {
14             "payment_id": payment.id,
15             "statusCode": 200,
16             "body": payment
17         }
18     except Exception as e:
19         raise ValueError(f"Failed to create payment: {e}")
20
21 def create_subscription_payment(event: dict):
22     # Create payment
23     pass
```

Context awareness

```
1 import os
2 from dataclasses import dataclass
3
4 from aws_lambda_powertools.utilities.idempotency import (
5     DynamoDBPersistenceLayer,
6     IdempotencyConfig,
7     idempotent_function,
8 )
9 from aws_lambda_powertools.utilities.typing import LambdaContext
10
11 table = os.getenv("IDEMPOTENCY_TABLE", "")
12 dynamodb = DynamoDBPersistenceLayer(table_name=table)
13 config = IdempotencyConfig(event_key_jmespath="order_id")
14
15
16 @idempotent_function(data_keyword_argument="order", config=config, persistence_store=dynamodb)
17 def process_order(order: Order):
18     return f"processed order {order.order_id}"
19
20
21 def lambda_handler(event: dict, context: LambdaContext):
22     config.register_lambda_context(context)
23
24     order_item = OrderItem(sku="fake", description="sample")
25     order = Order(item=order_item, order_id=1)
26
27     process_order(order=order)
```

Fine-tune settings

```
1 import os
2
3 from aws_lambda_powertools.utilities.idempotency import (
4     DynamoDBPersistenceLayer,
5     IdempotencyConfig,
6     idempotent,
7 )
8 from aws_lambda_powertools.utilities.typing import LambdaContext
9
10 table = os.getenv("IDEMPOTENCY_TABLE", "")
11 persistence_layer = DynamoDBPersistenceLayer(table_name=table)
12 config = IdempotencyConfig(
13     event_key_jmespath="powertools_json(body)",
14     # by default, it holds 256 items in a Least-Recently-Used (LRU) manner
15     use_local_cache=True,
16 )
17
18
19 @idempotent(config=config, persistence_store=persistence_layer)
20 def lambda_handler(event, context: LambdaContext):
21     return event
```


Swap persistence layer

```
1 import os
2
3 from aws_lambda_powertools.utilities.idempotency import (
4     idempotent,
5 )
6 from aws_lambda_powertools.utilities.typing import LambdaContext
7 from aws_lambda_powertools.utilities.idempotency.persistence.redis import (
8     RedisCachePersistenceLayer,
9 )
10
11 valkey_endpoint = os.getenv("VALKEY_ENDPOINT")
12 persistence_layer = RedisCachePersistenceLayer(host=valkey_endpoint, port=6379)
13
14
15 @idempotent(persistence_store=persistence_layer)
16 def lambda_handler(event, context: LambdaContext):
17     return event
```


Batch processing



Batch processing




Amazon Simple Queue Service (Amazon SQS)




Amazon Kinesis



Amazon DynamoDB

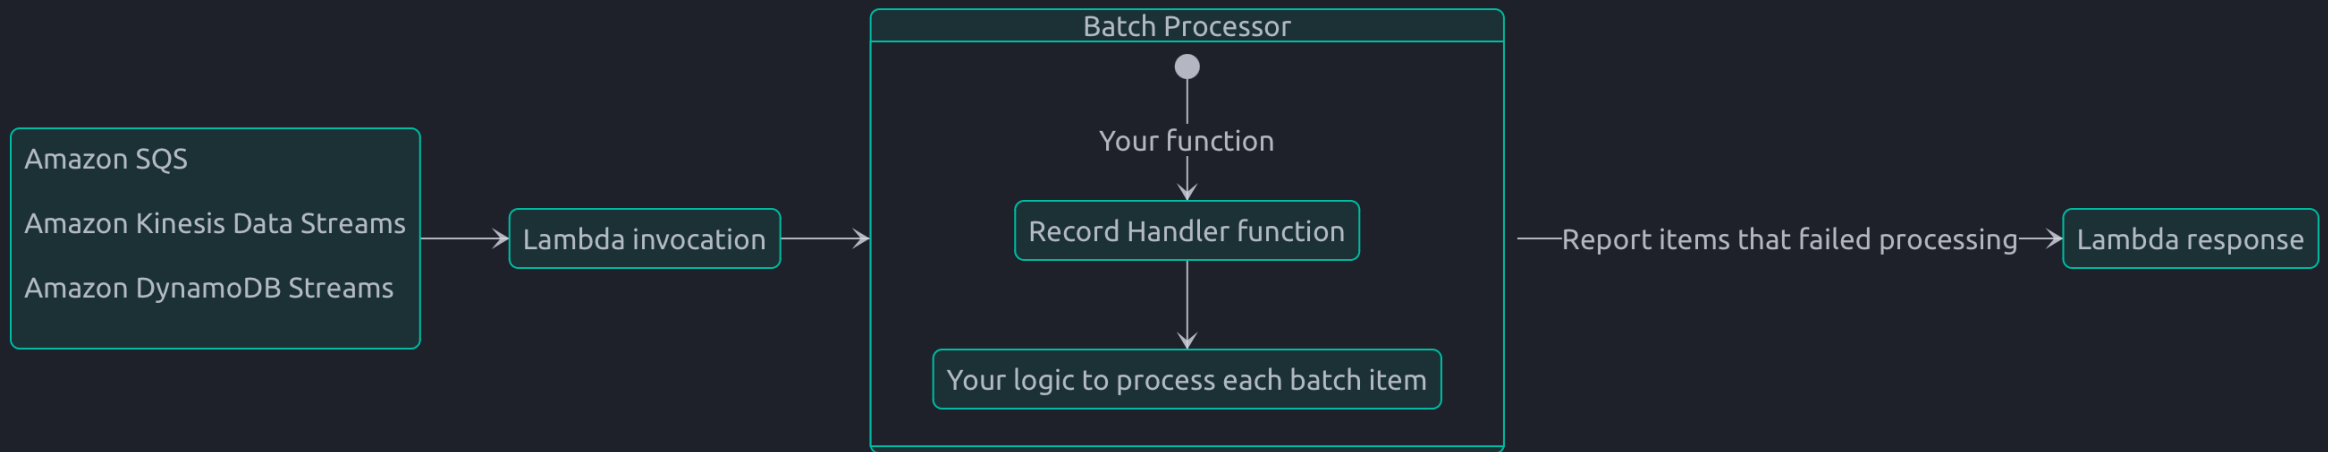


AWS Lambda



```
1 {  
2   "batchItemFailures": [  
3     {  
4       "itemIdentifier": "244fc6b4-87a3-44ab-83d2-361172410c3a"  
5     }  
6   ]  
7 }
```

Batch processing



Batch processing

```
1 from aws_lambda_powertools import Logger, Tracer
2 from aws_lambda_powertools.utilities.batch import (
3     BatchProcessor,
4     EventType,
5     process_partial_response,
6 )
7 from aws_lambda_powertools.utilities.data_classes.sqs_event import SQSRecord
8 from aws_lambda_powertools.utilities.typing import LambdaContext
9
10 processor = BatchProcessor(event_type=EventType.SQS)
11 tracer = Tracer()
12 logger = Logger()
13
14
15 @tracer.capture_method
16 def record_handler(record: SQSRecord):
17     payload: str = record.json_body # if json string data, otherwise record.body for str
18     logger.info(payload)
19
20
21 @logger.inject_lambda_context
22 @tracer.capture_lambda_handler
23 def lambda_handler(event, context: LambdaContext):
24     return process_partial_response(
25         event=event,
26         record_handler=record_handler,
27         processor=processor,
28         context=context,
29     )
```

Batch processing

```
1 from aws_lambda_powertools import Logger, Tracer
2 from aws_lambda_powertools.utilities.batch import (
3     BatchProcessor,
4     EventType,
5     process_partial_response,
6 )
7 from aws_lambda_powertools.utilities.data_classes.kinesis_stream_event import (
8     KinesisStreamRecord,
9 )
10 from aws_lambda_powertools.utilities.typing import LambdaContext
11
12 processor = BatchProcessor(event_type=EventType.KinesisDataStreams)
13 tracer = Tracer()
14 logger = Logger()
15
16
17 @tracer.capture_method
18 def record_handler(record: KinesisStreamRecord):
19     logger.info(record.kinesis.data_as_text)
20     payload: dict = record.kinesis.data_as_json()
21     logger.info(payload)
22
23 @logger.inject_lambda_context
24 @tracer.capture_lambda_handler
25 def lambda_handler(event, context: LambdaContext):
26     return process_partial_response(
27         event=event,
28         record_handler=record_handler,
29         processor=processor,
30         context=context
31     )
```


Roadmap



© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Roadmap



Powertools for
AWS Lambda
(Python)



Powertools for
AWS Lambda
(TypeScript)



Powertools for
AWS Lambda
(.NET)

Thank you!

Andrea Amorosi

X @dreamorosi



Please complete the session
survey in the mobile app

