
AWS
re:Invent
NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

ARC301-R

Building a monitoring strategy

Chris Kozlowski (he/him)

Specialist Technical Account Manager, Enterprise Support
AWS

Paul Moran (he/him)

Principal Technical Account Manager, Enterprise Support
AWS

What we are going to talk about today

- Well-Architected
- Why you need a monitoring strategy
- What goes into a monitoring strategy
- Building a monitoring strategy
- Conclusion and next steps

Case study

Links

- Brief: https://bit.ly/arc301brief
- Workbook: https://bit.ly/arc301workbook

Well-Architected

Well-Architected history

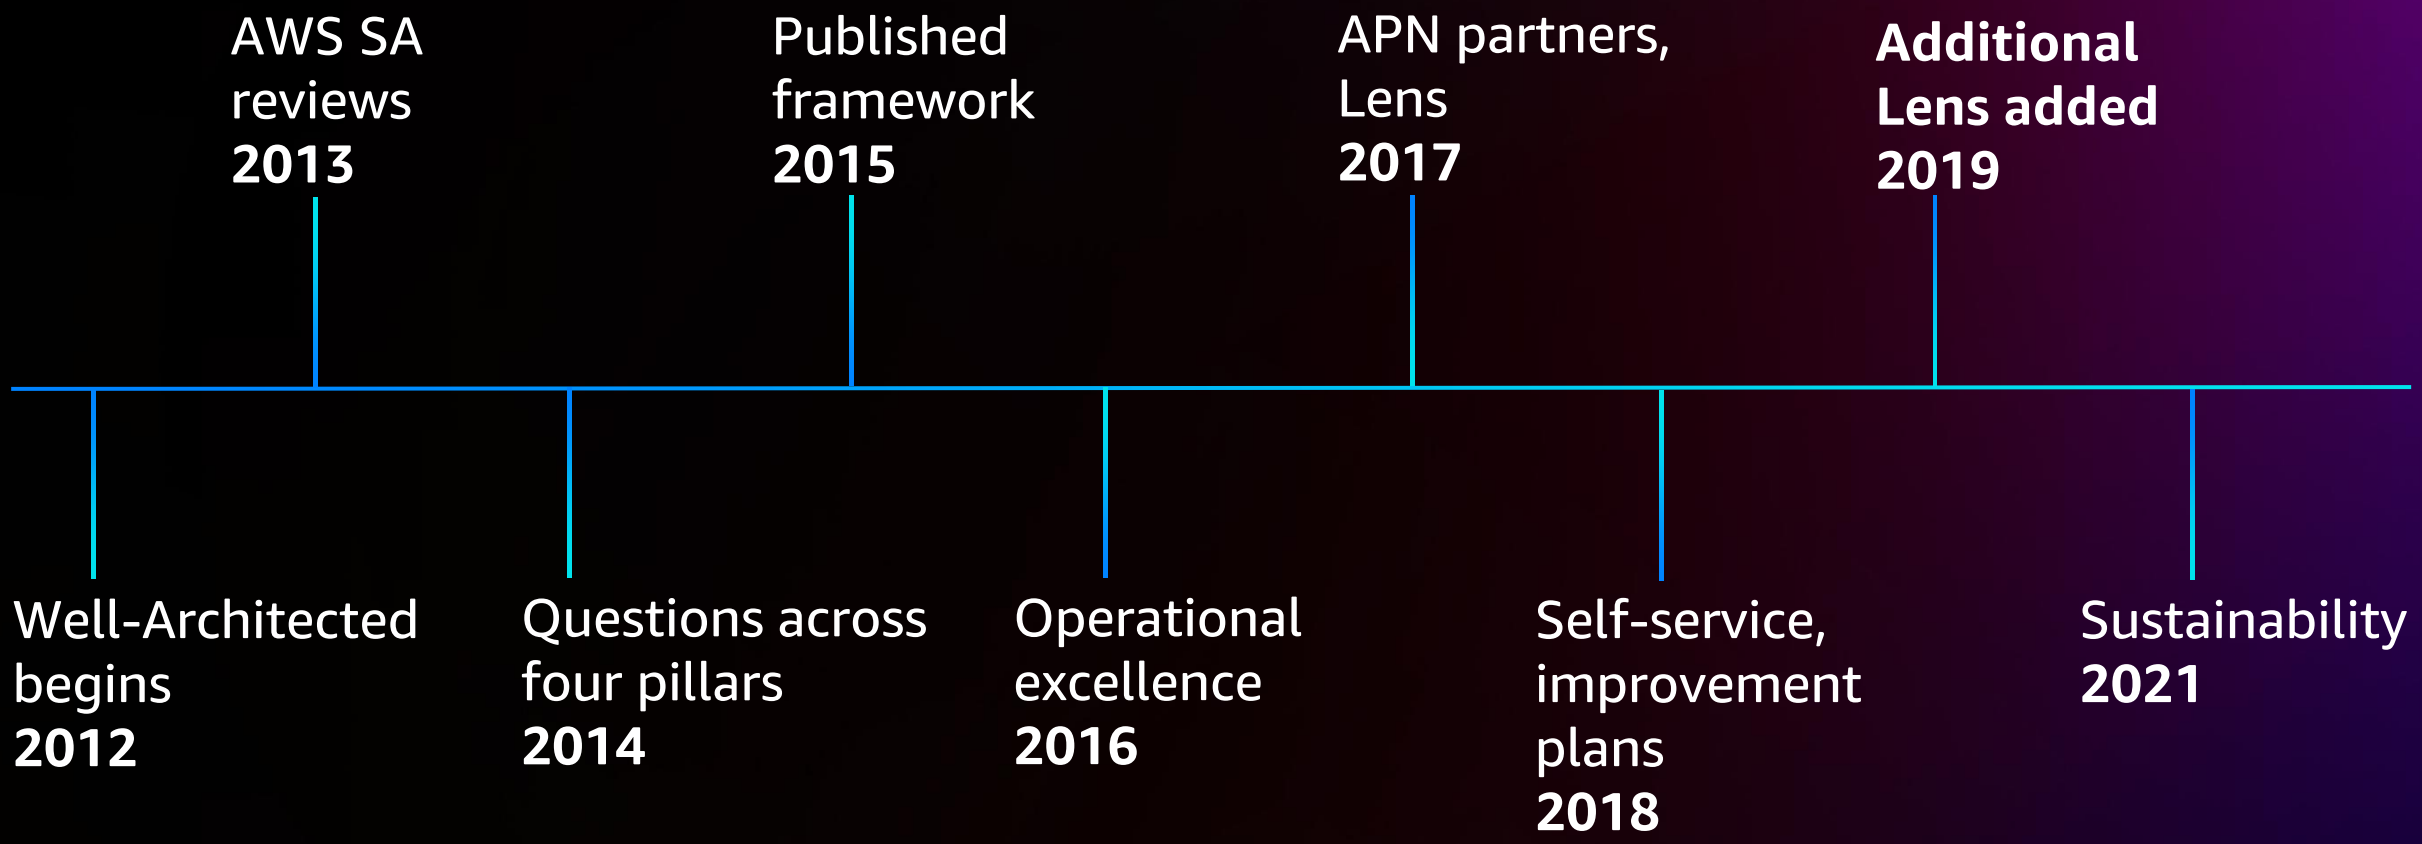

Well-Architected pillars

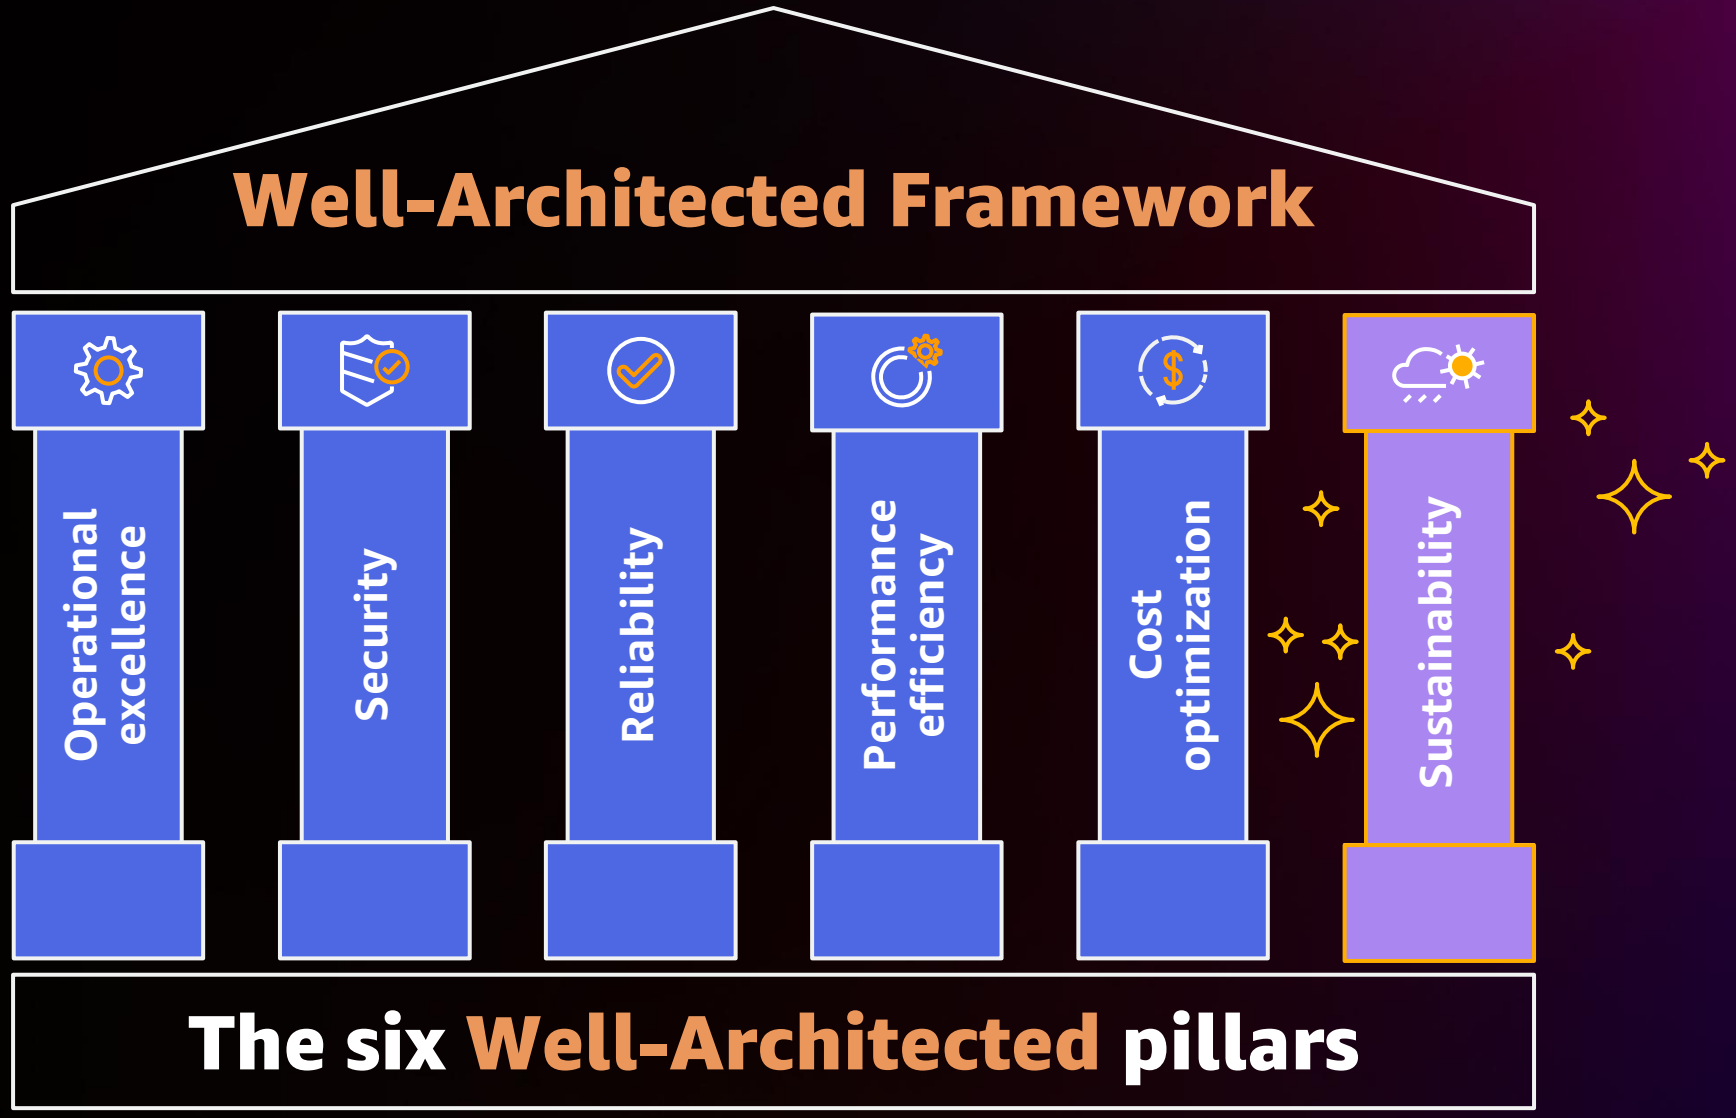

Why Well-Architected?

Build and deploy faster

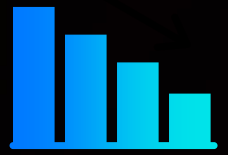

Lower or mitigate risks

Make informed decisions

Learn AWS best practices

Why do you need a monitoring strategy?

What are trying to do and why?

- Gain insights for business and operations
- Achieve situational awareness
- Enable proactive courses of action
- Provide timely and effective responses
- Support the achievement of business outcomes

Alarm fatigue

Single pane of glass

Observability

E-bonding

Buzzword bingo

“No ops”

Signal to noise ratio

Operational intelligence

Composable monitoring

Does monitoring need your attention?

Have you ever been asked . . .

how you are supposed to monitor in the cloud?

Does monitoring need your attention?

Have you ever been asked . . .

how to get visibility into serverless workloads?

Does monitoring need your attention?

Have you ever . . .

struggled with not knowing why a workload has issues?

Does monitoring need your attention?

Have you ever been asked . . .

which monitoring tool is better?

Does monitoring need your attention?

Have you ever . . .

seen the same issue happen in production repeatedly?

Why do you need a monitoring strategy?

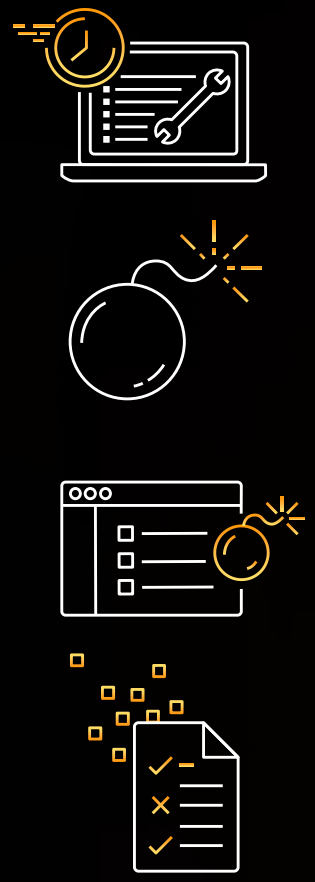

- Your strategy informs your plans
- Failing to plan is planning to fail
- You cannot manage what you do not measure
- A monitoring plan defines what is measured

What goes into a monitoring strategy?

What goes into a monitoring strategy?

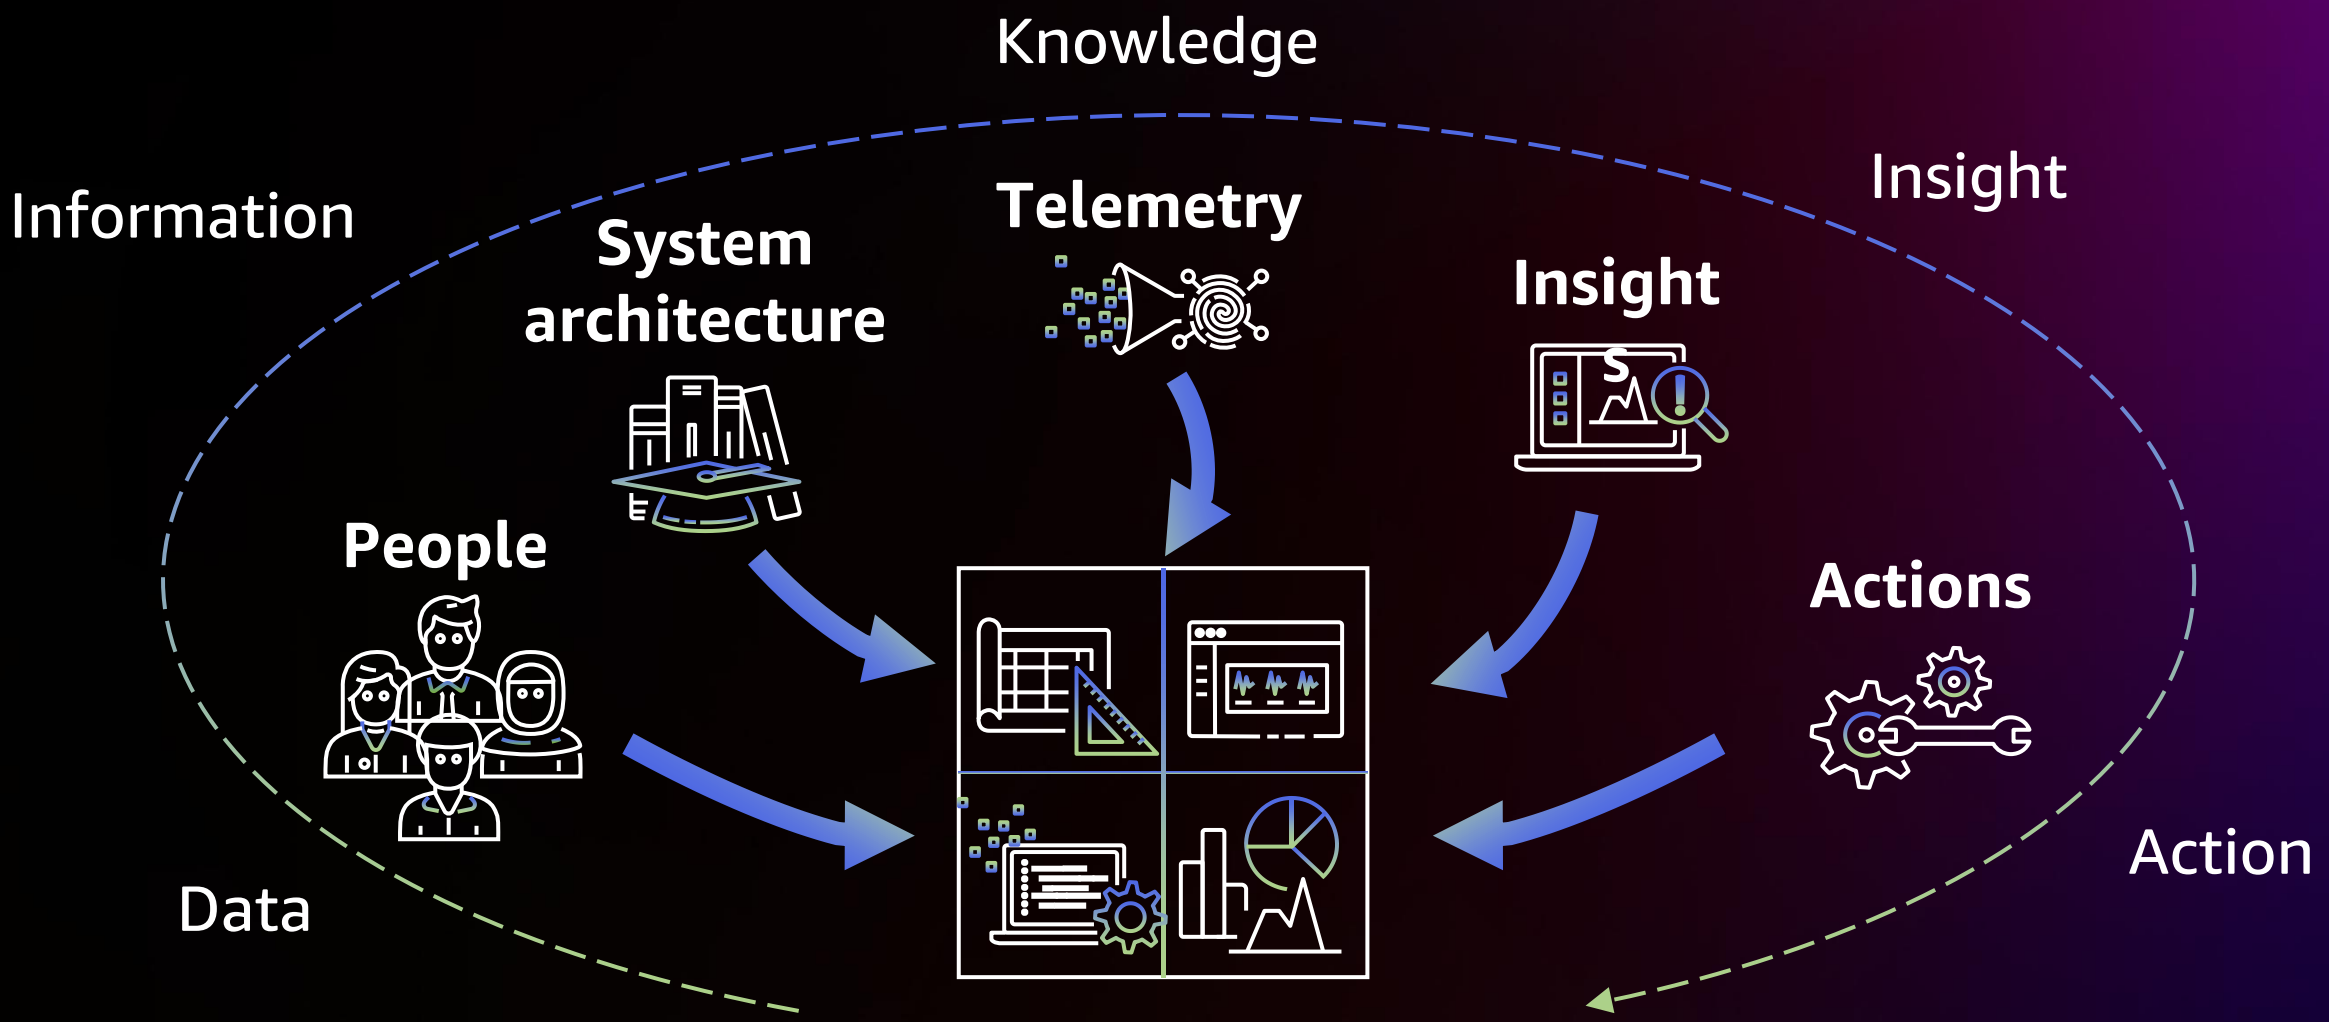

Categories of insights

FCAPS

Who has heard of FCAPS?

- Fault management
- Configuration management
- Accounting management
- Performance management
- Security management

Categories of insights

Faults: monitoring to avert or respond to incidents

Categories of insights

Configuration: monitoring and tracking configurations and changes

Categories of insights

Accounting: monitoring utilization and enabling attribution

Categories of insights

Performance: monitoring for constrained components and the impact of changes

Categories of insights

Security: monitoring access, security controls, and identifying inappropriate or malicious activity

We may detect failed components,
understand their configuration,
attribute the cost of their usage,
understand their individual performance,
and secure and monitor access to them . . .

This is not enough

Categories of insights – “What good looks like”

Understanding the achievement of business outcomes

Categories of insights

- Faults
- Security
- Configuration
- Performance
- Accounting
- **Business Outcomes**

Advanced categories of insights

Understanding user behavior

Advanced categories of insights

Understanding
workload behavior

Sources of insights

Sources of insights

Logs: Immutable records of discrete timestamped events

Sources of insights

Metrics: Measurements performed at specific time intervals

Sources of insights

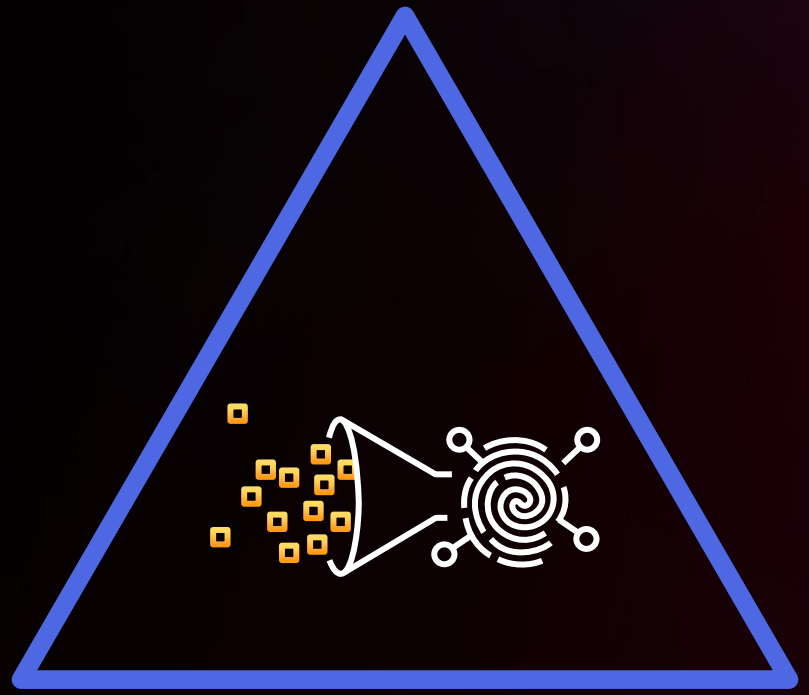

Traces: Sets of data tracking individually identified requests end-to-end through your application and services

Sources of insights – “What good looks like”

Observability: The ability to successfully use collected telemetry to discover and learn what is going on inside your workload

Observability

- Logs
- Metrics
- Traces

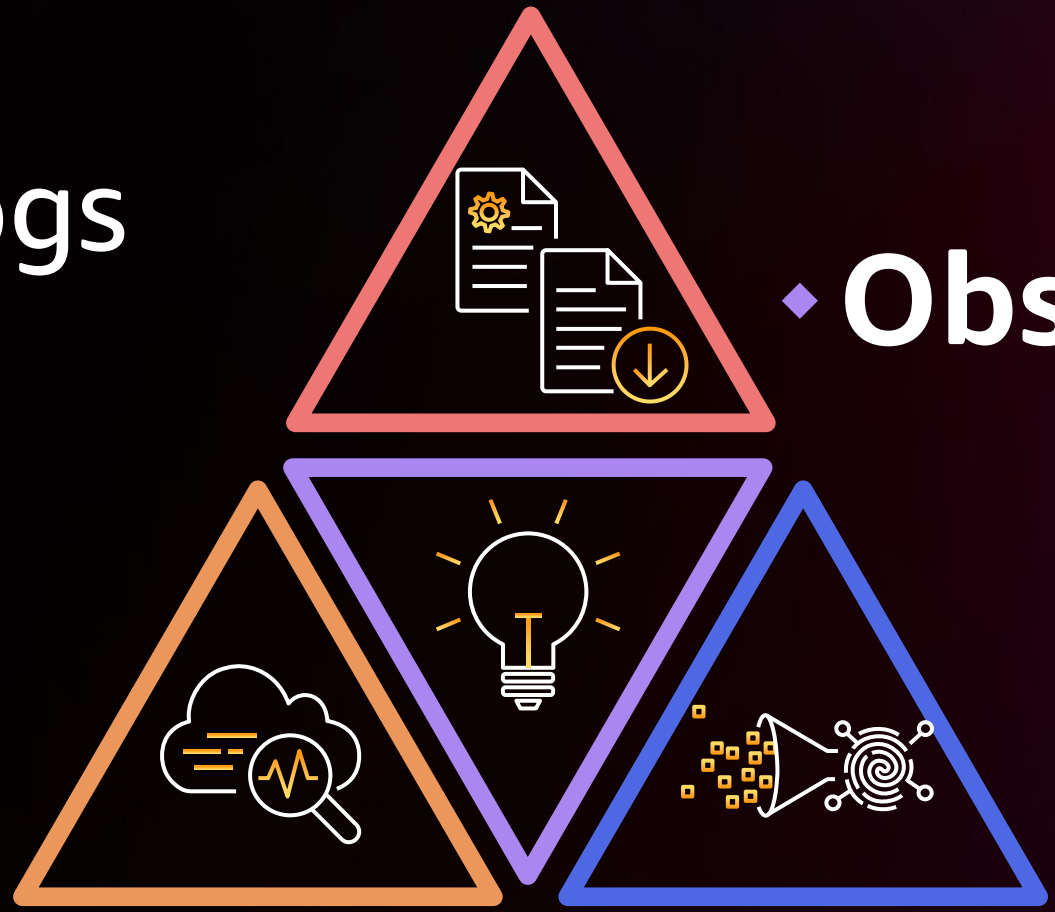

- **Observability**

AWS Observability

Services and features

Foundation

CloudWatch logs

CloudWatch metrics

X-Ray traces

CloudWatch alarms

Observability and monitoring

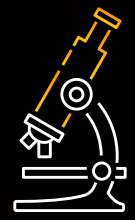

Observability

A measure of how well you can understand what is happening inside your workload

Monitoring

Analyzing telemetry to understand the health of your workload, identify risks to the achievement of business outcomes, enable responses, and inform decision making

Building a monitoring strategy

Monitoring strategy workflow overview

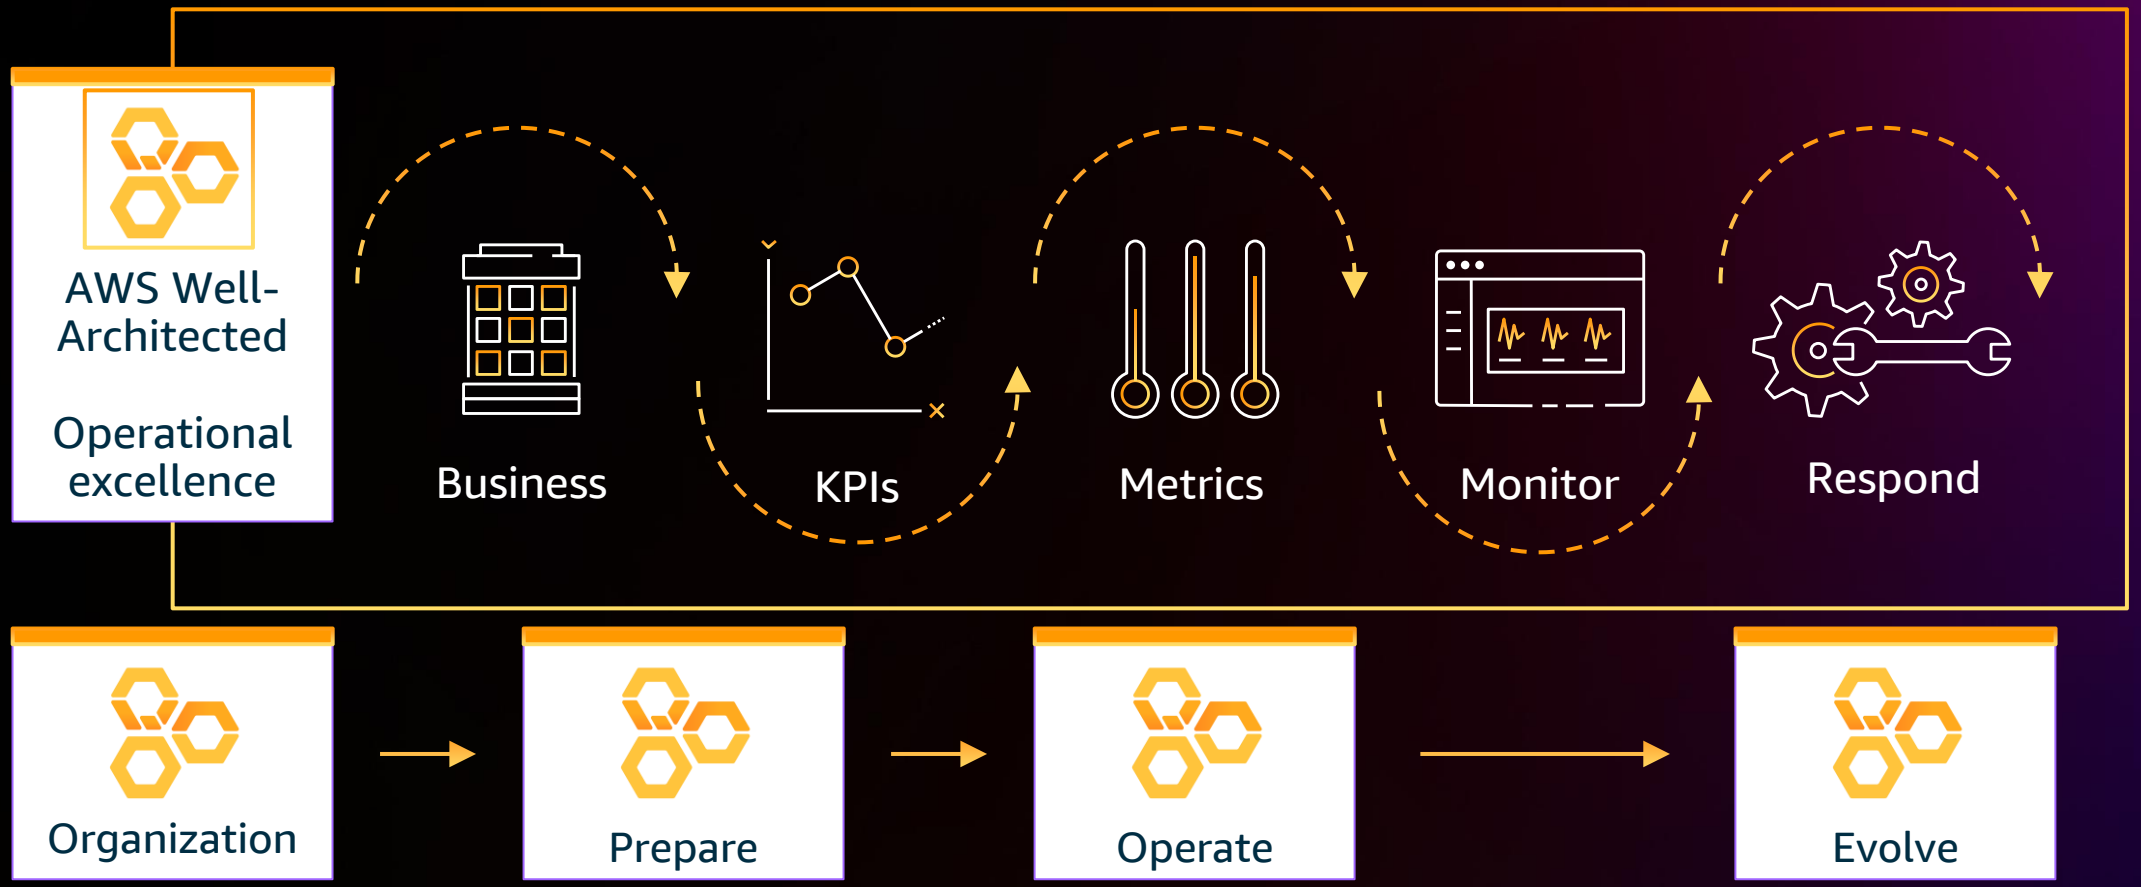

Start with people and understand their needs

PERSONAS AND THEIR REQUIREMENTS

Who will consume your insights?

Internal customers

- Business owners
- Product owners
- Operations
- Developers
- Security
- Compliance

Who will consume your insights?

External customers

- End-users
- Partners
- Vendors
- Suppliers
- Providers

Who will consume your insights?

Mandatory reporting

- Government regulators
- Accreditation bodies
- Industry organizations

Personas and requirements

Role	Areas of interest	Required insights
Chief financial officer	Accounting	Profit, margin, and cost; revenue versus spend
Accounting	Accounting	Status of payroll systems
Chief compliance officer	Security	Status of compliance/noncompliance requiring mandatory reporting
Product owner – retail site	Performance, availability	Health of product workloads, quality of customer experience
Retail site customer	Performance, availability	Availability and status of requests and transactions
and many more . . .		

Personas and requirements

Role	Areas of interest	Required insights
Information security	Security	Status of security controls, events, and metrics
Operations	Faults, performance	Health of workloads, components, services, and external dependencies
Developer	Configuration	Integrity of product workloads
Partners	Availability	Availability and status of requests and transactions
Government regulators	Security, accounting	Mandatory noncompliance reports
and many more . . .		

Exercise #1: Who are my customers?

Case study:
Brief: https://bit.ly/arc301brief
Workbook: https://bit.ly/arc301workbook

- Who are my customers? What are their roles?
- What are their areas of interest?
  There may be more than one!
- Which insights do they require?

Understand what the business is trying to achieve

BUSINESS OUTCOMES, GOALS, AND KPIS

Key performance indicators (KPIs)

KPIs are a metric that indicates how the business is doing

- Set by senior leaders across the organization
- They reflect whether the team or organization is achieving its business goals
- They are subject to interpretation; satisfactory values can change over time

The SMART test should apply – KPIs should be specific, measurable, achievable, relevant, and timely

Business outcomes, goals, and key performance indicators

Business outcomes	Goals	KPIs
Quality user experience	User is able to access the site	Site availability
Quality user experience	User receives populated page quickly enough that they do not abandon their transaction	Site responsiveness
Quality user experience	Users reach a detailed product page with minimal clicks and data entry	Steps required to navigate to a detailed product page
Quality user experience	Users complete purchases with minimal clicks and data entry	Steps required to complete purchase
Quality user experience	User is able to track the fulfillment status and delivery of their purchases	Availability of purchase tracking data
Quality user experience	Fulfillment and delivery information is current and accurate within committed delivery time	Fulfillment and delivery tracking data is available
and many more . . .		

Business outcomes, goals, and key performance indicators

Business outcomes	Goals	KPI
Timely regulatory reporting	Noncompliance identified quickly enough that reporting noncompliance is possible	Time to identify noncompliance
Timely regulatory reporting	Noncompliance reported quickly enough that legal penalties are not incurred	Time to report noncompliance
Timely vulnerability management	Vulnerable components should be identified within minimum time after vulnerability announcement	Time to identify vulnerable components
Timely vulnerability management	Vulnerabilities should be mitigated or remediated within minimum time after vulnerability is identified	Time to mitigate or remediate vulnerabilities

and many more . . .

Exercise #2 – What are my KPIs?

Refer to your case study:

- What are the business outcomes our leaders desire?
  Refer to their roles and insights – tie KPIs to roles!
- What are the goals required to meet those outcomes?
- How can we measure if we're achieving those goals?

Understand what telemetry is available and what you need

KNOWLEDGE OF THE SYSTEM AND SOURCES OF INSIGHT

The drive toward achieving business insights

Business insights!

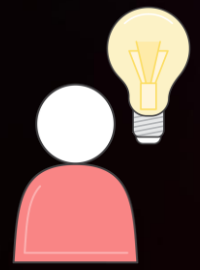

Customer sentiment, SLAs

Business-level alerts

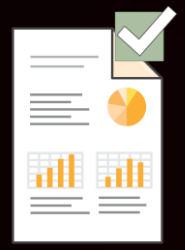

Webpage response time, job run length

System-level alerts

CPU wait %, disk queue depth

Telemetry – Sources of insights

- Logs - Immutable records of discrete timestamped events

- Metrics - Measurements performed at specific time intervals

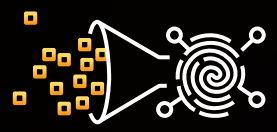

- Traces - Sets of data tracking individual identified requests end-to-end through your application and services



Knowledge of the system and sources of insight

Workload	Component	Subcomponent	Insights	Telemetry	Source
Retail site	Web server	Operating system	Faults	Logs	/var/log/message
Retail site	Web server	Operating system	Faults	Logs	/var/log/boot.log
Retail site	Web server	CPU	Performance	Standard metrics	CPU utilization
Retail site	Web server	Memory	Performance	Standard metrics	Memory utilization
Partner site	Shipping system	Disk I/O	Performance	Standard metrics	Disk I/O utilization
Retail site	Application Server	Warnings	Performance	Logs	/opt/retail-site-logs/ui/application.log
Retail site	Application Server	User activity	User behavior	Logs	/opt/retail-site-logs/ui/application.log
Retail site	Application Server	Production logs	Accounting	Logs	/opt/retail-site-logs/ui/production.log
Retail site	Application Server	Network activity	Security	VPC flow logs	VPC flow logs

and many more . . .

Exercise #3 - Telemetry

Refer to your case study:

- What are the sources of telemetry we have? What insights can they provide?
- Can we identify a source of insight for all of our roles and their associated KPIs?
- Do we have blind spots in our system for which we have no sources of insight?
  What might we need to add?

What does good look like?

METRICS AND THRESHOLDS

Determine thresholds for

Investigation

- Successful but constrained resources or unexpected behaviors, for example, an abrupt decrease in performance
- Identify contributing factors and take action as necessary

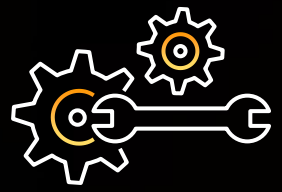

Intervention

- There is an impact on the achievement of business outcomes
- Mitigate or remediate and restore normal operations

Improvement

- Trend towards exceeding capacity, consistent high utilization, or decrease in performance attributed to changes to application or configuration of components
- Identify, plan, and prioritize improvements

Setting the right thresholds

S.M.A.R.T.

- Specific
- Measurable
- Achievable
- Relevant
- Time-bound

Metrics and thresholds

Workload	Component	Subcomponent	Metric	Test	Thresholds	Source
Retail site	Web server	Web Service	Page latency	Average page load time	>=5ms	CloudWatch log metrics /var/log/apache
Retail site	Web server	Web Service	Page errors	5xx HTML errors	>=1% error rate	CloudWatch log metrics /var/log/apache
Retail site	Web server	Web Client	Click count	Number of clicks to loaded product page	>=4 clicks count	X-Ray traces on transaction ID
Retail site	Web server	Web Client	Data entry	Count of data entry actions	>=2 actions	X-Ray traces on transaction ID
Partner site	Shipping system	Shipping API	Availability	API errors	>=5 errors in a 5 minute period	Application custom CloudWatch metric
Retail site	Amazon RDS	Amazon RDS	IOPS	Consumed IOPS	>95% provisions IOPS p99	CloudWatch metrics WriteIOPS
and many more . . .						

Exercise #4 – Thresholds

Refer to your case study:

- What does “good look like” as a S.M.A.R.T. objective?
- Map these objectives to your information sources in the previous exercise
- Are there any sources you cannot define what good looks like?

Taking action in response to events

ALERTS AND ACTIONS

How do you prioritize?

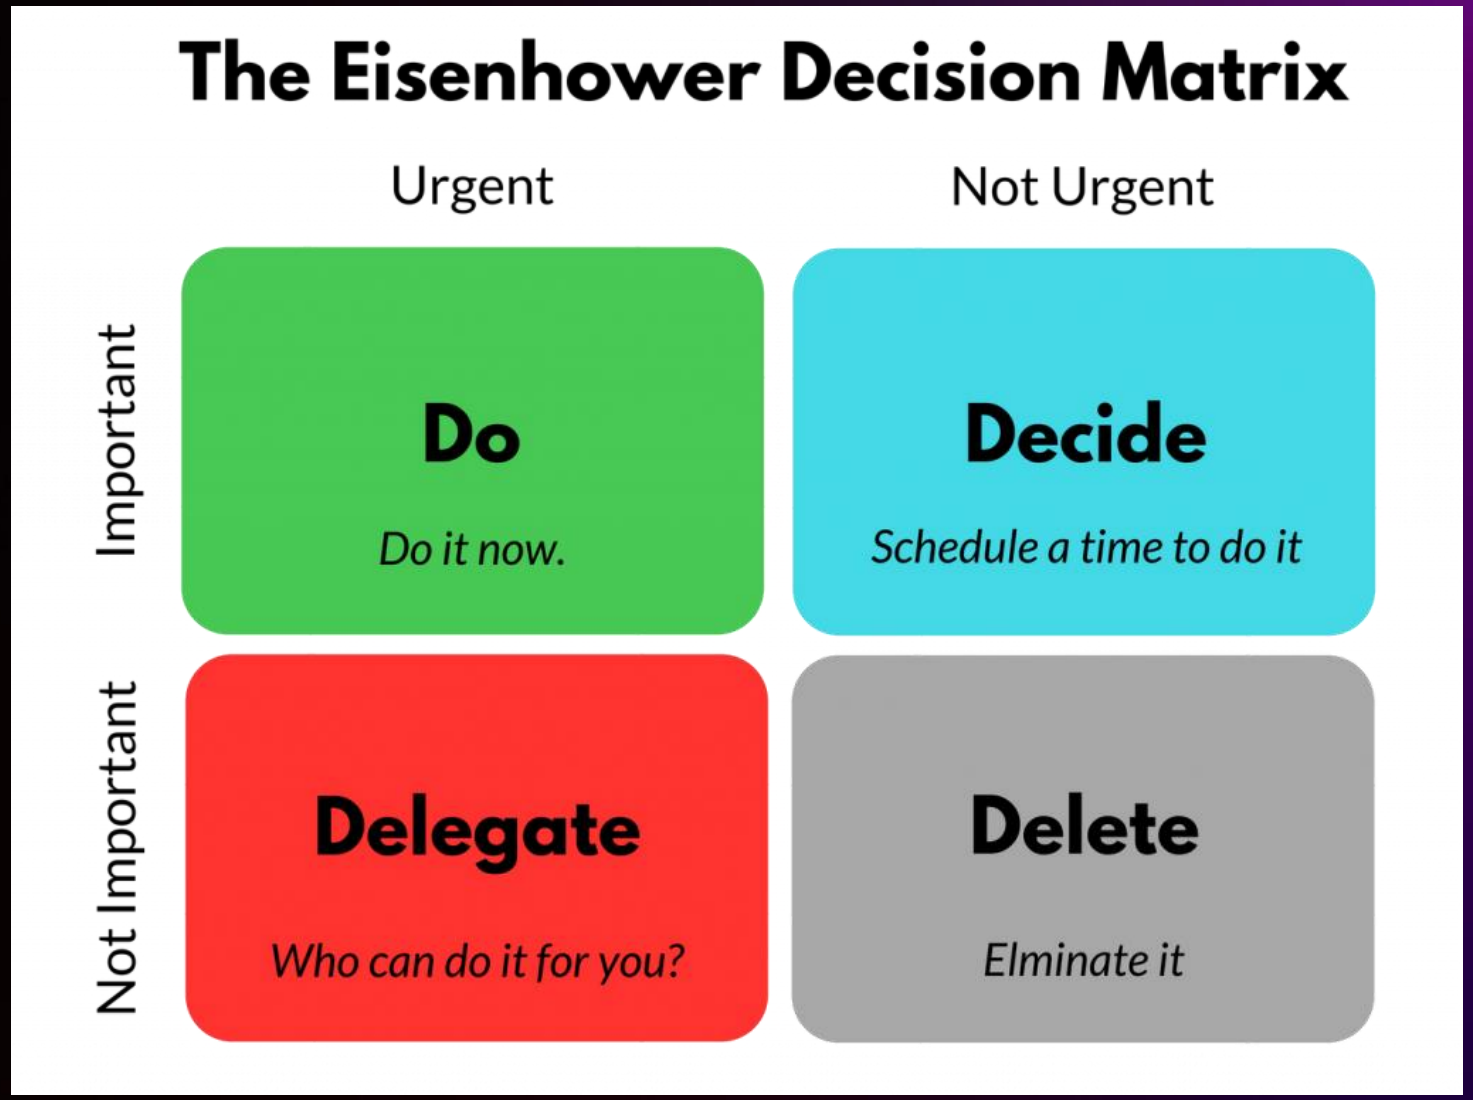

Criticality

Critical
Essential to the operation of the business

Urgent
Significant to the operation of the business

High
Supporting or providing important functions to the business

Medium
Supporting or providing noncritical functions

Low
Having limited or no impact on the operation of the business

The “when is lunch?” criticality sanity check

Critical

Who eats at 3am? You are on call and the time to deal with the incident is now

Urgent

Order delivery or get someone to bring you food; prepare to reschedule evening plans and work the incident now

High

Get food and eat at your desk while you continue to work on the issue

Medium

Get food when you are hungry; continue to work the issue when you are done

Low

Enjoy lunch; this can wait

Alerts and actions

Focus	Area	Criticality	Condition to alert on	Owner	Action
Workload	Fault	Critical	Retail site not available	Ops	Investigate incident, identify fault, and work to resolve issue immediately
Application	Fault	High	Newly deployed functionality generating errors	Dev	Rollback change if viable, remediate in place otherwise
Workload	Security	Critical	Sudden spike in outbound traffic volume from webserver	Security	Perform security incident response, isolate system and investigate
Application	Performance	Medium	Degraded performance	Dev/ops	Identify cause and remediate
Compliance	Compliance	High or critical	Noncompliance mandatory reporting requirement	Chief compliance officer	Prepare mandatory reporting documentation and deliver to regulators within required time

and many more . . .

These need to be dealt with but do not merit alerting

Focus	Area	Criticality	Condition to Alert on	Owner	Action
Application	Fault	Noncritical	Unhealthy instance in ASG	Ops	ASG automatically replaces instance
Application	Fault	Noncritical	Intermittent errors and retries from application	Dev	Investigate the errors and determine course of action
Application	Security	Low	Multiple failed login attempts	Security	If it was a single account over a few minutes and then successful or password recovered, no action
Application	Accounting	Medium	Trending to exceed budget	Dev/Ops/LOB	Determine why and then determine action
Application	Performance	Non-Critical	Minor decrease in application performance following implementation of new features	Dev/LOB	Create issue to address the change and prioritize as part of the development backlog

and many more . . .

Exercise #5 – Taking action

Refer to your case study:

- What will we do when the information is received?
- Who will perform the action?
- What will be the expected outcome?

Informing decisions

REPORTING AND TRENDING

Reporting and trending

- Monitoring need not always result in an alert
- Items for investigation may not reveal themselves immediately
- Trending over time may reveal opportunities for improvement

Reporting and trending

Personas and subscribers	Topic	Queries How we will access the data	Format(s) How the information is delivered
CFO, finance team	Billing and projected usage	AWS Cost and Usage Reports	Monthly report of budget and spend by department
CIO, CISO, security, chief compliance officer	Security controls and compliance status	AWS Systems Manager Compliance configuration and AWS Config reports	Dashboard Summary report delivered monthly
Product owner, ops teams	Monitoring and incident response	Mean time to detect (MTTD) Mean time to identify (MTTI) Mean time to recovery (MTTR)	Reports present in an operations metrics review
and many more . . .			

Reporting and trending

Personas and subscribers	Topic	Queries How we will access the data	Format(s) How the information is delivered
CIO, product owner, ops teams	Application health	Current health check status on components	Dashboard
Product owner, ops teams	Total availability	Requests fulfilled/requests made	Dashboard Summary report delivered monthly
Ops teams, dev teams, QA teams	Performance trending	Time to satisfy user transactions at points of observation in traces	Heat map of component utilization by time of day and day of week
CEO, product owner, ops teams	Total utilization: User activity and volume	Active sessions and utilization	Dashboard Summary report delivered monthly
and many more . . .			

Exercise #6 – Alerting, reporting, and dashboarding

Refer to your case study:

- Who needs this information?
  Refer to our roles!
- When do they need it?
- What is the best way to consume it?

What have we learned?

What have we learned?

- Why you need a monitoring strategy
- What goes into a monitoring strategy
- The contrast between observability and monitoring
- The steps to take to develop your strategy

What have we learned?

A simple mechanism to establish a strategy and create a plan focusing on

- Who will consume your insights and what are their needs
- Outcomes and KPIs
- Telemetry, metrics, and thresholds
- Time to action, reporting, alerting, and criticality

The output of that effort is the plan for implementing monitoring

Helpful resources

The One Observability Demo workshop

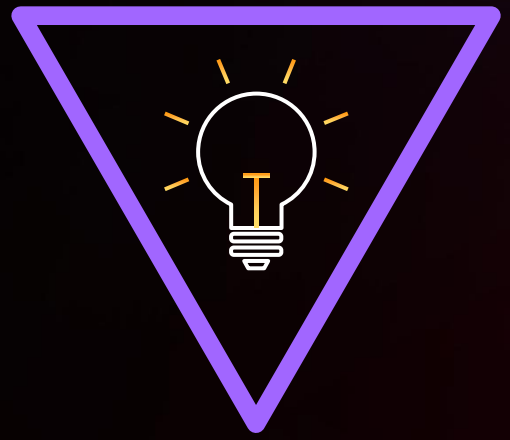

Providing an hands-on experience on the wide variety of toolsets AWS offers to setup monitoring and observability on your applications

https://observability.workshop.aws/

Architecture resources

- **AWS Well-Architected Framework**
  The official best practices for architecting in the AWS Cloud
  aws.amazon.com/architecture/well-architected
- **AWS Well-Architected Labs**
  Hands-on labs to help you learn, measure, and build using architectural best practices
  wellarchitectedlabs.com
- **AWS Architecture Center**
  Official AWS repository for all architecture resources
  aws.amazon.com/architecture
- **AWS Solutions Library**
  Vetted reference implementations and Well-Architected patterns
  aws.amazon.com/solutions

Thank you!

Paul Moran
@pdmoran_aws

Chris Kozlowski
kozlowck@amazon.com

Please complete the session survey in the **mobile app**