Workload Assessment Report

Assessment Name : [redacted]

Assessment Framework : [redacted]

Workload Name : [redacted]

Report Date : [redacted]

Workload Assessment Summary

Summary by Questions

Pillar	Total Questions	Verified	High Risk Issues (HRI)	Medium Risk Issues (MRI)	Low Risk Issues (LRI)	Not Applicable
Cost Optimization	11	0	10	0	1	0
Operational Excellence	11	0	11	0	0	0
Performance Efficiency	8	0	8	0	0	0
Reliability	13	0	13	0	0	0
Security	11	0	11	0	0	0
Sustainability	6	0	0	6	0	0
Total (Includes all Pillars)	**60**	**0**	**53**	**6**	**1**	**0**

Summary by Best Practices

Pillar	Total Best Practices	Verified	High Risk Issues (HRI)	Medium Risk Issues (MRI)	Low Risk Issues (LRI)	Not Applicable
Cost Optimization	49	9	16	13	11	0
Operational Excellence	83	7	30	29	17	0
Performance Efficiency	42	9	16	11	6	0
Reliability	67	6	31	29	1	0
Security	65	11	24	18	12	0
Sustainability	28	3	0	17	8	0
Total (Includes all Pillars)	**334**	**45**	**117**	**117**	**55**	**0**

Disclaimer : If there are any deactivated or deleted cloud accounts that belongs to any workload for which there is an assessment, then resources violated (belongs to cloud accounts) that are shown under the policy violations will not be shown after accounts deactivation/deletion, this could result in policy shown as violation but with out any resources. If there are any policies that are removed from the list of available policies for that tenant and if these policies are mapped to any framework where there is an active assessment, then violated resources under that policy will be shown as empty.

Workload Assessment Summary By Best Practices

Pillar Name: Cost Optimization				
Question: 01. How do you implement cloud financial management?				
Question Risk: High				
Best Practice	**Nature**	**Severity**	**Status**	**Violated Resources**
01. Establish a cost optimization function	Manual	High	Verified	
02. Establish a partnership between finance and technology	Manual	High	Open	
Recommendation: Define key members : Ensure that all relevant members of your finance and technology teams participate in the partnership. Relevant finance members will be those having interaction with the cloud bill, this will typically be CFOs, financial controllers, financial planners, business analysts, procurement and sourcing. Technology members will typically be product and application owners, technical managers and representatives from all teams that build on the cloud. Other members may include business unit owners, such as marketing that will influence usage of products, and third parties such as consultants to ensure alignment to your goals and mechanisms, and assist reporting. Define topics for discussion : Define the topics that are common across the teams, or will need a shared understanding. Follow cost from that time it is created, until the bill is paid. Note any members involved, and organizational processes that are required to be applied. Understand each step or process it goes through and the associated information, such as pricing models available, tiered pricing, discount models, budgeting, and financial requirements. Establish regular cadence : The group needs to come together regularly against their goals and metrics. A typical cadence involves reviewing the state of the organization, reviewing any programs currently running, then review overall financial and optimization metrics. Then key workloads are then reported on in greater detail.				
03. Establish cloud budgets and forecasts	Automated	High	Verified	
04. Implement cost awareness in your organizational processes	Manual	High	Open	
Recommendation: Identify relevant organizational processes : Each organizational unit reviews their processes and identifies processes that impact cost and usage. Any processes that result in the creation or termination of a resource need to be included for review. Also look for processes that can support cost awareness in your business, such as incident management and training. Update processes with cost awareness: Each process is modified to be made cost aware. The process may require additional pre checks, such as assessing the impact of cost, or post checks validating that the expected changes in cost and usage occurred. Supporting processes such as training and incident management can be extended to include items for cost and usage.				
05. Report and notify on cost optimization	Manual	Low	Open	
Recommendation: Configure AWS Budgets : Configure AWS Budgets on all accounts for your workload. Set a budget for the overall account spend, and a budget for the workload by using tags.Well Architected Labs: Cost and Governance Usage Report on cost optimization : Set up a regular cycle to discuss and analyze the efficiency of the workload. Using the metrics established, report on the metrics achieved and the cost of achieving them. Look to identify any negative trends to be able to fix them. Also look for positive trends that you can promote across your organization. Reporting should involve representatives from the application teams and owners, finance, and management. Well Architected Labs: Visualization				
06. Monitor cost proactively	Automated	Medium	Verified	
07. Keep up to date with new service releases	Manual	Medium	Open	
Recommendation: Subscribe to blogs: Go to the AWS blogs pages and subscribe to the What's New Blog and other relevant blogs. You can sign up on the communication preference page with your email address.				

Pillar Name: Cost Optimization

Question: 01. How do you implement cloud financial management?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Subscribe to AWS News: Regularly review the AWS News Blog and What's New with AWS for information on new service and feature releases. Subscribe to the RSS feed, or with your email to follow announcements and releases. Follow AWS Price Reductions: Regular price cuts on all our services has been a standard way for AWS to pass on the economic efficiencies to our customers gained from our scale. As of April 2022, AWS has reduced prices 115 times since it was launched in 2006. If you have any pending business decisions due to price concerns, you can review them again after price reductions and new service integrations. You can learn about the previous price reductions efforts, including Amazon Elastic Compute Cloud (Amazon EC2) instances, in the price reduction category of the AWS News Blog. AWS events and meetups: Attend your local AWS summit, and any local meetups with other organizations from your local area. If you cannot attend in person, try to attend virtual events to hear more from AWS experts and other customers business cases. Meet with your account team: Schedule a regular cadence with your account team, meet with them and discuss industry trends and AWS services. Speak with your account manager, Solutions Architect, and support team.				
08. Create a cost-aware culture	Manual	Low	Open	

Recommendation:
Report cloud costs to technology teams: To raise cost awareness, and establish efficiency KPIs for finance and business stakeholders.
Inform stakeholders or team members about planned changes: Create an agenda item to discuss planned changes and the cost benefit impact on the workload during weekly change meetings.
Meet with your account team: Establish a regular meeting cadence with your account team, and discuss industry trends and AWS services. Speak with your account manager, architect, and support team.
Share success stories: Share success stories about cost reduction for any workload, AWS account, or organization to create a positive attitude and encouragement around cost optimization.
Training: Ensure technical teams or team members are trained for awareness of resource costs on AWS Cloud.
AWS events and meetups: Attend local AWS summits, and any local meetups with other organizations from your local area.
Subscribe to blogs: Go to the AWS blogs pages and subscribe to the What's New Blog and other relevant blogs to follow new releases, implementations, examples, and changes shared by AWS.
Resources

09. Quantify business value from cost optimization	Manual	Medium	Open	

Recommendation:
Executing cost optimization best practices: For example, resource lifecycle management reduces infrastructure and operational costs and creates time and unexpected budget for experimentation. This increases organization agility and uncovers new opportunities for revenue generation.
Implementing automation: For example, Auto Scaling, which ensures elasticity at minimal effort, and increases staff productivity by eliminating manual capacity planning work. For more details on operational resiliency, refer to the Well Architected Reliability Pillar whitepaper.
Forecasting future AWS costs: Forecasting helps finance stakeholders to set expectations with other internal and external organization stakeholders, and helps improve your organization's financial predictability. AWS Cost Explorer can be used to perform forecasting for your cost and usage.

Pillar Name: Cost Optimization

Question: 02. How do you govern usage?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Develop policies based on your organization requirements	Automated	High	Verified	
02. Implement goals and targets	Manual	High	Open	

Recommendation:
Define expected usage levels : Focus on usage levels to begin with. Engage with the application owners, marketing and greater business teams to understand what the expected usage levels will be for the workload. How will customer demand change over time, and will there be any changes due to seasonal increases or marketing campaigns.

Pillar Name: Cost Optimization

Question: 02. How do you govern usage?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Define workload resourcing and costs : With the usage levels defined, quantify the changes in workload resources required to meet these usage levels. You may need to increase the size or number of resources for a workload component, increase data transfer, or change workload components to a different service at a specific level. Specify what the costs will be at each of these major points, and what the changes in cost will be when there are changes in usage. Define business goals : Taking the output from the expected changes in usage and cost, combine this with expected changes in technology, or any programs that you are running, and develop goals for the workload. Goals must address usage, cost and the relation between the two. Ensure that there are organizational programs, for example capability building like training and education, if there are expected changes in cost without changes in usage. Define targets : For each of the defined goals specify a measurable target. If a goal is to increase efficiency in the workload, the target will quantify the amount of improvement, typical in business outputs per dollar spent, and when it will be delivered.				
03. Implement an account structure	Automated	High	Verified	
04. Implement groups and roles	Manual	Low	Open	
Recommendation: Implement groups : Using the groups of users defined in your organizational policies, implement the corresponding groups, if necessary. Refer to the security pillar for best practices on users, groups, and authentication. Implement roles and policies : Using the actions defined in your organizational policies, create the required roles and access policies. Refer to the security pillar for best practices on roles and policies.				
05. Implement cost controls	Manual	Medium	Open	
Recommendation: Implement notifications on spend : Using your defined organization policies, create AWS budgets to provide notifications when spending is outside of your policies. Configure multiple cost budgets, one for each account, which notifies you about overall account spending. Then configure additional cost budgets within each account for smaller units within the account. These units vary depending on your account structure. Some common examples are AWS Regions, workloads (using tags), or AWS services. Ensure that you configure an email distribution list as the recipient for notifications, and not an individual's email account. You can configure an actual budget for when an amount is exceeded, or use a forecasted budget for notifying on forecasted usage. Implement controls on usage : Using your defined organization policies, implement IAM policies and roles to specify which actions users can perform and which actions they cannot perform. Multiple organizational policies may be included in an AWS policy. In the same way that you defined policies, start broadly and then apply more granular controls at each step. Service limits are also an effective control on usage. Implement the correct service limits on all your accounts.				
06. Track project lifecycle	Automated	Low	Open	12
Recommendation: Perform workload reviews : As defined by your organizational policies, audit your existing projects. The amount of effort spent in the audit should be proportional to the approximate risk, value, or cost to the organization. Key areas to include in the audit would be risk to the organization of an incident or outage, value, or contribution to the organization (measured in revenue or brand reputation), cost of the workload (measured as total cost of resources and operational costs), and usage of the workload (measured in number of organization outcomes per unit of time). If these areas change over the lifecycle, adjustments to the workload are required, such as full or partial decommissioning.				

Pillar Name: Cost Optimization

Question: 03. How do you monitor usage and cost?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Configure detailed information sources	Manual	High	Open	

Pillar Name: Cost Optimization

Question: 03. How do you monitor usage and cost?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Recommendation: Configure the cost and usage report : Using the billing console, configure at least one cost and usage report. Configure a report with hourly granularity that includes all identifiers and resource IDs. You can also create other reports with different granularities to provide higher level summary information. Configure hourly granularity in cost explorer : Using the billing console, enable Hourly and Resource Level Data. NOTE: There will be associated costs with enabling this feature, please refer to the pricing. Configure application logging : Ensure that your application logs each business outcome that it delivers so it can be tracked and measured. Ensure that the granularity of this data is at least hourly to be matched with the cost and usage data. Refer to the Operational Excellence pillar for more detail on logging and monitoring.				
02. Identify cost attribution categories	Manual	High	Open	
Recommendation: Define your organization categories : Meet with stakeholders to define categories that reflect your organization's structure and requirements. These will directly map to the structure of existing financial categories, such as business unit, budget, cost center, or department. Look at the outcomes the cloud delivers for your business, such as training or education, as these are also organization categories. Multiple categories can be assigned to a resource, and a resource can be in multiple different categories, so define as many categories as needed. Define your functional categories : Meet with stakeholders to define categories that reflect the functions that you have within your business. This may be the workload or application names, and the type of environment, such as production, testing, or development. Multiple categories can be assigned to a resource, and a resource can be in multiple different categories, so define as many categories as needed.				
03. Establish organization metrics	Manual	High	Open	
Recommendation: Define workload outcomes : Meet with the stakeholders in the business and define the outcomes for the workload. These are a primary measure of customer usage and must be business metrics and not technical metrics. There should be a small number of high level metrics (less than five) per workload. If the workload produces multiple outcomes for different use cases, then group them into a single metric. Define workload component outcomes : Optionally, if you have a large and complex workload, or can easily break your workload into components (such as microservices) with well defined inputs and outputs, define metrics for each component. The effort should reflect the value and cost of the component. Start with the largest components and work towards the smaller components.				
04. Configure billing and cost management tools	Manual	High	Open	
Recommendation: Create a Cost Optimization group : Configure your account and create a group that has access to the required Cost and Usage reports. This group must include representatives from all teams that own or manage an application. This ensures that every team has access to their cost and usage information. Configure AWS Budgets : Configure AWS Budgets on all accounts for your workload. Set a budget for the overall account spend, and a budget for the workload by using tags. Configure AWS Cost Explorer : Configure AWS Cost Explorer for your workload and accounts. Create a dashboard for the workload that tracks overall spend, and key usage metrics for the workload. Configure advanced tooling : Optionally, you can create custom tooling for your organization that provides additional detail and granularity. You can implement advanced analysis capability using Amazon Athena, and dashboards using Amazon QuickSight.				
05. Add organization information to cost and usage	Manual	Medium	Open	
Recommendation: Define a tagging schema : Gather all stakeholders from across your business to define a schema. This typically includes technical, financial, and management people. Define a list of tags that all resources must have, as well as a list of tags that resources should have. Ensure that the tag names and values are consistent across your organization. Tag resources : Using your defined cost attribution categories, place tags on all resources in your workloads according to the categories. Use tools such as the CLI, Tag Editor, or Systems Manager, to increase efficiency. Implement Cost Categories : You can create Cost Categories without implementing tagging, Cost Categories use the existing cost and usage dimensions. Create category rules from your schema and implement it into cost categories. Automate tagging : To ensure that you maintain high levels of tagging across all resources, automate tagging so that resources are automatically tagged when they are created. Use the features within the service, or services such as AWS CloudFormation, to ensure that resources are tagged when created. You can also create a custom microservice that scans the workload periodically and removes any resources that are not tagged, which is ideal for test and development environments. Monitor and report on tagging : To ensure that you maintain high levels of tagging across your organization, report and monitor the tags across your workloads. You can use AWS Cost Explorer to view the cost of tagged and untagged resources, or use services such as Tag Editor. Regularly review the number of untagged resources and take action to add tags until you reach the desired level of tagging.				

Pillar Name: Cost Optimization

Question: 03. How do you monitor usage and cost?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
06. Allocate costs based on workload metrics	Manual	Low	Open	

Recommendation:
Allocate costs to workload metrics : Using the defined metrics and tagging configured, create a metric that combines the workload output and workload cost. Use the analytics services such as Athena and QuickSight to create an efficiency dashboard for the overall workload, and any components.

Pillar Name: Cost Optimization

Question: 04. How do you decommission resources?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Track resources over their life time	Automated	High	Verified	
02. Implement a decommissioning process	Automated	High	Open	12

Recommendation:
Create and implement a decommissioning process. : Working with the workload developers and owners, build a decommissioning process for the workload and its resources. The process should cover the method to verify if the workload is in use, and also if each of the workload resources are in use. The process also covers the steps necessary to decommission the resource, removing them from service while ensuring compliance with any regulatory requirements. Any associated resources are also covered, such as licenses or attached storage. Finally the process provides notification to the workload owners that the decommissioning process has been executed.

03. Decommission resources	Automated	Medium	Open	12

Recommendation:
Decommission resources : Using the decommissioning process, decommission each of the resources that have been identified as orphaned.

04. Decommission resources automatically	Automated	Low	Open	12

Recommendation:
Implement AWS Auto Scaling : For resources that are supported, configure them with AWS Auto Scaling.
Configure CloudWatch to Terminate Instances : Instances can be configured to terminate using CloudWatch alarms. Using the metrics from the decommissioning process, implement an alarm with an EC2 action. Ensure you verify the operation in a non production environment before rolling out.
Implement code within the workload : You can use the AWS SDK or AWS CLI to decommission workload resources. Implement code within the application that integrates with AWS and terminates or removes resources that are no longer used.

05. Enforce data retention policies	Manual	Medium	Open	

Recommendation:
Implement AWS Auto Scaling : For resources that are supported, configure them with AWS Auto Scaling.
Configure CloudWatch to Terminate Instances : Instances can be configured to terminate using CloudWatch alarms. Using the metrics from the decommissioning process, implement an alarm with an EC2 action. Ensure you verify the operation in a non production environment before rolling out.
Implement code within the workload : You can use the AWS SDK or AWS CLI to decommission workload resources. Implement code within the application that integrates with AWS and terminates or removes resources that are no longer used.

Pillar Name: Cost Optimization

Question: 05. How do you evaluate cost when you select services?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Identify organization requirements for cost	Manual	High	Open	
Recommendation: Identify organization requirements for cost : Meet with team members from your organization, they include product management, application owners, development and operational teams, management, and finance. Prioritize the Well Architected pillars for this workload and its components, the output is be a list of the pillars in order. You can also add a weighting to each, this can help indicate how much additional focus a pillar has, or how similar the focus is between two pillars.				
02. Analyze all components of the workload	Automated	High	Open	11
Recommendation: List the workload components : Build the list of all the workload components, this is used as verification to check that each component was analyzed. The effort spent should reflect the criticality to the workload as defined by organization priorities. Grouping together resources functionally improves efficiency, i.e. production database storage, if there are multiple databases. Prioritize component list : Take the component list and prioritize it in order of effort. This is typically in order of the cost of the component from most expensive to least expensive, or the criticality as defined by organization priorities. Perform the analysis : For each component on the list, review the options and services available and chose the option that aligns best with your organizational priorities.				
03. Perform a thorough analysis of each component	Manual	High	Open	
Recommendation: Perform a thorough analysis : Using the component list, work through each component from the highest priority to lowest priority. For the higher priority and more costly components, perform additional analysis and assess all available options and their long term impact. For lower priority components, assess if changes in usage would change the priority of the component, and then perform an analysis of appropriate effort.				
04. Select software with cost-effective licensing	Manual	Low	Open	
Recommendation: Analyze license options : Review the licensing terms of available software. Look for open source versions that have the required functionality, and whether the benefits of licensed software outweigh the cost. Favorable terms will align the cost of the software to the benefit it provides. Analyze the software provider : Review any historical pricing or licensing changes from the vendor. Look for any changes that do not align to outcomes, such as punitive terms for running on specific vendors hardware or platforms. Also look for how they execute audits and penalties that could be imposed.				
05. Select components of this workload to optimize cost in line with organization priorities	Manual	Medium	Open	
Recommendation: Select each service to optimize cost : Using your prioritized list and analysis, select each option that provides the best match with your organizational priorities.				
06. Perform cost analysis for different usage over time	Automated	Medium	Verified	

Pillar Name: Cost Optimization

Question: 06. How do you meet cost targets when you select resource type, size and number?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Perform cost modeling	Automated	High	Open	

Recommendation:
Perform cost modelling : Deploy the workload or a proof of concept, into a separate account with the specific resource types and sizes to test. You run the workload with the test data and record the output results, along with the cost data for the time the test was run. You then redeploy the workload or change the resource types and sizes and re run the test.

02. Select resource type, size, and number based on data	Automated	Medium	Verified	
03. Select resource type, size, and number automatically based on metrics	Automated	Medium	Verified	

Pillar Name: Cost Optimization

Question: 07. How do you use pricing models to reduce cost?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Perform pricing model analysis	Manual	High	Open	

Recommendation:
Perform a commitment discount analysis : Using Cost Explorer in your account review the Savings Plans and Reserved Instance recommendations. To ensure you implement the correct recommendations with the required discounts and risk, follow the Well Architected labs.
Analyze workload elasticity : Using the hourly granularity in Cost Explorer, or a custom dashboard analyze the workload elasticity. Look for regular changes in the number of instances that are running, short duration instances are candidates for spot instances or spot fleet.

02. implement Regions based on cost	Manual	Medium	Open	

Recommendation:
Review region pricing : Analyze the workload costs in the current region. Starting with the highest costs by service and usage type, calculate the costs in other regions that are available. If the forecasted saving outweighs the cost of moving the component or workload, migrate to the new region.

03. Select third-party agreements with cost-efficient terms	Manual	Medium	Open	

Recommendation:
Analyze third party agreements and terms : Review the pricing in third party agreements. Perform modelling for different levels of your usage, and factor in new costs such as new service usage, or increases in current services due to workload growth. Decide if the additional costs provide the required benefits to your business.

04. Implement pricing models for all components of this workload	Manual	Low	Open	

Recommendation:
Implement pricing models : Using your analysis results purchase Savings Plans (SP's), Reserved Instances (RI's) or implement spot. If it is your first RI purchase then choose the top 5 or 10 recommendations in the list, then monitor and analyze the results over the next month or two. Purchase small numbers of commitment discounts regular cycles, for example every 2 weeks or monthly. Implement spot instances for workloads that can be interrupted or are stateless.
Workload Review Cycle : Implement a review cycle for the workload that specifically analyzes pricing model coverage. Once the workload has the required coverage, purchase additional commitment discounts every 2 4 weeks or as your organization usage changes.

Pillar Name: Cost Optimization

Question: 07. How do you use pricing models to reduce cost?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
05. Perform pricing model analysis at the management account level	Manual	Low	Open	

Recommendation:
Perform a commitment discount analysis : Using Cost Explorer in your account review the Savings Plans and Reserved Instance recommendations. To ensure you implement the correct recommendations with the required discounts and risk, follow the Well Architected labs.

Pillar Name: Cost Optimization

Question: 08. How do you plan for data transfer charges?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Perform data transfer modeling	Manual	High	Open	

Recommendation:
Calculate data transfer costs : Use the AWS pricing pages and calculate the data transfer costs for the workload. Also calculate the data transfer costs at different usage levels, for both increases and reductions in workload usage. Where there are multiple options for the workload architecture cost each option for comparison.
Link costs to outcomes : For each data transfer cost incurred, specify the outcome that it achieves for the workload. If it is transfer between components it may be for decoupling, if it is between availability zones it may be for redundancy.

02. Select components to optimize data transfer cost	Manual	Medium	Open	

Recommendation:
Select components for data transfer : Using the data transfer modelling, focus on where the largest data transfer costs are or where they would be if the workload usage changes. Look for alternative architectures or additional components that remove or reduce the need for data transfer, or lower its cost.

03. Implement services to reduce data transfer costs	Manual	Medium	Open	

Recommendation:
Implement services : Using the data transfer modelling, look at where the largest costs and highest volume flows are. Review the AWS services and assess whether there is a service that reduces or removes the transfer, specifically networking and content delivery. Also look for caching services where there is repeated access to data, or large amounts of data.

Pillar Name: Cost Optimization

Question: 09. How do you manage demand, and supply resources?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Perform an analysis on the workload demand	Manual	High	Open	

Pillar Name: Cost Optimization

Question: 09. How do you manage demand, and supply resources?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Recommendation: Analyze existing workload data : Analyze data from the existing workload, previous versions of the workload, or predicted usage patterns. Use log files and monitoring data to gain insight on how customers use the workload. Typical metrics are the actual demand, in requests per second, the times when the rate of demand changes or when it is at different levels, and the rate of change of demand. Ensure you analyze a full cycle of the workload, ensuring you collect data for any seasonal changes such as end of month or end of year events. The effort reflected in the analysis should reflect the workload characteristics. The largest effort should be placed on high value workloads that have the largest changes in demand. The least effort should be placed on low value workloads that have minimal changes in demand. Common metrics for value are risk, brand awareness, revenue or workload cost. Forecast outside influence : Meet with team members from across the organization that can influence or change the demand in the workload. Common teams would be sales, marketing or business development. Work with them to know the cycles they operate with, and if there are any events that would change the demand of the workload. Forecast the workload demand with this data.				
02. Implement a buffer or throttle to manage demand	Manual	Medium	Open	
Recommendation: Analyze the client requirements : Analyze the client requests to determine if they are capable of performing retries. For clients that cannot perform retries, buffers will need to be implemented. Analyze the overall demand, rate of change, and required response time to determine the size of throttle or buffer required. Implement a buffer or throttle : Implement a buffer or throttle in the workload. A queue such as SQS can provide a buffer to your workload components. Amazon API Gateway can provide throttling for your workload components.				
03. Supply resources dynamically	Manual	Low	Open	
Recommendation: Configure time based scheduling : For predictable changes in demand, time based scaling can provide the correct amount of resources in a timely manner. It is also useful if resource creation and configuration is not fast enough to respond to changes in demand. Using the workload analysis configure scheduled scaling using AWS Auto Scaling. Configure Auto Scaling : To configure scaling based on active workload metrics, use Amazon Auto Scaling. Use the analysis and configure auto scaling to trigger on the correct resource levels, and ensure that the workload scales in the required time.				

Pillar Name: Cost Optimization

Question: 10. How do you evaluate new services?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Develop a workload review process	Manual	High	Open	
Recommendation: Define review frequency : Define how frequently the workload and its components should be reviewed. This is a combination of factors and may differ from workload to workload within your organization, it may also differ between components in the workload. Common factors include: the importance to the organization measured in terms of revenue or brand, the total cost of running the workload (including operation and resource costs), the complexity of the workload, how easy is it to implement a change, any software licensing agreements, and if a change would incur significant increases in licensing costs due to punitive licensing. Components can be defined functionally or technically, such as web servers and databases, or compute and storage resources. Balance the factors accordingly and develop a period for the workload and its components. You may decide to review the full workload every 18 months, review the web servers every 6 months, the database every 12 months, compute and short term storage every 6months, and long term storage every 12 months. Define review thoroughness : Define how much effort is spent on the review of the workload or workload components. Similar to the review frequency this is a balance of multiple factors. You may decide to spend 1 week of analysis on the database component, and 4 hours for storage reviews.				
02. Review and analyze this workload regularly	Manual	Medium	Open	

Pillar Name: Cost Optimization

Question: 10. How do you evaluate new services?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources

Recommendation:
Regularly review the workload : Using your defined process, perform reviews with the frequency specified. Ensure you spend the correct amount of effort on each component. This process would be similar to the initial design process where you selected services for cost optimization. Analyze the services and the benefits they would bring, this time factor in the cost of making the change, not just the long term benefits.
Implement new services : If the outcome of the analysis is to implement changes, first perform a baseline of the workload to know the current cost per output. Implement the changes, then perform an analysis to confirm the new cost per output.

Pillar Name: Cost Optimization

Question: 11. How do you evaluate the cost of effort?

Question Risk: Low

Best Practice	Nature	Severity	Status	Violated Resources
01. Perform automations for operations	Manual	Low	Open	

Recommendation:
Regularly review the workload : Using your defined process, perform reviews with the frequency specified. Ensure you spend the correct amount of effort on each component. This process would be similar to the initial design process where you selected services for cost optimization. Analyze the services and the benefits they would bring, this time factor in the cost of making the change, not just the long term benefits.
Implement new services : If the outcome of the analysis is to implement changes, first perform a baseline of the workload to know the current cost per output. Implement the changes, then perform an analysis to confirm the new cost per output.

Pillar Name: Operational Excellence

Question: 01. How do you determine what your priorities are?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Evaluate external customer needs	Manual	High	Open	

Recommendation:
Understand business needs: Business success is enabled by shared goals and understanding across stakeholders, including business, development, and operations teams.
Review business goals, needs, and priorities of external customers: Engage key stakeholders, including business, development, and operations teams, to discuss goals, needs, and priorities of external customers. This ensures that you have a thorough understanding of the operational support that is required to achieve business and customer outcomes.
Establish shared understanding: Establish shared understanding of the business functions of the workload, the roles of each of the teams in operating the workload, and how these factors support your shared business goals across internal and external customers.

02. Evaluate internal customer needs	Manual	High	Open	

Pillar Name: Operational Excellence

Question: 01. How do you determine what your priorities are?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Recommendation: Understand business needs: Business success is enabled by shared goals and understanding across stakeholders including business, development, and operations teams. Review business goals, needs, and priorities of internal customers: Engage key stakeholders, including business, development, and operations teams, to discuss goals, needs, and priorities of internal customers. This ensures that you have a thorough understanding of the operational support that is required to achieve business and customer outcomes. Establish shared understanding: Establish shared understanding of the business functions of the workload, the roles of each of the teams in operating the workload, and how these factors support shared business goals across internal and external customers.				
03. Evaluate governance requirements	Manual	High	Open	
Recommendation: Understand governance requirements: Evaluate internal governace factors, such as program or organizational policy, program policies, issue or system specific policies, standards, procedures, baselines, and guidelines. Validate that you have mechanisms to identify changes to governance. If no governance requirements are identified, ensure that you have applied due diligence to this determination				
04. Evaluate compliance requirements	Manual	High	Open	
Recommendation: Understand compliance requirements: Evaluate external factors, such as regulatory compliance requirements and industry standards, to ensure that you are aware of guidelines or obligations that might mandate or emphasize specific focus. If no compliance requirements are identified, ensure that due diligence was applied to the determination. Understand regulatory compliance requirements: Identify regulatory compliance requirements that you are legally obligated to satisfy. Use these requirements to focus your efforts. Examples include obligations from privacy and data protection acts. Understand industry standards and best practices: Identify industry standards and best practice requirements that apply to your workload, such as the Payment Card Industry Data Security Standard (PCI DSS). Use these requirements to focus your efforts. Understand internal compliance requirements: Identify compliance requirements and best practices that are established by your organization. Use these requirements to focus your efforts. Examples include information security policies and data classification standards.				
05. Evaluate threat landscape	Manual	Medium	Open	
Recommendation: Evaluate threat landscape: Evaluate threats to the business (for example, competition, business risk and liabilities, operational risks, and information security threats), so that you can include their impact when determining where to focus efforts. Maintain a threat model: Establish and maintain a threat model identifying potential threats, planned and in place mitigations, and their priority. Review the probability of threats manifesting as incidents, the cost to recover from those incidents and the expected harm caused, and the cost to prevent those incidents. Revise priorities as the contents of the threat model change.				
06. Evaluate tradeoffs	Manual	Medium	Open	
Recommendation: Evaluate tradeoffs: Evaluate the impact of tradeoffs between competing interests, to help make informed decisions when determining where to focus efforts. For example, accelerating speed to market for new features might be emphasized over cost optimization.				
07. Manage benefits and risks	Manual	Low	Open	
Recommendation: Manage benefits and risks: Balance the benefits of decisions against the risks involved. Identify benefits: Identify benefits based on business goals, needs, and priorities. Examples include time to market, security, reliability, performance, and cost. Identify risks: Identify risks based on business goals, needs, and priorities. Examples include time to market, security, reliability, performance, and cost. Assess benefits against risks and make informed decisions: Determine the impact of benefits and risks based on goals, needs, and priorities of your key stakeholders, including business, development, and operations. Evaluate the value of the benefit against the probability of the risk being realized and the cost of its impact. For example, emphasizing speed to market over reliability might provide competitive advantage. However, it may result in reduced uptime if there are reliability issues.				

Pillar Name: Operational Excellence

Question: 02. How do you structure your organization to support your business outcomes?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Resources have identified owners	Automated	High	Open	1604

Recommendation:
Resources have identified owners: Define what ownership means for the resource use cases in your environment. Specify and record owners for resources including at a minimum name, contact information, organization, and team. Store resource ownership information with resources using metadata such as tags or resource groups. Use AWS Organizations to structure accounts and implement policies to ensure ownership and contact information are captured.
Define forms of ownership and how they are assigned: Ownership may have multiple definitions in your organization with different uses cases. You may wish to define a "workload owner" as the individual who owns the risk and liability for the operation of a workload, and whom ultimately has authority to make decisions about the workload. You may wish to define ownership in terms of financial or administrative responsibility where ownership rolls up to a parent organization. A developer may be the owner of their development environment and be responsible for incidents that its operation causes. Their product lead may own responsibility for the financial costs associated to the operation of their development environments.
Define who owns an organization, account, collection of resources, or individual components: Define and record ownership in an appropriately accessible location organized to support discovery. Update definitions and ownership details as they change.
Capture ownership in the metadata for the resources: Capture resource ownership using metadata such as tags or resource groups, specifying ownership and contact information. Use AWS Organizations to structure accounts and ensure ownership and contact information are captured.

02. Processes and procedures have identified owners	Manual	High	Open	

Recommendation:
Process and Procedures have identified owners responsible for their definition: Capture the processes and procedures used in your environment and the individual or team responsible for their definition.
Identify process and procedures: Identify the operations activities conducted in support of your workloads. Document these activities in a discoverable location.
Define who owns the definition of a process or procedure: Uniquely identify the individual or team responsible for the specification of an activity. They are responsible to ensure it can be successfully performed by an adequately skilled team member with the correct permissions, access, and tools. If there are issues with performing that activity, the team members performing it are responsible to provide the detailed feedback necessary for the activitiy to be improved.
Capture ownership in the metadata of the activity artifact: Procedures automated in services like AWS Systems Manager, through documents, and AWS Lambda, as functions, support capturing metadata information as tags. Capture resource ownership using tags or resource groups, specifying ownership and contact information. Use AWS Organizations to create tagging polices and ensure ownership and contact information are captured.

03. Operations activities have identified owners responsible for their performance	Automated	High	Open	1604

Recommendation:
Operations activities have identified owners responsible for their performance: Capture the responsibility for performing processes and procedures used in your environment
Identify process and procedures: Identify the operations activities conducted in support of your workloads. Document these activities in a discoverable location.
Define who is responsible to perform each activity: Identify the team responsible for an activity. Ensure they have the details of the activity, and the necessary skills and correct permissions, access, and tools to perform the activity. They must understand the condition under which it is to be perform (for example, on an event or schedule). Make this information discoverable so that members of your organization can identify who they need to contact, team or individual, for specific needs.

04. Team members know what they are responsible for	Automated	High	Open	1605

Recommendation:
Enure team members understand their roles and responsibilities: Identify team members roles and responsibilities and ensure they understand the expectations of their role. Make this information discoverable so that members of your organization can identify who they need to contact, team or individual, for specific needs.

05. Mechanisms exist to identify responsibility and ownership	Automated	High	Open	1

Recommendation:
Mechanisms exist to identify responsibility and ownership: Provide accessible mechanisms for members of your organization to discover and identify ownership and responsibility. This will enable them to identify who to contact, team or individual, for specific needs.

06. Mechanisms exist to request additions, changes, and exceptions	Manual	Medium	Open	

Pillar Name: Operational Excellence

Question: 02. How do you structure your organization to support your business outcomes?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Recommendation: Mechanisms exist to request additions, changes, and exceptions: When standards are rigid innovation is constrained. Provide mechanisms for members of your organization to make requests to owners of processes, procedures, and resources in support of their business needs.				
07. Responsibilities between teams are predefined or negotiated	Manual	Low	Open	
Recommendation: Responsibilities between teams are predefined or negotiated: Specifying the methods by which teams interact, and the information necessary for them to support each other, can help minimize the delay introduced as requests are iteratively reviewed and clarified. Having specific agreements that define expectations (for example, response time, or fulfillment time) enables teams to make effective plans and resource appropriately.				

Pillar Name: Operational Excellence

Question: 03. How does your organizational culture support your business outcomes?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Executive Sponsorship	Manual	High	Open	
Recommendation: Executive Sponsorship: Senior leadership clearly sets expectations for the organization and evaluates success. Senior leadership is the sponsor, advocate, and driver for the adoption of best practices and evolution of the organization Set expectations: Define and publish goals for your organizations including how they will be measured. Track achievement of goals: Measure the incremental achievement of goals regularly and share the results so that appropriate action can be taken if outcomes are at risk. Provide the resources necessary to achieve your goals: Regularly review if resources are still appropriate, of if additional resources are needed based on: new information, changes to goals, responsibilities, or your business environment. Advocate for your teams: Remain engaged with your teams so that you understand how they are doing and if there are external factors affecting them. When your teams are impacted by external factors, reevaluate goals and adjust targets as appropriate. Identify obstacles that are impeding your teams progress. Act on behalf of your teams to help address obstacles and remove unnecessary burdens. Be a driver for adoption of best practices: Acknowledge best practices that have provide quantifiable benefits and recognize the creators and adopters. Encourage further adoption to magnify the benefits achieved. Be a driver for evolution of for your teams: Create a culture of continual improvement. Encourage both personal and organizational growth and development. Provide long term targets to strive for that will require incremental achievement over time. Adjust this vision to compliment your needs, business goals, and business environment as they change.				
02. Team members are empowered to take action when outcomes are at risk	Automated	High	Verified	
03. Escalation is encouraged	Automated	High	Verified	
04. Communications are timely, clear, and actionable	Automated	High	Open	
Recommendation: Communications are timely, clear, and actionable: Mechanisms are in place to provide notification of risks or planned events in a clear and actionable way with enough notice to allow appropriate responses. Document planned activities on a change calendar and provide notifications: Provide an accessible source of information where planned events can be discovered. Provide notifications of planned events from the same system. Track events and activity that may have an impact on your workload: Monitoring vulnerability notifications and patch information to understand vulnerabilities in the wild and potential risks associated to your workload components. Provide notification to team members so that they can take action.				

Pillar Name: Operational Excellence

Question: 03. How does your organizational culture support your business outcomes?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
05. Experimentation is encouraged	Manual	Medium	Open	

Recommendation:
Experimentation is encouraged: Encourage experimentation to support learning and innovation.
Experiment with a variety of technologies: Encourage experimentation with technologies that may have applicability now or in the future to the achievement of your business outcomes. This knowledge may inform future innovation.
Experiment with a goal in mind: Encourage experimentation with specific goals for team members to reach for, or with technologies that may have applicability in the near future. This knowledge may inform your innovation.
Provide structured time to experiment: Dedicate specific times when team members can be free of their normal responsibilities, so that they can focus on their experiments.
Provide the resources to support experimentation: Fund the resources required to conduct experiments (for example, software, or cloud resources).
Acknowledge success: Recognize the value yielded by experimentation. Understand that experiments with undesired outcomes are successful and have identified a path that will not lead to success. Team members are not punished for undesired outcomes from experiments.

06. Team members are enabled and encouraged to maintain and grow their skill sets	Manual	Medium	Open	

Recommendation:
Team members are enabled and encouraged to maintain and grow their skill sets: To adopt new technologies, support innovation, and to support changes in demand and responsibilities in support of your workloads continuing education is necessary.
Provide resources for education: Provided dedicated structured time, access to training materials, lab resources, and support participation in conferences and professional organizations that provide opportunities for learning from both educators and peers. Provide junior team members' access to senior team members as mentors or allow them to shadow their work and be exposed to their methods and skills. Encourage learning about content not directly related to work in order to have a broader perspective.
Team education and cross team engagement: Plan for the continuing education needs of your team members. Provide opportunities for team members to join other teams (temporarily or permanently) to share skills and best practices benefiting your entire organization
Support pursuit and maintenance of industry certifications: Support your team members acquiring and maintaining industry certifications that validate what they have learned, and acknowledge their accomplishments.

07. Resource teams appropriately	Manual	Medium	Open	

Recommendation:
Resource teams appropriately: Ensure you have an understanding of the success of your teams and the factors that contribute to their success or lack of success. Act to support teams with appropriate resources.
Understand team performance: Measure the achievement of operational outcomes and the development of assets by your teams. Track changes in output and error rate over time. Engage with teams to understand the work related challenges that impact them (for example, increasing responsibilities, changes in technology, loss of personnel, or increase in customers supported).
Understand impacts on team performance: Remain engaged with your teams so that you understand how they are doing and if there are external factors affecting them. When your teams are impacted by external factors, reevaluate goals and adjust targets as appropriate. Identify obstacles that are impeding your teams progress. Act on behalf of your teams to help address obstacles and remove unnecessary burdens.
Provide the resources necessary for teams to be successful: Regularly review if resources are still appropriate, of if additional resources are needed, and make appropriate adjustments to support teams.

08. Diverse opinions are encouraged and sought within and across teams	Manual	Low	Open	

Recommendation:
Seek diverse opinions and perspectives: Don't make the women take the notes. Hire more diverse candidate Make resources Safe space make people welcome
Expand roles and responsibilities: Provide opportunity for team members to take on roles that they might not otherwise. They will gain experience and perspective from the role, and from interactions with new team members with whom they might not otherwise interact. They will bring their experience and perspective to the new role and team members they interact with. As perspective increases additional business opportunities may emerge, or new opportunities for improvement may be identified. Have members within a team take turns at common tasks that others typically perform to understand the demands and impact of performing them.
Provide a safe and welcoming environment: Have policy and controls that protect team members' mental and physical safety within your organization. Team members should be able to interact without fear of reprisal. When team members feel safe and welcome they are more likely to be engaged and productive. The more diverse your organization the better your understanding can be of the people you support including your customers. When your team members are comfortable, feel free to speak, and are confident they will be heard, they are more likely to share valuable insights (for example, marketing opportunities, accessibility needs, unserved market segments, unacknowledged risks in your environment).
Enable team members to participate fully: Provide the resources necessary for your employees to participate fully in all work related activities. Team members that face daily challenges have developed skills for working around them. These uniquely developed skills can provide significant benefit to your organization. Supporting team members with necessary accommodations will increase the benefits you can receive from their contributions.

Pillar Name: Operational Excellence

Question: 04. How do you design your workload so that you can understand its state?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Implement application telemetry	Automated	High	Open	7

Recommendation:
Implement log and metric telemetry: Instrument your application code to emit information about their internal state, status, and the achievement of business outcomes. Use this information to determine when a response is required.
Implement application telemetry: Design your application code to emit information about its internal state, status, and achievement of business outcomes (for example, queue depth, error messages, and response times).

02. Implement and configure workload telemetry	Automated	High	Open	41

Recommendation:
Implement log and metric telemetry: Instrument your workload to emit information about its internal state, status, and the achievement of business outcomes. Use this information to determine when a response is required.
Implement and configure workload telemetry: Design and configure your workload to emit information about its internal state and current status (for example, API call volume, HTTP status codes, and scaling events).

03. Implement user activity telemetry	Manual	Medium	Open	

Recommendation:
Implement user activity telemetry: Design your application code to emit information about user activity (for example, click streams, or started, abandoned, and completed transactions). Use this information to help understand how the application is used, patterns of usage, and to determine when a response is required.

04. Implement dependency telemetry	Automated	Medium	Verified	
05. Implement transaction traceability	Automated	Low	Open	7

Recommendation:
Implement transaction traceability: Design your application and workload to emit information about the flow of transactions across system components, such as transaction stage, active component, and time to complete activity. Use this information to determine what is in progress, what is complete, and what the results of completed activities are. This helps you determine when a response is required. For example, longer than expected transaction response times within a component can indicate issues with that component.

Pillar Name: Operational Excellence

Question: 05. How do you reduce defects, ease remediation, and improve flow into production?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Use version control	Manual	High	Open	

Recommendation:
Use version control: Maintain assets in version controlled repositories. Doing so supports tracking changes, deploying new versions, detecting changes to existing versions, and reverting to prior versions (for example, rolling back to a known good state in the event of a failure). Integrate the version control capabilities of your configuration management systems into your procedures.

02. Test and validate changes	Manual	High	Open	

Pillar Name: Operational Excellence

Question: 05. How do you reduce defects, ease remediation, and improve flow into production?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Recommendation: Test and validate changes: Changes should be tested and the results validated at all lifecycle stages (for example, development, test, and production). Use testing results to confirm new features and mitigate the risk and impact of failed deployments. Automate testing and validation to ensure consistency of review, to reduce errors caused by manual processes, and reduce the level of effort.				
03. Use configuration management systems	Manual	Medium	Open	
Recommendation: Use configuration management systems to track and implement changes, to reduce errors caused by manual processes, and reduce the level of effort.				
04. Use build and deployment management systems	Manual	Medium	Open	
Recommendation: Use build and deployment management systems: Use build and deployment management systems to track and implement change, to reduce errors caused by manual processes, and reduce the level of effort. Fully automate the integration and deployment pipeline from code check in through build, testing, deployment, and validation. This reduces lead time, enables increased frequency of change, and reduces the level of effort.				
05. Perform patch management	Automated	Medium	Verified	
06. Share design standards	Manual	Medium	Open	
Recommendation: Share design standards: Share existing best practices, design standards, checklists, operating procedures, and guidance and governance requirements across teams to reduce complexity and maximize the benefits from development efforts. Ensure that procedures exist to request changes, additions, and exceptions to design standards to support continual improvement and innovation. Ensure that teams are aware of published content so that they can take advantage of content, and limit rework and wasted effort.				
07. Implement practices to improve code quality	Manual	Medium	Open	
Recommendation: Implement practices to improve code quality: Implement practices to improve code quality to minimize defects and the risk of their being deployed. For example, test driven development, pair programming, code reviews, and standards adoption.				
08. Use multiple environments	Manual	Medium	Open	
Recommendation: Use multiple environments: Provide developers sandbox environments with minimized controls to enable experimentation. Provide individual development environments to enable work in parallel, increasing development agility. Implement more rigorous controls in the environments approaching production to allow developers necessary freedom for innovation. Use infrastructure as code and configuration management systems to deploy environments that are configured consistent with the controls present in production to ensure systems operate as expected when deployed. When environments are not in use, turn them off to avoid costs associated with idle resources (for example, development systems on evenings and weekends). Deploy production equivalent environments when load testing to enable valid results.				
09. Make frequent, small, reversible changes	Manual	Low	Open	
Recommendation: Make frequent, small, reversible changes: Frequent, small, and reversible changes reduce the scope and impact of a change. This eases troubleshooting, enables faster remediation, and provides the option to roll back a change. It also increases the rate at which you can deliver value to the business.				
10. Fully automate integration and deployment	Manual	Low	Open	
Recommendation:				

Pillar Name: Operational Excellence				
Question: 05. How do you reduce defects, ease remediation, and improve flow into production?				
Question Risk: High				
Best Practice	**Nature**	**Severity**	**Status**	**Violated Resources**
Use build and deployment management systems: Use build and deployment management systems to track and implement change, to reduce errors caused by manual processes, and reduce the level of effort. Fully automate the integration and deployment pipeline from code check in through build, testing, deployment, and validation. This reduces lead time, enables increased frequency of change, and reduces the level of effort.				

Pillar Name: Operational Excellence				
Question: 06. How do you mitigate deployment risks?				
Question Risk: High				
Best Practice	**Nature**	**Severity**	**Status**	**Violated Resources**
01. Plan for unsuccessful changes	Manual	High	Open	
Recommendation: Plan for unsuccessful changes: Plan to revert to a known good state (that is, roll back the change), or remediate in the production environment (that is, roll forward the change) if a change does not have the desired outcome. When you identify changes that you cannot roll back if unsuccessful, apply due diligence prior to committing the change.				
02. Test and validate changes	Manual	High	Open	
Recommendation: Test and validate changes: Test changes and validate the results at all lifecycle stages (for example, development, test, and production), to confirm new features and minimize the risk and impact of failed deployments.				
03. Use deployment management systems	Manual	Medium	Open	
Recommendation: Use deployment management systems: Use deployment management systems to track and implement change. This will reduce errors caused by manual processes, and reduce the level of effort to deploy changes. Automate the integration and deployment pipeline from code check in through testing, deployment, and validation. This reduces lead time, enables increased frequency of change, and further reduces the level of effort.				
04. Test using limited deployments	Manual	Medium	Open	
Recommendation: Test using limited deployments: Test with limited deployments alongside existing systems to confirm desired outcomes prior to full scale deployment. For example, use deployment canary testing or one box deployments.				
05. Deploy using parallel environments	Manual	Medium	Open	
Recommendation: Deploy using parallel environments: Implement changes onto parallel environments, and transition or cut over to the new environment. Maintain the prior environment until there is confirmation of successful deployment. This minimizes recovery time by enabling rollback to the previous environment. For example, use immutable infrastructures with blue/green deployments.				
06. Deploy frequent, small, reversible changes	Manual	Low	Open	
Recommendation: Deploy frequent, small, reversible changes: Use frequent, small, and reversible changes to reduce the scope of a change. This results in easier troubleshooting and faster remediation with the option to roll back a change.				
07. Fully automate integration and deployment	Manual	Low	Open	

Pillar Name: Operational Excellence

Question: 06. How do you mitigate deployment risks?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources

Recommendation:
Use build and deployment management systems: Use build and deployment management systems to track and implement change, to reduce errors caused by manual processes, and reduce the level of effort. Fully automate the integration and deployment pipeline from code check in through build, testing, deployment, and validation. This reduces lead time, enables increased frequency of change, and reduces the level of effort.

Best Practice	Nature	Severity	Status	Violated Resources
08. Automate testing and rollback	Manual	Low	Open	

Recommendation:
Automate testing and rollback: Automate testing of deployed environments to confirm desired outcomes. Automate rollback to previous known good state when outcomes are not achieved to minimize recovery time and reduce errors caused by manual processes. For example, perform detailed synthetic user transactions following deployment, verify the results, and roll back on failure.

Pillar Name: Operational Excellence

Question: 07. How do you know that you are ready to support a workload?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Ensure personnel capability	Manual	High	Open	

Recommendation:
Personnel capability: Validate that there are sufficient trained personnel to effectively support the workload.
Team size: Ensure that you have enough team members to cover operational activities, including on call duties.
Team skill: Ensure that your team members have sufficient training on AWS, your workload, and your operations tools to perform their duties.
Review capabilities: Review team size and skill as operating conditions and workloads change, to ensure there is sufficient capability to maintain operational excellence. Make adjustments to ensure that team size and skill match the operational requirements for the workloads that the team supports.

Best Practice	Nature	Severity	Status	Violated Resources
02. Ensure a consistent review of operational readiness	Manual	High	Open	

Recommendation:
Ensure consistent review of operational readiness: Ensure you have a consistent review of your readiness to operate a workload. Review must include at a minimum the operational readiness of the teams and the workload, and security considerations. Review elements can be hard requirements or you can make a risk based decision to operate a workload that does not satisfy all requirements. Review elements can be specific to a workload, architecture, or can be implementation dependent. Implement reviews as code and trigger reviews in response to events where appropriate, to ensure consistency, speed of execution, and reduce errors caused by manual processes.
Create checklists: Ensure you have a consistent review of your readiness to operate a workload. Create operational readiness checklists and validate them against your business, development, operations, and governance requirements. Ensure they address: governance, best practices, configuration standards, restoration procedures, monitoring, maintenance procedures, IT operations procedures, and staffing.
Use checklists: Make checklists accessible to developers so that they can develop to the appropriate standards. Evaluate checklists when moving between lifecycle stages and environments so that you can identify issues early, when the level of effort to remediate issues is lower. Use the results of checklists to make informed decisions about benefits and risks when considering promoting changes between environments.
Implement checklists as code and trigger checklist execution in response to events: Implement checklists as code and trigger checklist execution in response to events where possible, to enhance speed, ensure consistency, and reduce errors caused by manual processes. Integrate automated checklist execution into deployment pipelines.

Best Practice	Nature	Severity	Status	Violated Resources
03. Use runbooks to perform procedures	Manual	Medium	Open	

Recommendation:
Use runbooks to perform standard procedures: Runbooks are documented procedures to achieve specific outcomes. Enable consistent and prompt responses to well understood events by documenting procedures in runbooks. Runbooks must contain the minimum information for an adequately skilled person to achieve the desired outcome. For example, required permissions, required tools, constraints on performing the procedure

Pillar Name: Operational Excellence				
Question: 07. How do you know that you are ready to support a workload?				
Question Risk: High				
Best Practice	**Nature**	**Severity**	**Status**	**Violated Resources**
(for example, specific maintenance windows), and execution steps. Implement runbooks as code: Perform your operations as code by implementing your runbooks as code to ensure consistency and reduce errors caused by manual processes Trigger runbooks in response to events: Trigger the execution of runbook code in response to observed events when appropriate. This increases the speed of the response and reduces the level of effort to respond.				
04. Use playbooks to investigate issues	Manual	Medium	Open	
Recommendation: Use playbooks to identify issues: Playbooks are documented processes to investigate issues. Enable consistent and prompt responses to failure scenarios by documenting processes in playbooks. Playbooks must contain the information and guidance necessary for an adequately skilled person to gather applicable information, identify potential sources of failure, isolate faults, and determine contributing factors (i.e. perform root cause analysis). Implement playbooks as code: Perform your operations as code by scripting your playbooks to ensure consistency and limit reduce errors caused by manual processes. Playbooks can be composed of multiple scripts representing the different steps that might be necessary to identify the contributing factors to an issue. Runbook activities can be triggered or performed as part of playbook activities, or may prompt for execution of a playbook in response to identified events.				
05. Make informed decisions to deploy systems and changes	Manual	Low	Open	
Recommendation: Make informed decisions to deploy workloads and changes: Evaluate the capabilities of the team to support the workload and the workload's compliance with governance. Evaluate these against the benefits of deployment when determining whether to transition a system or change into production. Understand the benefits and risks, and make informed decisions.				
06. Create support plans for production workloads	Manual	Low	Open	
Recommendation: Make informed decisions to deploy workloads and changes: Evaluate the capabilities of the team to support the workload and the workload's compliance with governance. Evaluate these against the benefits of deployment when determining whether to transition a system or change into production. Understand the benefits and risks, and make informed decisions.				

Pillar Name: Operational Excellence				
Question: 08. How do you understand the health of your workload?				
Question Risk: High				
Best Practice	**Nature**	**Severity**	**Status**	**Violated Resources**
01. Identify key performance indicators	Manual	High	Open	
Recommendation: Identify key performance indicators: Identify key performance indicators (KPIs) based on desired business and customer outcomes. Evaluate KPIs to determine workload success.				
02. Define workload metrics	Automated	High	Open	
Recommendation: Define workload metrics: Define workload metrics to measure the achievement of KPIs. Define workload metrics to measure the health of the workload and its individual components. Evaluate metrics to determine if the workload is achieving desired outcomes, and to understand the health of the workload.				
03. Collect and analyze workload metrics	Automated	High	Open	34

Pillar Name: Operational Excellence

Question: 08. How do you understand the health of your workload?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Recommendation: Collect and analyze workload metrics: Perform regular proactive reviews of metrics to identify trends and determine where appropriate responses are needed.				
04. Establish workload metrics baselines	Manual	Medium	Open	
Recommendation: Establish baselines for workload metrics : Establish baselines for workload metrics to provide expected values as the basis for comparison.				
05. Learn expected patterns of activity for workload	Manual	Medium	Open	
Recommendation: Learn expected patterns of activity for workload: Establish patterns of workload activity to determine when behavior is outside of the expected values so that you can respond appropriately if required.				
06. Alert when workload outcomes are at risk	Automated	Medium	Verified	
07. Alert when workload anomalies are detected	Automated	Low	Verified	
08. KPIs and metrics	Manual	Low	Open	
Recommendation: Validate the achievement of outcomes and the effectiveness of KPIs and metrics : Create a business level view of your workload operations to help you determine if you are satisfying needs and to identify areas that need improvement to reach business goals. Validate the effectiveness of KPIs and metrics and revise them if necessary.				

Pillar Name: Operational Excellence

Question: 09. How do you understand the health of your operations?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Identify key performance indicators	Manual	High	Open	
Recommendation: Identify key performance indicators: Identify key performance indicators (KPIs) based on desired business and customer outcomes. Evaluate KPIs to determine operations success.				
02. Define operations metrics	Manual	High	Open	
Recommendation: Define operations metrics: Define operations metrics to measure the achievement of KPIs. Define operations metrics to measure the health of operations and its activities. Evaluate metrics to determine if operations are achieving desired outcomes, and to understand the health of the operations.				
03. Collect and analyze operations metrics	Automated	High	Open	34
Recommendation: Collect and analyze operations metrics: Perform regular proactive reviews of metrics to identify trends and determine where appropriate responses are needed.				

Pillar Name: Operational Excellence

Question: 09. How do you understand the health of your operations?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
04. Establish operations metrics baselines	Manual	Medium	Open	
Recommendation: Learn expected patterns of activity for operations: Establish patterns of operations activity to determine when behavior is outside of the expected values so that you can respond appropriately if required.				
05. Learn the expected patterns of activity for operations	Manual	Medium	Open	
Recommendation: Learn expected patterns of activity for operations: Establish patterns of operations activity to determine when behavior is outside of the expected values so that you can respond appropriately if required.				
06. Alert when operations outcomes are at risk	Automated	Medium	Open	34
Recommendation: Alert when operations outcomes are at risk: Raise an alert when operations outcomes are at risk so that you can respond appropriately if required.				
07. Alert when operations anomalies are detected	Automated	Low	Open	34
Recommendation: Alert when operations anomalies are detected: Raise an alert when operations anomalies are detected so that you can respond appropriately if required.				
08. KPIs and metrics	Automated	Low	Open	34
Recommendation: Validate the achievement of outcomes and the effectiveness of KPIs and metrics : Create a business level view of your operations activities to help you determine if you are satisfying needs and to identify areas that need improvement to reach business goals. Validate the effectiveness of KPIs and metrics and revise them if necessary.				

Pillar Name: Operational Excellence

Question: 10. How do you manage workload and operations events?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Use a process for event, incident, and problem management	Automated	High	Verified	
02. Have a process per alert	Manual	High	Open	
Recommendation: Process per alert: Any event for which you raise an alert should have a well defined response (runbook or playbook) with a specifically identified owner (for example, individual, team, or role) accountable for successful execution. Performance of the response may be automated or conducted by another team but the owner is accountable for ensuring the process delivers the expected outcomes. By having these processes, you ensure effective and prompt responses to operations events and you can prevent actionable events from being obscured by less valuable notifications. For example, automatic scaling might be applied to scale a web front end, but the operations team might be accountable to ensure that the automatic scaling rules and limits are appropriate for workload needs.				
03. Prioritize operational events based on business impact	Manual	Medium	Open	

Pillar Name: Operational Excellence

Question: 10. How do you manage workload and operations events?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Recommendation: Prioritize operational events based on business impact: Ensure that when multiple events require intervention, those that are most significant to the business are addressed first. For example, impacts can include loss of life or injury, financial loss, regulatory violations, or damage to reputation or trust.				
04. Define escalation paths	Manual	Medium	Open	
Recommendation: Define escalation paths: Define escalation paths in your runbooks and playbooks, including what triggers escalation, and procedures for escalation. For example, escalation of an issue from support engineers to senior support engineers when runbooks cannot resolve the issue, or when a predefined period of time has elapsed. Another example of an appropriate escalation path is from senior support engineers to the development team for a workload when the playbooks are unable to identify a path to remediation, or when a predefined period of time has elapsed. Specifically identify owners for each action to ensure effective and prompt responses to operations events. Escalations can include third parties. For example, a network connectivity provider or a software vendor. Escalations can include identified authorized decision makers for impacted systems.				
05. Define a customer communication plan for outages	Manual	Medium	Open	
Recommendation: Enable push notifications: Communicate directly with your users (for example, with email or SMS) when the services they use are impacted, and when the services return to normal operating conditions, to enable users to take appropriate action.				
06. Communicate status through dashboards	Manual	Medium	Open	
Recommendation: Communicate status through dashboards: Provide dashboards tailored to their target audiences (for example, internal technical teams, leadership, and customers) to communicate the current operating status of the business and provide metrics of interest. Providing a self service option for status information reduces the disruption of fielding requests for status by the operations team. Examples include Amazon CloudWatch dashboards, and AWS Personal Health Dashboard.				
07. Automate responses to events	Manual	Low	Open	
Recommendation: Automate responses to events: Automate responses to events to reduce errors caused by manual processes, and to ensure prompt and consistent responses.				

Pillar Name: Operational Excellence

Question: 11. How do you evolve operations?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Have a process for continuous improvement	Manual	High	Open	
Recommendation: Define processes for continuous improvement: Regularly evaluate and prioritize opportunities for improvement to focus efforts where they provide the greatest benefits. Implement changes to improve and evaluate the outcomes to determine success. If the outcomes do not satisfy the goals, and the improvement is still a priority, iterate using alternative courses of action. Your operations processes should include dedicated time and resources to make continuous incremental improvements possible.				

Pillar Name: Operational Excellence

Question: 11. How do you evolve operations?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
02. Perform post-incident analysis	Manual	High	Open	

Recommendation:
Use a process to determine contributing factors: Review all customer impacting incidents. Have a process to identify and document the contributing factors of an incident so that you can develop mitigations to limit or prevent recurrence and you can develop procedures for prompt and effective responses. Communicate root cause as appropriate, tailored to target audiences.

03. Implement feedback loops	Manual	High	Open	

Recommendation:
Feedback loops: Have procedures embedded in your operations activities to capture feedback from their execution and to identify areas for improvement.
Immediate feedback: Immediate feedback comes from the execution of operations activities where, through review of the execution and outcomes, it is recognized that the process could be improved. Feedback can come from customers, team members, or automated output of an activity. When the improvement has a low level of effort, or significant benefit, consider implementing it immediately. Track opportunities for improvement in your backlog or issue system as appropriate. For example, a process where data is staged on an intermediate device could be optimized by instead placing the data directly into the target environment. This would eliminate a step in the process and the requirement for the intermediate resources.
Retrospective analysis: Perform retrospective analysis regularly to capture feedback from the review of operational outcomes and metrics over time. Use trends to identify areas that need improvement. For example, review the rate of deployment failures to identify when potential issues with development and deployment activities have emerged.

04. Perform knowledge management	Manual	High	Open	

Recommendation:
Knowledge Management: Ensure mechanisms exist for your team members to discover the information that they are looking for in a timely manner, access it, and identify that its current and complete. Maintain mechanisms to identify needed content, content in need of refresh, and content that should be archived so that it is no longer referenced.

05. Define drivers for improvement	Manual	Medium	Open	

Recommendation:
Understand drivers for improvement: You should only make changes to a system when a desired outcome is supported.
Desired capabilities: Evaluate desired features and capabilities when evaluating opportunities for improvement.
Unacceptable issues: Evaluate unacceptable issues, bugs, and vulnerabilities when evaluating opportunities for improvement.
Compliance requirements: Evaluate updates and changes required to maintain compliance with regulation, policy, or to remain under support from a third party, when reviewing opportunities for improvement.

06. Validate insights	Manual	Medium	Open	

Recommendation:
Validate Insights: Engage with business owners and subject matter experts to ensure there is common understanding and agreement of the meaning of the data you have collected. Identify additional concerns, potential impacts, and determine a courses of action.

07. Perform operations metrics reviews	Manual	Medium	Open	

Recommendation:
Operations metrics reviews: Regularly perform retrospective analysis of operations metrics with cross team participants from different areas of the business. Engage stakeholders, including the business, development, and operations teams, to validate your findings from immediate feedback and retrospective analysis, and to share lessons learned. Use their insights to identify opportunities for improvement and potential courses of action.

08. Document and share lessons learned	Manual	Low	Open	

Recommendation:
Document and share lessons learned: Have procedures to document the lessons learned from the execution of operations activities and retrospective analysis so that they can be used by other teams.

Pillar Name: Operational Excellence

Question: 11. How do you evolve operations?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Share learnings: Have procedures to share lessons learned and associated artifacts across teams. For example, share updated procedures, guidance, governance, and best practices through an accessible wiki; share scripts, code, and libraries through a common repository.				
09. Allocate time to make improvements	Manual	Low	Open	

Recommendation:
Allocate time to make improvements: Dedicate time and resources within your processes to make continuous incremental improvements possible. Implement changes to improve and evaluate the results to determine success. If the results do not satisfy the goals, and the improvement is still a priority, pursue alternative courses of action.

Pillar Name: Performance Efficiency

Question: 01. How do you select the best performing architecture?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Understand the available services and resources	Manual	High	Open	

Recommendation:
Inventory your workload software and architecture for related services: Gather an inventory of your workload and decide which category of products to learn more about.Gather an Identify workload components that can be replaced with managed services to increase performance and reduce operational complexity.

02. Define a process for architectural choices	Manual	High	Open	

Recommendation:
Select an architectural approach: Identify the kind of architecture that meets your performance requirements.Identify constraints, such as the media for delivery (desktop, web, mobile, IoT), legacy requirements, and integrations.Identify opportunities for reuse, including refactoring. Consult other teams, architecture diagrams, and resources such as AWS Solution Architects, AWS Reference Architectures, and APN Partners to help you choose an architecture.
Define performance requirements: Use the customer experience to identify the most important metrics. For each metric, identify the target, measurement approach, and priority.Define the customer experience. Document the performance experience required by customers, including how customers will judge the performance of the workload.Prioritize experience concerns for critical user stories. Include performance requirements and implement scripted user journeys to ensure that you know how the stories perform against your requirements.

03. Factor cost requirements into decisions	Automated	High	Verified	
04. Use policies or reference architectures	Manual	Medium	Open	

Recommendation:
Deploy your workload using existing policies or reference architectures: Integrate the services into your cloud deployment, then use your performance tests to ensure that you can continue to meet your performance requirements.

05. Use guidance from your cloud provider or an appropriate partner	Automated	Medium	Verified	
06. Benchmark existing workloads	Automated	Medium	Verified	

Pillar Name: Performance Efficiency

Question: 01. How do you select the best performing architecture?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
07. Load test your workload	Manual	Low	Open	

Recommendation:
Validate your approach with load testing: Load test a proof of concept to find out if you meet your performance requirements. You can use AWS services to run production scale environments to test your architecture. Because you only pay for the test environment when it is needed, you can carry out full scale testing at a fraction of the cost of using an on premises environment. Amazon EC2 testing policy
Monitor metrics: Amazon CloudWatch can collect metrics across the resources in your architecture. You can also collect and publish custom metrics to surface business or derived metrics. Use CloudWatch or third party solutions to set alarms that indicate when thresholds are breached.
Test at scale: Load testing uses your actual workload so you can see how your solution performs in a production environment. You can use AWS services to run production scale environments to test your architecture. Because you only pay for the test environment when it is needed, you can run full scale testing at a lower cost than using an on premises environment. Take advantage of the AWS Cloud to test your workload to discover where it fails to scale, or if it scales in a non linear way. For example, use Spot Instances to generate loads at low cost and discover bottlenecks before they are experienced in production.

Pillar Name: Performance Efficiency

Question: 02. How do you select your compute solution?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Evaluate the available compute options	Automated	High	Verified	
02. Understand the available compute configuration options	Automated	High	Verified	
03. Collect compute-related metrics	Automated	High	Open	7

Recommendation:
Collect compute related metrics: Amazon CloudWatch can collect metrics across the compute resources in your environment. Use a combination of CloudWatch and other metrics recording tools to track the system level metrics within your workload. Record data such as CPU usage levels, memory, disk I/O, and network to gain insight into utilization levels or bottlenecks. This data is crucial to understand how the workload is performing and how effectively it is using resources. Use these metrics as part of a data driven approach to actively tune and optimize your workload's resources.

04. Determine the required configuration by right-sizing	Automated	Medium	Verified	
05. Use the available elasticity of resources	Automated	Medium	Open	11

Recommendation:
Take advantage of elasticity: Elasticity matches the supply of resources you have against the demand for those resources. Instances, containers, and functions provide mechanisms for elasticity either in combination with automatic scaling or as a feature of the service. Use elasticity in your architecture to ensure that you have sufficient capacity to meet performance requirements at all scales of use. Ensure that the metrics for scaling up or down elastic resources are validated against the type of workload being deployed. If you are deploying a video transcoding application, 100% CPU is expected and should not be your primary metric. Alternatively, you can measure against the queue depth of transcoding jobs waiting to scale your instance types. Ensure that workload deployments can handle both scale up and scale down events. Scaling down workload components safely is as critical as scaling up resources when demand dictates. Create test scenarios for scale down events to ensure that the workload behaves as expected.

06. Continually evaluate compute needs based on metrics	Automated	Low	Verified	

Pillar Name: Performance Efficiency

Question: 03. How do you select your storage solution?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Understand storage characteristics and requirements	Manual	High	Open	

Recommendation:
Define storage performance requirements: Identify your workload most important storage performance metrics and implement improvements as part of a data driven approach, using benchmarking or load testing. Use this data to identify where your storage solution is constrained, and examine configuration options to improve the solution. Determine the expected growth rate for your workload and choose a storage solution that will meet those rates. Object and file storage solutions, such as Amazon S3 and Amazon Elastic File System, enable unlimited storage.

02. Evaluate available configuration options	Manual	Medium	Open	

Recommendation:
Determine storage characteristics: When you evaluate a storage solution, determine which storage characteristics you require, such as ability to share, file size, cache size, latency, throughput, and persistence of data. Then match your requirements to the AWS service that best fits your needs.

03. Make decisions based on access patterns and metrics	Manual	Low	Open	

Recommendation:
Optimize your storage usage and access patterns: Choose storage systems based on your workload's access patterns and the characteristics of the available storage options. Determine the best place to store data that will enable you to meet your requirements while reducing overhead. Use performance optimizations and access patterns when configuring and interacting with data based on the characteristics of your storage (for example, striping volumes or partitioning data).
Select appropriate metrics for storage options: Ensure that you select the appropriate storage metrics for the workload. Each storage option offers various metrics to track how your workload performs over time. Ensure that you are measuring against any storage burst metrics (for example, monitoring burst credits for Amazon EFS). For storage systems that are fixed sized, such as Amazon Elastic Block Store or Amazon FSx, ensure that you are monitoring the amount of storage used versus the overall storage size. Create automation when possible to increase the storage size when reaching a threshold.
Monitor metrics: Amazon CloudWatch can collect metrics across the resources in your architecture. You can also collect and publish custom metrics to surface business or derived metrics. Use CloudWatch or third party solutions to set alarms that indicate when thresholds are breached.

Pillar Name: Performance Efficiency

Question: 04. How do you select your database solution?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Understand data characteristics	Manual	High	Open	

Recommendation:
Research and document data characteristics: Before choosing a database solution, understand the functional requirements of your workload and how it interacts with data. When evaluating a database solution, determine if it is best suited to meet your requirements (for example, transactions or high availability) so that you can select the best combination of databases to use for your workload. Evaluate alternative databases that could better meet your workload requirements. For example, if you are building an IoT application it may be better to select a timeseries database, such as Amazon Timestream, to easily store and analyze trillions of events per day at 1/10th the cost of relational databases.

02. Evaluate the available options	Manual	High	Open	

Recommendation:
Select the appropriate database type for your workload: AWS allows you to choose from multiple purpose built database engines including relational, key value, document, in memory, graph, time series, and ledger databases. The AWS portfolio of purpose built databases supports diverse data models and allows you to build use case driven, highly scalable, distributed applications. By picking the best database to solve a specific

Pillar Name: Performance Efficiency

Question: 04. How do you select your database solution?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
problem or a group of problems, you can break away from restrictive one size fits all monolithic databases and focus on building applications to meet the needs of your business. Define database performance requirements: Identify the database performance metrics that matter for your workload, and implement the requirements as part of a data driven approach, using benchmarking or load testing. Use this data to identify where your database solution is constrained, and examine configuration options to solve the issue. Enable database caching options: Evaluate database caching options, such as Amazon ElastiCache for Redis for caching relational database or Amazon DynamoDB Accelerator (DAX) for a fully managed, highly available, in memory cache for DynamoDB. These options can deliver improved performance, in some cases from milliseconds to microseconds even at millions of requests per second.				
03. Collect and record database performance metrics	Manual	High	Open	
Recommendation: Collect database related metrics: Design your workload to record metrics related to database activity. This data is crucial for understanding how your database systems are impacting the overall performance of your workload and where you can make changes to improve performance and efficiency. For example, tracking data points such as query times, the number of transactions, disk usage, index usage, or slow queries, enables you to optimize your database systems. Monitor metrics: Amazon CloudWatch can collect metrics across the resources in your architecture. You can also collect and publish custom metrics to surface business or derived metrics. Use CloudWatch or third party solutions to set alarms that indicate when thresholds are breached.				
04. Choose data storage based on access patterns	Manual	Medium	Open	
Recommendation: Use access patterns to determine data storage: Evaluate your workload access patterns to find an appropriate data storage pattern. For example, if your workload requires ad hoc query access, you may select a relational database such as Amazon RDS. If your workload is driven by a high growth rate or high traffic events, you should select a key value database, such as Amazon DynamoDB.				
05. Optimize data storage based on access patterns and metrics	Manual	Low	Open	
Recommendation: Optimize data storage based on metrics and patterns: Use reported metrics to identify any underperforming areas in your workload and optimize your database components. Each database system has different performance related characteristics to evaluate, such as how data is indexed, cached, or distributed among multiple systems. Measure the impact of your optimizations.				

Pillar Name: Performance Efficiency

Question: 05. How do you configure your networking solution?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Understand how networking impacts performance	Manual	High	Open	
Recommendation: Define networking performance requirements: Identify your workload important networking performance metrics. Implement requirements as part of a data driven approach, using benchmarking or load testing. Use this data to identify where your network solution is constrained, and examine configuration options that could improve the solution. Measure the network impact on your workload: Analyze your workload's networking requirements to understand how network performance impacts overall performance				
02. Evaluate available networking features	Automated	High	Verified	
03. Choose appropriately sized dedicated connectivity or VPN for hybrid workloads	Manual	High	Open	
Recommendation:				

Pillar Name: Performance Efficiency

Question: 05. How do you configure your networking solution?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Develop a hybrid networking architecture based on your bandwidth requirements: Estimate the bandwidth and latency requirements of your hybrid applications. Based on your bandwidth requirements, a single VPN or Direct Connect connection might not be enough, and you must architect a hybrid setup to enable traffic load balancing across multiple connections. Direct Connect may be required. It offers a more predictable and consistence performance because it does not involve internet. It is great for production workloads that require consistent latency and almost zero jitter.				
04. Leverage load-balancing and encryption offloading	Automated	High	Verified	
05. Choose network protocols to improve performance	Manual	Medium	Open	
Recommendation: Optimize network traffic: Select the appropriate protocol to optimize the performance of your workload. There is a relationship between latency and bandwidth to achieve throughput. If your file transfer is using TCP, higher latencies reduce overall throughput. There are approaches to fix latency with TCP tuning and optimized transfer protocols, some which use UDP.				
06. Choose your workload’s location based on network requirements	Manual	Medium	Open	
Recommendation: Reduce latency by selecting the correct locations: Identify where your users and data are located. Take advantage of AWS Regions, Availability Zones, placement groups, and edge locations to reduce latency.				
07. Optimize network configuration based on metrics	Automated	Medium	Open	17
Recommendation: Enable VPC Flow Logs: VPC Flow Logs enable you to capture information about the IP traffic going to and from network interfaces in your VPC. VPC Flow Logs help you with a number of tasks, such as troubleshooting why specific traffic is not reaching an instance, which can help you diagnose overly restrictive security group rules. You can use flow logs as a security tool to monitor the traffic that is reaching your instance, to profile your network traffic, and to look for abnormal traffic behaviors. Enable appropriate metrics for network options: Ensure that you select the appropriate network metrics for your workload. You can enable metrics for VPC NAT gateway, transit gateways, and VPN tunnels.				

Pillar Name: Performance Efficiency

Question: 06. How do you evolve your workload to take advantage of new releases?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Stay up-to-date on new resources and services	Manual	High	Open	
Recommendation: Evaluate updates and releases: Define a process to evaluate updates, new features, and services from AWS. For example, building proof of concepts that use new technologies. When trying new ideas or services, run performance tests to measure the impact on the efficiency or performance of the workload. Take advantage of the flexibility that you have in AWS to test new ideas or technologies frequently with minimal cost or risk.				
02. Define a process to improve workload performance	Manual	Medium	Open	
Recommendation: Identify the key performance constraints for your workload: Document your workload performance constraints so that you know what kinds of innovation might improve the performance of your workload.				
03. Evolve workload performance over time	Automated	Low	Open	
Recommendation:				

Pillar Name: Performance Efficiency				
Question: 06. How do you evolve your workload to take advantage of new releases?				
Question Risk: High				
Best Practice	**Nature**	**Severity**	**Status**	**Violated Resources**
Evolve your workload over time: Use the information you gather when evaluating new services or technologies to drive change. As your business or workload changes, performance needs also change. Use data gathered from your workload metrics to evaluate areas where you can achieve the biggest gains in efficiency or performance, and proactively adopt new services and technologies to keep up with demand.				

Pillar Name: Performance Efficiency				
Question: 07. How do you monitor your resources to ensure they are performing?				
Question Risk: High				
Best Practice	**Nature**	**Severity**	**Status**	**Violated Resources**
01. Record performance-related metrics	Automated	High	Open	7
Recommendation: Record performance data: Identify the relevant performance metrics for your workload and record them. This data helps identify which components are impacting overall performance or efficiency of your workload. Identify performance metrics: Use the customer experience to identify the most important metrics. For each metric, identify the target, measurement approach, and priority. Use these data points to build alarms and notifications to proactively address performance related issues.				
02. Analyze metrics when events or incidents occur	Automated	High	Open	7
Recommendation: Prioritize experience concerns for critical user stories: When you write critical user stories for your architecture, include performance requirements, such as specifying how quickly each critical story should execute. For these critical stories, implement additional scripted user journeys to ensure that you know how the user stories perform against your requirements.				
03. KPIs) to measure workload performance	Manual	High	Open	
Recommendation: Define the customer experience: Document the performance experience required by customers, including how customers judge the performance of the workload. Use these requirements to establish your KPIs, which indicate how the system is performing overall. Test user journeys: Use synthetic or sanitized versions of production data (remove sensitive or identifying information) for load testing. Exercise your entire architecture by using replayed or pre programmed user journeys through your application at scale.				
04. Use monitoring to generate alarm-based notifications	Automated	Medium	Open	7
Recommendation: Monitor metrics: Amazon CloudWatch can collect metrics across the resources in your architecture. You can collect and publish custom metrics to surface business or derived metrics. Use CloudWatch or a third party monitoring service to set alarms that indicate when thresholds are breached.				
05. Review metrics at regular intervals	Automated	Medium	Open	7
Recommendation: Constantly improve metric collection and monitoring: As part of responding to incidents or events, evaluate which metrics were helpful in addressing the issue and which metrics could have helped that are not currently being tracked. Use this method to improve the quality of metrics you collect so that you can prevent or more quickly resolve future incidents.				
06. Monitor and alarm proactively	Automated	Low	Open	7
Recommendation:				

Pillar Name: Performance Efficiency				
Question: 07. How do you monitor your resources to ensure they are performing?				
Question Risk: High				
Best Practice	**Nature**	**Severity**	**Status**	**Violated Resources**
Monitor performance during operations: Implement processes that provide visibility into performance as your workload is running. Build monitoring dashboards and establish a baseline for performance expectations.				

Pillar Name: Performance Efficiency				
Question: 08. How do you use tradeoffs to improve performance?				
Question Risk: High				
Best Practice	**Nature**	**Severity**	**Status**	**Violated Resources**
01. Understand the areas where performance is most critical	Manual	High	Open	
Recommendation: Identify constrained areas of the workload: Use load testing or monitoring to identify constrained areas (memory, CPU, custom metrics and other key performance indicators).				
02. Learn about design patterns and services	Manual	High	Open	
Recommendation: Understand the available product options: Learn which performance configuration options are available and how they could impact the workload. Optimizing the performance of your workload depends on understanding how these options interact with your architecture, and the impact they have on measured performance and user perceived performance. Evaluate design patterns from the Amazon Builders Library: The Amazon Builders Library provides a detailed information about how Amazon builds and operates technology. The free articles in the library are written by Amazon Builders senior engineers and cover topics across architecture, software delivery, and operations. For example, you can see how Amazon automates software delivery to achieve over 150 million deployments a year or how Amazon builder engineers implement principles such as shuffle sharding to build resilient systems that are highly available and fault tolerant.				
03. Identify how tradeoffs impact customers and efficiency	Manual	High	Open	
Recommendation: Identify tradeoffs: Use metrics and monitoring to identify areas of poor performance in your system. Determine how to make improvements, and how tradeoffs will impact the system and the user experience. For example, implementing caching data can help dramatically improve performance, but it requires a clear strategy for how and when to update or invalidate cached data to prevent incorrect system behavior.				
04. Measure the impact of performance improvements	Manual	Medium	Open	
Recommendation: Use a combination of strategies: A well architected system uses a combination of performance related strategies. Determine which strategy will have the largest positive impact on a given hotspot or bottleneck. For example, sharding data across multiple relational database systems could improve overall throughput while retaining support for transactions and, within each shard, caching can help to reduce the load.				
05. Use various performance-related strategies	Manual	Low	Open	
Recommendation: Use a data driven approach to evolve your architecture: As you make changes to the workload, collect and evaluate metrics to determine the impact of those changes. Measure the impacts to the system and to the end user to understand how your tradeoffs impact your workload. Use a systematic approach, such as load testing, to explore whether the tradeoff improves performance.				

Pillar Name: Reliability

Question: 01. How do you manage service quotas and constraints?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Aware of service quotas and constraints	Automated	High	Verified	
02. Manage service quotas across accounts and regions	Manual	High	Open	

Recommendation:
Identify service quotas across all relevant accounts, Regions, and Availability Zones: The limits are scoped to account and Region.
Select relevant accounts and Regions based on your service requirements, latency, regulatory, and disaster recovery (DR) requirements
Review AWS service quotas in the published documentation and Service Quotas Determine all the services your workload requires by looking at the deployment code Use AWS Config to find all AWS resources used in your AWS accounts You can also use your AWS CloudFormation to determine your AWS resources used. Look at the resources that were created either in the AWS console or via the list stack resources CLI command. You can also see resources configured to be deployed in the template itself. Determine the service quotas that apply. Use the programmatically accessible information via Trusted Advisor and Service Quotas.

03. Accommodate fixed service quotas and constraints through architecture	Manual	Medium	Open	

Recommendation:
Be aware of fixed service quotas: Be aware of fixed service quotas and constraints and architect around these.

04. Monitor and manage quotas	Manual	Medium	Open	

Recommendation:
Monitor and manage your quotas: Evaluate your potential usage on AWS, increase your regional service quotas appropriately, and allow planned growth in usage.
Capture current resource consumption (for example, buckets, instances, etc.): Use service API operations, such as the Amazon EC2 DescribeInstances API, to collect current resource consumption.
Capture your current quotas: Use AWS Service Quotas, AWS Trusted Advisor, and AWS documentation Use AWS Service Quotas, an AWS service that helps you manage your quotas for over 100 AWS services from one location Use Trusted Advisor service limits to determine your current service limits Use service API operations to determine current service quotas where supported Keep a record of quota increases that have been requested, and their status: After a quota increase has been approved, ensure that you update your records to reflect the change to the quota.

05. Automate quota management	Manual	Medium	Open	

Recommendation:
Set up automated monitoring: Implement tools using SDKs to alert you when thresholds are being approached.
Use Service Quotas and augment the service with an automated quota monitoring solution, such as AWS Limit Monitor or an offering from AWS Marketplace
Set up triggered responses based on quota thresholds, using Amazon SNS and AWS Service Quotas APIs Configure limit thresholds Integrate with change events from AWS Config, deployment pipelines, Amazon EventBridge, or third parties Set up triggers to take appropriate action on notifications and contact AWS Support when necessary
Test automation Artificially set low quota thresholds to test responses Manually trigger change events Run a game day to test the quota increase change process

06. Ensure that a sufficient gap exists between the current quotas and the maximum usage to accommodate failover	Manual	Medium	Open	

Recommendation:
Ensure that there is a sufficient gap between your service quota and your maximum usage to accommodate for a failover
Determine your service quotas, accounting for your deployment patterns, availability requirements, and consumption growth Determine your reliability requirements (also known as your "number of 9's") Establish your fault scenarios (for example, loss of a component, an Availability Zone, or a Region) Establish your deployment methodology (for example, Canary, Blue/Green, Red/Black, or rolling) Include an appropriate buffer (for example, 15%) to the current limit Plan consumption growth (for example, monitor your trends in consumption)
Request quota increases if necessary: Plan for necessary time for quota increase requests to be fulfilled

Pillar Name: Reliability

Question: 02. How do you plan your network topology?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Use highly available network connectivity for your workload public endpoints	Automated	High	Verified	
02. Provision redundant connectivity between private networks in the cloud and on-premises environments	Manual	High	Open	

Recommendation:
Ensure that you have highly available connectivity between AWS and on premises environment: Use multiple AWS Direct Connect (DX) connections or VPN tunnels between separately deployed private networks. Use multiple DX locations for high availability. If using multiple AWS Regions, ensure redundancy in at least two of them. You might want to evaluate AWS Marketplace appliances that terminate VPNs. If you use AWS Marketplace appliances, deploy redundant instances for high availability in different Availability Zones.
Ensure that you have a redundant connection to your on premises environment: You may need redundant connections to multiple AWS Regions to achieve your availability needs. Use service API operations to identify correct use of Direct Connect circuits If only one Direct Connect connection exists or you have none, set up redundant VPN tunnels to your virtual private gateways
Capture your current connectivity (for example, Direct Connect, virtual private gateways, AWS Marketplace appliances) Use service API operations to query configuration of Direct Connect connections Use service API operations to collect virtual private gateways where route tables use them Use service API operations to collect AWS Marketplace applications where route tables use them

03. Ensure IP subnet allocation accounts for expansion and availability	Automated	Medium	Verified	
04. Prefer hub-and-spoke topologies over many-to-many mesh	Manual	Medium	Open	

Recommendation:
Prefer hub and spoke topologies over many to many mesh: If more than two network address spaces (VPCs, on premises networks) are connected via VPC peering, AWS Direct Connect, or VPN, then use a hub and spoke model like that provided by AWS Transit Gateway.
For only two such networks, you can simply connect them to each other, but as the number of networks grows, the complexity of such meshed connections becomes untenable. AWS Transit Gateway provides an easy to maintain hub and spoke model, allowing routing of traffic across your multiple networks.

05. Enforce non-overlapping private IP address ranges in all private address spaces where they are connected	Manual	Medium	Open	

Recommendation:
Monitor and manage your CIDR use: Evaluate your potential usage on AWS, add CIDR ranges to existing VPCs, and create VPCs to allow planned growth in usage.
Capture current CIDR consumption (for example, VPCs, subnets, etc.) Use service API operations to collect current CIDR consumption
Capture your current subnet usage Use service API operations to collect subnets per VPC in each Region Record the current usage Determine if you created any overlapping IP ranges Calculate the spare capacity
Note overlapping IP ranges: You can either migrate to a new range of addresses or use Network and Port Translation (NAT) appliances from AWS Marketplace if you need to connect the overlapping ranges.

Pillar Name: Reliability

Question: 03. How do you design your workload service architecture?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Choose how to segment your workload	Manual	High	Open	

Recommendation:
Choose your architecture type based on how you will segment your workload: Choose a service oriented architecture (SOA) or microservices architecture (or in some cases a monolithic architecture).

Pillar Name: Reliability

Question: 03. How do you design your workload service architecture?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
SOA and microservices offer respectively smaller segmentation, which is preferred as a modern scalable and reliable architecture. However, if you have a good reason to choose a monolithic architecture, then you must ensure itss modular and has the ability to ultimately evolve to SOA or microservices as your product scales with user adoption. SOA can be a good compromise for achieving smaller segmentation while avoiding some of the complexities of microservices. If your workload is amenable to it, and your organization can support it, you should use a microservices architecture to achieve the best agility and reliability. AWS App Mesh can be used with service oriented architectures to provide reliable discovery and access of services.				
02. Build services focused on specific business domains and functionality	Manual	High	Open	

Recommendation:
Design your workload based on your business domains and their respective functionality: Focusing on specific functionality enables you to differentiate the reliability requirements of different services, and target investments more specifically. A concise business problem and having a small team associated with each service also enables easier organizational scaling.
Perform Domain Analysis to map out a domain driven design (DDD) for your workload. Then you can choose an architecture type to meet your workload needs
Decompose your services into smallest possible components: With microservices architecture you can separate your workload into components with the minimal functionality to enable organizational scaling and agility.
Define the API for the workload and its design goals, limits, and any other considerations for use. Define the API. The API definition should allow for growth and additional parameters. Define the designed availabilities.
Your API may have multiple design goals for different features. Establish limits Use testing to define the limits of your workload capabilities.

03. Provide service contracts per API	Manual	Medium	Open	

Recommendation:
Provide service contracts per API: Service contracts are documented agreements between teams on service integration and include a machine readable API definition, rate limits, and performance expectations.
A versioning strategy allows clients to continue using the existing API and migrate their applications to the newer API when they are ready.
Amazon API Gateway is a fully managed service that makes it easy for developers to create APIs at any scale. Using the OpenAPI Specification (OAS), formerly known as the Swagger Specification, you can define your API contract and import it intro API Gateway. With API Gateway, you can then version and deploy the APIs.

Pillar Name: Reliability

Question: 04. How do you design interactions in a distributed system to prevent failures?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Identify which kind of distributed system is required	Manual	High	Open	

Recommendation:
Identify which kind of distributed system is required: Challenges with distributed systems involved latency, scaling, understanding networking APIs, marshalling and unmarshalling data, and the complexity of algorithms such as Paxos. As the systems grow larger and more distributed, what had been theoretical edge cases turn into regular occurrences.
Hard real time distributed systems require responses to be given synchronously and rapidly.
Soft real time systems have a more generous time window of minutes or greater for response
Offline systems handle responses through batch or asynchronous processing.
Hard real time distributed systems have the most stringent reliability requirements.

02. Implement loosely coupled dependencies	Manual	High	Open	

Recommendation:
Implement loosely coupled dependencies: Dependencies such as queuing systems, streaming systems, workflows, and load balancers are loosely coupled. Loose coupling helps isolate behavior of a component from other components that depend on it, increasing resiliency and agility.

Pillar Name: Reliability

Question: 04. How do you design interactions in a distributed system to prevent failures?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Amazon EventBridge allows you to build event driven architectures, which are loosely coupled and distributed If changes to one component force other components that rely on it to also change, then they are tightly coupled. Loose coupling breaks this dependency so that dependency components only need to know the versioned and published interface. Make component interactions asynchronous where possible. This model is suitable for any interaction that does not need an immediate response and where an acknowledgement that a request has been registered will suffice.				
03. Make all responses idempotent	Manual	Medium	Open	
Recommendation: Make all responses idempotent: An idempotent service promises that each request is completed exactly once, such that making multiple identical requests has the same effect as making a single request. Clients can issue API requests with an idempotency token the same token is used whenever the request is repeated. An idempotent service API uses the token to return a response identical to the response that was returned the first time that the request was completed.				
04. Do constant work	Manual	Low	Open	
Recommendation: Do constant work: Systems can fail when there are large, rapid changes in load Engineer workloads so that payload sizes remain constant regardless of number of successes or failures. For example, if the health check system is monitoring 100,000 servers, the load on it is nominal under the normally light server failure rate. However, if a major event makes half those servers unhealthy, then the health check system would be overwhelmed trying to update notification systems and communicate state to its clients. Instead, the health check system should send the full snapshot of the current state each time. 100,000 server health states, each represented by a bit, would only be a 12.5 KB payload. Whether no servers or failing, or all of them, the health check system is doing constant work, and large, rapid changes are not a threat to the system stability. This is actually how the control plane is designed for Amazon Route 53 health checks.				

Pillar Name: Reliability

Question: 05. How do you design interactions in a distributed system to mitigate or withstand failures?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Implement graceful degradation to transform applicable hard dependencies into soft dependencies	Manual	High	Open	
Recommendation: Implement graceful degradation to transform applicable hard dependencies into soft dependencies: When a component's dependencies are unhealthy, the component itself can still function, although in a degraded manner. For example, when a dependency call fails, failover to a predetermined static response. By returning a static response, your workload mitigates failures that occur in its dependencies Detect when the retry operation is likely to fail, and prevent your client from making failed calls with the circuit breaker pattern				
02. Throttle requests	Manual	High	Open	
Recommendation: Throttle requests: This is a mitigation pattern to respond to an unexpected increase in demand. Some requests are honored but those over a defined limit are rejected and return a message indicating they have been throttled. The expectation on clients is that they will back off and abandon the request or try again at a slower rate. Use Amazon API Gateway				

Pillar Name: Reliability

Question: 05. How do you design interactions in a distributed system to mitigate or withstand failures?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
03. Control and limit retry calls	Manual	High	Open	

Recommendation:
Control and limit retry calls: Use exponential backoff to retry after progressively longer intervals. Introduce jitter to randomize those retry intervals, and limit the maximum number of retries.
Amazon SDKs implement this by default. Implement similar logic in your dependency layer when calling your own dependent services. Decide what the timeouts are and when to stop retrying based on your use case.

04. Fail fast and limit queues	Manual	High	Open	

Recommendation:
Fail fast and limit queues: If the workload is unable to respond successfully to a request, then fail fast. This allows the releasing of resources associated with a request, and permits the service to recover if its running out of resources. If the workload is able to respond successfully but the rate of requests is too high, then use a queue to buffer requests instead. However, do not allow long queues that can result in serving stale requests that the client has already given up on.
Implement fail fast when service is under stress
Limit queues: In a queue based system, when processing stops but messages keep arriving, the message debt can accumulate into a large backlog, driving up processing time. Work can be completed too late for the results to be useful, essentially causing the availability hit that queueing was meant to guard against.

05. Set client timeouts	Manual	High	Open	

Recommendation:
Set both a connection timeout and a request timeout on any remote call, and generally on any call across processes: Many frameworks offer built in timeout capabilities, but be careful as many have default values that are infinite or too high. A value that is too high reduces the usefulness of the timeout because resources continue to be consumed while the client waits for the timeout to occur. A too low value can generate increased traffic on the backend and increased latency because too many requests are retried. In some cases, this can lead to complete outages because all requests are being retried.

06. Make services stateless where possible	Manual	Medium	Open	

Recommendation:
Make your applications stateless: Stateless applications enable horizontal scaling and are tolerant to the failure of an individual node.
Remove state that could actually be stored in request parameters. Some data (like cookies) can be passed in headers or query parameters. Refactor to remove state that can be quickly passed in requests.
After examining whether the state is required, move any state tracking to a resilient Multi zone cache or data store like Amazon ElastiCache, Amazon RDS, Amazon DynamoDB, or a third party distributed data solution: Store a state that could not be moved to resilient data stores. Some data may not actually be needed per request and can be retrieved on demand. Remove data that can be asynchronously retrieved. Decide on a data store that meets the requirements for a required state. Consider a NoSQL database for non relational data.

07. Implement emergency levers	Manual	Medium	Open	

Recommendation:
Implement emergency levers: These are rapid processes that may mitigate availability impact on your workload. They can be operated in the absence of a root cause. An ideal emergency lever reduces the cognitive burden on the resolvers to zero by providing fully deterministic activation and deactivation criteria. Example levers include blocking all robot traffic or serving a static response. Levers are often manual, but they can also be automated
Tips for implementing and using emergency levers When levers are activated, do LESS, not more Keep it simple, avoid bimodal behavior Test your levers periodically
These are examples of actions that are NOT emergency levers Add capacity Call up service owners of clients that depend on your service and ask them to reduce calls Making a change to code and releasing it

Pillar Name: Reliability

Question: 06. How do you monitor workload resources?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Monitor all components for the workload (Generation)	Automated	High	Open	

Recommendation:
Enable logging where available: AWS has monitoring and log information available for consumption. Monitoring and logs can be used to define alerts, change, and recovery processes
Define all the AWS services you are using
Enable logging for all services: AWS has logging for many services. If the service doesn't have the logging at the level you wish, you can add logging from your workloads Enable logging of Amazon S3 Enable logging of Elastic Load Balancing Enable VPC Flow Logs Enable CloudTrail logs Use the Amazon CloudWatch Agent to stream log data from instance to CloudWatch Logs Use the awslogs log driver with Amazon ECS to stream log data to CloudWatch Logs AWS Lambda automatically streams log data to CloudWatch Logs
Consume all default metrics: Every service generates default metrics. Evaluate the metrics to decide which metrics on each service need alerts.
Metrics can be evaluated individually or in aggregate Go to the CloudWatch console and explore the metrics collected Refer to the documentation for which metrics and dimensions are collected
CloudWatch Synthetics enables you to get up Canary tests
Create custom metrics for your own use: AWS won't generate some metrics and combinations of metrics, but you can create them using custom metrics
If you need memory usage or disk consumption, use the CloudWatch Agent or PutMetricData API
Aggregate your logs: Log aggregation gives you a single place where you can look at log data and set alerts
Use CloudWatch Logs for common log files You can use CloudWatch Logs for most common log aggregation use cases
Store all logs in Amazon S3, or in Amazon S3 Glacier for longer term storage You can export CloudWatch Logs to Amazon S3. CloudTrail and Elastic Load Balancing logs are sent to Amazon S3

02. Define and calculate metrics (Aggregation)	Automated	High	Open	34

Recommendation:
Define and calculate metrics (Aggregation): Store log data and apply filters where necessary to calculate metrics, such as counts of a specific log event, or latency calculated from log event timestamps
Metric filters define the terms and patterns to look for in log data as it is sent to CloudWatch Logs. CloudWatch Logs uses these metric filters to turn log data into numerical CloudWatch metrics that you can graph or set an alarm on
Use a trusted third party to aggregate logs Follow the instructions of the third party. Most third party products integrate with CloudWatch and Amazon S3
Some AWS services can publish logs directly to Amazon S3. This way, if your main requirement for logs is storage in Amazon S3, you can easily have the service producing the logs send them directly to Amazon S3 without setting up additional infrastructure

03. Send notifications (Real-time processing and alarming)	Automated	High	Open	

Recommendation:
Perform real time processing and alarming: Organizations that need to know, receive notifications when significant events occur
Amazon CloudWatch dashboards are customizable home pages in the CloudWatch console that you can use to monitor your resources in a single view, even those resources that are spread across different Regions
Create an alarm when the metric surpasses a limit

04. Automate responses (Real-time processing and alarming)	Manual	Medium	Open	

Recommendation:
Use AWS Systems Manager to perform automated actions: AWS Config continuously monitors and records your AWS resource configurations, and can trigger AWS Systems Manager Automation to remediate issues
Create and use Systems Manager Automation documents. These define the actions that Systems Manager performs on your managed instances and other AWS resources when an automation execution runs
Amazon CloudWatch sends alarm state change events to Amazon EventBridge. Create EventBridge rules to automate responses
Create and execute a plan to automate responses
Inventory all your alert response procedures: You must plan your alert responses before you rank the tasks
Inventory all the tasks with specific actions that must be taken: Most of these actions are documented in runbooks. You must also have playbooks for alerts of unexpected events
Examine the runbooks and playbooks for all automatable actions: In general, if an action can be defined, it most likely can be automated
Rank the error prone or time consuming activities first: It is most beneficial to remove sources of errors and reduce time to resolution
Establish a plan to complete automation: Maintain an active plan to automate and update the automation
Examine manual requirements for opportunities for automation: Challenge your manual process for opportunities to automate

05. Analytics	Automated	Medium	Open	6

Pillar Name: Reliability

Question: 06. How do you monitor workload resources?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Recommendation: CloudWatch Logs Insights enables you to interactively search and analyze your log data in Amazon CloudWatch Logs Use Amazon CloudWatch Logs send logs to Amazon S3 where you can use or Amazon Athena to query the data Create an S3 lifecycle policy for your server access logs bucket. Configure the lifecycle policy to periodically remove log files. Doing so reduces the amount of data that Athena analyzes for each query				
06. Conduct reviews regularly	Automated	Medium	Open	34
Recommendation: Create multiple dashboards for the workload: You must have a top level dashboard that contains the key business metrics, as well as the technical metrics you have identified to be the most relevant to the projected health of the workload as usage varies. You should also have dashboards for various application tiers and dependencies that can be inspected Schedule and conduct regular reviews of the workload dashboards: Conduct regular inspection of the dashboards. You may have different cadences for the depth at which you inspect Inspect for trends in the metrics: Compare the metric values to historic values to see if there are trends that may indicate that something that needs investigation. Examples of this include: increasing latency, decreasing primary business function, and increasing failure responses Inspect for outliers/anomalies in your metrics: Averages or medians can mask outliers. Look at the highest and lowest values during the time frame and investigate the causes of extreme scores. As you continue to eliminate these causes, lowering your definition of extreme allows you to continue to improve the consistency of your workload performance Look for sharp changes in behavior: An immediate change in quantity or direction of a metric may indicate that there has been a change in the application, or external factors that you may need to add additional metrics to track				
07. Monitor end-to-end tracing of requests through your system	Manual	Medium	Open	
Recommendation: Monitor end to end tracing of requests through your system: AWS X Ray is a service that collects data about requests that your application serves, and provides tools you can use to view, filter, and gain insights into that data to identify issues and opportunities for optimization. For any traced request to your application, you can see detailed information not only about the request and response, but also about calls that your application makes to downstream AWS resources, microservices, databases and HTTP web APIs				

Pillar Name: Reliability

Question: 07. How do you design your workload to adapt to changes in demand?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Use automation when obtaining or scaling resources	Automated	High	Open	11
Recommendation: Configure and use AWS Auto Scaling: This monitors your applications and automatically adjusts capacity to maintain steady, predictable performance at the lowest possible cost. Using AWS Auto Scaling, you can setup application scaling for multiple resources across multiple services. Configure Auto Scaling on your Amazon EC2 instances and Spot Fleets, Amazon ECS tasks, Amazon DynamoDB tables and indexes, Amazon Aurora Replicas, and AWS Marketplace appliances as applicable. Use service API operations to specify the alarms, scaling policies, warm up times, and cool down times. Use Elastic Load Balancing: Load balancers can distribute load by path or by network connectivity. Application Load Balancers can distribute load by path. Configure an Application Load Balancer to distribute traffic to different workloads based on the path under the domain name. Application Load Balancers can be used to distribute loads in a manner that integrates with AWS Auto Scaling to manage demand. Network Load Balancers can distribute load by connection. Configure a Network Load Balancer to distribute traffic to different workloads using TCP, or to have a constant set of IP addresses for your workload. Network Load Balancers can be used to distribute loads in a manner that integrates with AWS Auto Scaling to manage demand.				

Pillar Name: Reliability

Question: 07. How do you design your workload to adapt to changes in demand?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Use a highly available DNS provider: DNS names allow your users to enter names instead of IP address to access your workloads and distributes this information to a defined scope, usually globally for users of the workload. Use Amazon Route 53 or a trusted DNS provider Use Route 53 to manage your CloudFront distributions and load balancers. Determine the domains and subdomains you are going to manage. Create appropriate record sets using ALIAS or CNAME records. Use the AWS global network to optimize the path from your users to your applications.: AWS Global Accelerator continually monitors the health of your application endpoints and redirects traffic to healthy endpoints in less than 30 seconds AWS Global Accelerator is a service that improves the availability and performance of your applications with local or global users. It provides static IP addresses that act as a fixed entry point to your application endpoints in a single or multiple AWS Regions, such as your Application Load Balancers, Network Load Balancers or Amazon EC2 instances. Configure and use Amazon CloudFront or a trusted content delivery network: A content delivery network (CDN) can provide faster end user response times and can serve requests for content that may cause unnecessary scaling of your workloads. Configure Amazon CloudFront distributions for your workloads, or use a third party CDN. You can limit access to your workloads so that they are only accessible from CloudFront by using the IP ranges for CloudFront in your endpoint security groups or access policies.				
02. Obtain resources upon detection of impairment to a workload	Automated	Medium	Open	11
Recommendation: Obtain resources upon detection of impairment to a workload: Scale resources reactively when necessary if availability is impacted, to restore workload availability. Use scaling plans which are the core component of AWS Auto Scaling. It's where you configure a set of instructions for scaling your resources. If you work with AWS CloudFormation or add tags to AWS resources, you can set up scaling plans for different sets of resources, per application. AWS Auto Scaling provides recommendations for scaling strategies customized to each resource. After you create your scaling plan, AWS Auto Scaling combines dynamic scaling and predictive scaling methods together to support your scaling strategy. Amazon EC2 Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application. You create collections of EC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes above this size. Amazon DynamoDB auto scaling uses the AWS Application Auto Scaling service to dynamically adjust provisioned throughput capacity on your behalf, in response to actual traffic patterns. This enables a table or a global secondary index to increase its provisioned read and write capacity to handle sudden increases in traffic, without throttling.				
03. Obtain resources upon detection that more resources are needed for a workload	Automated	Medium	Open	11
Recommendation: Obtain resources upon detection that more resources are needed for a workload: Scale resources proactively to meet demand and avoid availability impact Calculate how many compute resources you will need (compute concurrency) to handle a given request rate When you have a historical pattern for usage, setup scheduled scaling for Amazon EC2 auto scaling Use AWS predictive scaling				
04. Load test your workload	Manual	Medium	Open	
Recommendation: Perform load testing to identify which aspect of your workload indicates that you must add or remove capacity: Load testing should have representative traffic similar to what you receive in production. Increase the load while watching the metrics you have instrumented, to determine which metric indicates when you must add or remove resources. Identify the mix of requests: You may have varied mixes of requests, so you should look at various time frames when identifying the mix of traffic. Implement a load driver: You can use custom code, open source, or commercial software to implement a load driver. Load test initially using small capacity: You see some immediate effects by driving load onto a lesser capacity, possibly as small as one instance or container. Load test against larger capacity: The effects will be different on a distributed load, so you must test against as close to a product environment as possible.				

Pillar Name: Reliability

Question: 08. How do you implement change?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Use runbooks for standard activities such as deployment	Manual	High	Open	

Recommendation:
Enable consistent and prompt responses to well understood events by documenting procedures in runbooks.
Use the principle of infrastructure as code to define your infrastructure: By using AWS CloudFormation (or a trusted third party) to define your infrastructure, you can use version control software to version and track changes.
Use AWS CloudFormation (or a trusted third party provider) to define your infrastructure.
Create templates that are singular and decoupled, using good software design principles. Determine the permissions, templates, and responsible parties for implementation Use source control, like AWS CodeCommit or a trusted third party tool, for version control.

02. Integrate functional testing as part of your deployment	Manual	High	Open	

Recommendation:
Integrate functional testing as part of your deployment: Functional tests are run as part of automated deployment. If success criteria are not met, the pipeline is halted or rolled back.
Invoke AWS CodeBuild during the Test Action of your software release pipelines modeled in AWS CodePipeline. This capability enables you to easily run a variety of tests against your code, such as unit tests, static code analysis and integration tests.
Use AWS Marketplace solutions for executing automated tests as part of your software delivery pipeline.

03. Integrate resiliency testing as part of your deployment	Manual	Medium	Open	

Recommendation:
Integrate resiliency testing as part of your deployment: Use Chaos Engineering, the discipline of experimenting on a workload in order to build confidence in the workload capability to withstand turbulent conditions in production.
Resiliency tests inject faults or resource degradation to assess that your workload responds with its designed resilience
These tests can be run regularly in pre production environments in automated deployment pipelines.
They should also be run in production, as part of scheduled game days.
Using Chaos Engineering principles, propose hypotheses about how your workload will perform under various impairments, then test your hypotheses using resiliency testing.

04. Deploy using immutable infrastructure	Manual	Medium	Open	

Recommendation:
Deploy using immutable infrastructure: Deploy using an immutable infrastructure approach such as Blue/Green or Canary Deployment

05. Deploy changes with automation	Manual	Medium	Open	

Recommendation:
Automate your deployment pipeline: Deployment pipelines allow you to invoke automated testing and detection of anomalies, and either halt the pipeline at a certain step before production deployment, or automatically roll back a change.
Use AWS CodePipeline (or a trusted third party product) to define and execute your pipelines. Configure the pipeline to start when a change is committed to your code repository. Use Amazon Simple Notification Service (SNS) and Amazon Simple Email Service (SES) to send notifications about problems in the pipeline or integrate with a team chat tool, like Amazon Chime.

Pillar Name: Reliability

Question: 09. How do you back up data?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Identify and back up all data that needs to be backed up, or reproduce the data from sources	Automated	High	Open	9

Recommendation:
Understand and use the backup capabilities of the AWS services and resources used by your workload: AWS provides capabilities to back up your workload data
Take snapshots of your encrypted Amazon EC2 EBS volumes: You can back up the data on your Amazon EBS volumes to Amazon S3 by taking point in time snapshots. Snapshots are incremental backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved. Set up snapshot schedules and retention policies that meet your requirements.
Use AWS CloudFormation (or a trusted third party provider) to create your EBS volume with KMS encryption.
Take regular snapshots: Determine the schedule and appropriate processes for taking regular snapshots of EBS volumes. It may vary by application. Determine the need to replicate these snapshots based on your disaster recovery requirements.: You may need to replicate these to other accounts or AWS Regions
Replicate your Amazon EFS file systems: AWS does not take snapshots of Amazon EFS file systems by default. Set up snapshot schedules and retention policies that meet your requirements.
Automate an EFS to EFS backup
Create a scheduled event to cause the backup solution to deploy and execute. After you have the copy, you can use AWS API operations (or a trusted third party solution) to copy the data to Amazon S3 for higher durability: If the AWS Region you are using doesn't support Amazon EFS, use AWS Data Pipeline.
Take snapshots of your encrypted Amazon RDS instances: By default, AWS takes a snapshot each day and retains it for that day. Set up snapshot schedules and retention policies that meet your requirements.
Use AWS CloudFormation (or a trusted third party provider) to create your RDS instance with KMS encryption.
Set the preferred backup window, back up retention period, KMS key ID, and storage encryption options. Determine the need to replicate these snapshots based on your disaster recovery requirements.
Take snapshots of your Amazon DynamoDB tables: AWS does not take snapshots of DynamoDB tables by default. Set up snapshot schedules and retention policies that meet your requirements.
Use CloudFormation or a trusted third party provider to create your DynamoDB table and its associated TTL settings, IAM policies, and CloudWatch Alarms: Use the AWS CLI or AWS SDKs or a trusted third party provider to create your DynamoDB Auto Scaling.
Use the DynamoDB API operations to create backups and restore backups: The snapshot is automatically encrypted using AWS managed keys.
Use the DynamoDB API operations to enable Point in Time recovery: This allows you to restore to any time in the last 35 days.
Copy the log files in CloudWatch Logs to Amazon S3 for retention and archiving: AWS stores CloudWatch Logs log files for as long as you specify in your retention policy. If you need the logs longer for analytics, forensics, and archiving, you can export them to Amazon S3.
Use CloudFormation or a trusted third party provider to create your CloudWatch Logs log groups and their associated retention period in days.
Create the scheduled event to invoke an AWS Lambda function that will use the CloudWatch Logs GetLogEvents API and put the log data into Amazon S3. Specify the lifecycle policy on the S3 bucket for when the logs will be put into Amazon S3 Glacier for archiving and eventual deletion.

02. Secure and encrypt backups	Automated	High	Open	

Recommendation:
Use encryption on each of your data stores: If your source data is encrypted, then the backup will also be encrypted.
Enable encryption in RDS: You can configure encryption at rest using AWS Key Management Service when you create an RDS instance.
Enable encryption on EBS Volumes: You can configure default encryption or specify a unique key upon volume creation.
Use the required Amazon DynamoDB encryption: DynamoDB encrypts all data at rest. You can either use an AWS owned Customer Master Key (CMK) or an AWS managed CMK, specifying a key that is stored in your account.
Encrypt your data stored in Amazon EFS: Configure the encryption when you create your file system.
Configure the encryption in the source and destination Regions: You can configure encryption at rest in S3 using keys stored in KMS, but the keys are Region specific. You can specify the destination keys when you configure the replication.
Implement least privilege permissions to access your backups: Follow best practices to limit the access to the backups, snapshots, and replicas in accordance with security best practices.

03. Perform data backup automatically	Automated	Medium	Open	9

Recommendation:
Use AWS Backup to schedule your backups of services they support.: AWS Backup is a fully managed backup service that makes it easy to centralize and automate the backup of data across AWS services in the cloud and on premises.
Create backup plans: A backup plan is a policy expression that defines when and how you want to back up your AWS resources.
Create backup vaults: A backup vault is a container for organizing your backups.
Create backups: A backup represents the content of the resources at a specified time.

Pillar Name: Reliability

Question: 09. How do you back up data?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Create a EventBridge Event that invokes a Step Function State Machine to perform your backups.: You can create a State Machine in AWS Step Functions that coordinates your backups.				
04. Perform periodic recovery of the data to verify backup integrity and processes	Manual	Medium	Open	

Recommendation:
Include restoration in the automation of your backups.: You can extend your State Machine in AWS Step Functions to perform the restoration of each of the systems that you perform a backup.

Pillar Name: Reliability

Question: 10. How do you use fault isolation to protect your workload?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Deploy the workload to multiple locations	Automated	High	Verified	
02. Select the appropriate locations for your multi-location deployment	Manual	High	Open	

Recommendation:
Implement self healing: Deploy your instances or containers using automatic scaling when possible. If you cannot use automatic scaling, use automatic recovery for EC2 instances or implement self healing automation based on Amazon EC2 or ECS container lifecycle events.
Use Auto Scaling groups for instances and container workloads that have no requirements for a single instance IP address, private IP address, Elastic IP address, and instance metadata. The launch configuration user data can be used to implement automation that can self heal most workloads.
Use automatic recovery of EC2 instances for workloads that require a single instance ID address, private IP address, Elastic IP address, and instance metadata. Automatic Recovery will send recovery status alerts to a SNS topic as the instance failure is detected.
Use EC2 instance lifecycle events, or ECS events, to automate self healing where automatic scaling or EC2 recovery cannot be used. Use the events to invoke automation that will heal your component according to the process logic you require.
Use automatic recovery of EC2 instances for workloads that require a single instance ID address, private IP address, Elastic IP address, and instance metadata.
Automatic Recovery will send recovery status alerts to a SNS topic as the instance failure is detected.
Use EC2 instance lifecycle events or ECS events to automate self healing where automatic scaling or EC2 recovery cannot be used.
Use the events to invoke automation that will heal your component according to the process logic you require.

03. Automate recovery for components constrained to a single location	Manual	Medium	Open	

Recommendation:
Use bulkhead architectures: Like the bulkheads on a ship, this pattern ensures that a failure is contained to a small subset of requests/users so the number of impaired requests is limited, and most will continue without error. Bulkheads for data are usually called partitions or shards, while bulkheads for services are known as cells
Evaluate cell based architecture for your workload: In a cell based architecture, each cell is a complete, independent instance of the service and has a fixed maximum size. As load increases, workloads grow by adding more cells. A partition key is used on incoming traffic to determine which cell will process the request. Any failure is contained to the single cell it occurs in, so that the number of impaired requests is limited as other cells continue without error. It is important to identify the proper partition key to minimize cross cell interactions and avoid the need to involve complex mapping services in each request. Services that require complex mapping end up merely shifting the problem to the mapping services, while services that require cross cell interactions reduce the independence of cells (and thus the assumed availability improvements of doing so).
In his AWS blog post, Colm MacCarthaigh explains how Amazon Route 53 uses the concept of shuffle sharding to isolate customer requests into shards

04. Use bulkhead architectures to limit scope of impact	Manual	High	Open	

Recommendation:

Pillar Name: Reliability

Question: 10. How do you use fault isolation to protect your workload?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources

Use bulkhead architectures: Like the bulkheads on a ship, this pattern ensures that a failure is contained to a small subset of requests/users so the number of impaired requests is limited, and most will continue without error. Bulkheads for data are usually called partitions or shards, while bulkheads for services are known as cells
Evaluate cell based architecture for your workload: In a cell based architecture, each cell is a complete, independent instance of the service and has a fixed maximum size. As load increases, workloads grow by adding more cells. A partition key is used on incoming traffic to determine which cell will process the request. Any failure is contained to the single cell it occurs in, so that the number of impaired requests is limited as other cells continue without error. It is important to identify the proper partition key to minimize cross cell interactions and avoid the need to involve complex mapping services in each request. Services that require complex mapping end up merely shifting the problem to the mapping services, while services that require cross cell interactions reduce the independence of cells (and thus the assumed availability improvements of doing so).
In his AWS blog post, Colm MacCarthaigh explains how Amazon Route 53 uses the concept of shuffle sharding to isolate customer requests into shards

Pillar Name: Reliability

Question: 11. How do you design your workload to withstand component failures?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Monitor all components of the workload to detect failures	Automated	High	Open	

Recommendation:
Determine the collection interval for your components based on your recovery goals.
Your monitoring interval is dependent on how quickly you must recover: Your recovery time is driven by the time it takes to recover, so you must determine the frequency of collection by accounting for this time and your recovery time objective (RTO).
Configure detailed monitoring for components.
Determine if detailed monitoring for EC2 instances and Auto Scaling is necessary: Detailed monitoring provides 1 min interval metrics, and default monitoring provides 5 min interval metrics.
Determine if enhanced monitoring for RDS is necessary: Enhanced monitoring uses an agent on the RDS instances to get useful information about different process or threads on an RDS instance.
Create custom metrics to measure business Key Performance Indicators (KPIs) : Workloads implement key business functions. These functions should be used as KPIs that help identify when an indirect problem happens.
Monitor the user experience for failures using user canaries: Synthetic transaction testing (also known as "canary testing", but not to be confused with canary deployments) that can run and simulate customer behavior is among the most important testing processes. Run these tests constantly against your workload endpoints from diverse remote locations.
Create custom metrics that track the user's experience: If you can instrument the experience of the customer, you can determine when the consumer experience degrades.
Set alarms to detect when any part of your workload is not working properly, and to indicate when to Auto Scale resources. : Alarms can be visually displayed on dashboards, send alerts via SNS or email, and work with Auto Scaling to scale up or down the resources for a workload.
Create dashboards to visualize your metrics: Dashboards can be used to visually see trends, outliers, and other indicators of potential problems, or to provide an indication of problems you may want to investigate.

02. Fail over to healthy resources	Automated	High	Open	11

Recommendation:
Fail over to healthy resources: Ensure that if a resource failure occurs, that healthy resources can continue to serve requests. For location failures (such as Availability Zone or AWS Region) ensure you have systems in place to fail over to healthy resources in unimpaired locations.
If your workload is using AWS services, such as Amazon S3 or Amazon DynamoDB, then they are automatically deployed to multiple Availability Zones. In case of failure, the AWS control plane automatically routes traffic to healthy locations for you.
For Amazon RDS you must choose Multi AZ as a configuration option, and then on failure AWS automatically directs traffic to the healthy instance.
For Amazon EC2 instances or Amazon ECS tasks, you choose which Availability Zones to deploy to. Elastic Load Balancing then provides the solution to detect instances in unhealthy zones and route traffic to the healthy ones. Elastic Load Balancing can even route traffic to components in your on premises data center.
For multi region approaches (which might also include on premises data centers), ensure that data and resources from healthy locations can continue to serve requests For example, cross region read replicas enable you to deploy your data to multiple AWS Regions, but you still must promote the read replica to master and point your traffic at it in the event of a primary location failure.

Pillar Name: Reliability

Question: 11. How do you design your workload to withstand component failures?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Amazon Route 53 provides a way to define internet domains, and assign routing policies, which might include health checks, to ensure that traffic is routed to healthy Regions. Alternately, AWS Global Accelerator provides static IP addresses that act as a fixed entry point to your application, then routes to endpoints in AWS Regions of your choosing, using the AWS global network instead of the public internet for better performance and reliability.				
03. Automate healing on all layers	Manual	High	Open	
Recommendation: Use Auto Scaling groups to deploy tiers in an Application: Auto scaling can perform self healing on stateless applications, and add and remove capacity. Implement automatic recovery on EC2 instances that have applications deployed that cannot be deployed in multiple locations, and can tolerate rebooting upon failures. : Automatic recovery can be used to replace failed hardware and restart the instance when the application is not capable of being deployed in multiple locations. The instance metadata and associated IP addresses are kept, as well as the Amazon EBS volumes and mount points to Elastic File Systems or File Systems for Lustre and Windows. Using AWS OpsWorks, you can configure Auto Healing of EC2 instances at the layer level Implement automated recovery using AWS Step Functions and AWS Lambda when you cannot use automatic scaling or automatic recovery, or when automatic recovery fails. : When you cannot use automatic scaling, and either cannot use automatic recovery or automatic recovery fails, you can automate the healing using AWS Step Functions and AWS Lambda. Amazon EventBridge can be used to monitor and filter for events such as CloudWatch Alarms or changes in state in other AWS services. Based on event information, it can then trigger AWS Lambda (or other targets) to execute custom remediation logic on your workload.				
04. Rely on the data plane and not the control plane during recovery	Automated	High	Verified	
05. Use static stability to prevent bimodal behavior	Manual	Medium	Open	
Recommendation: Alarms on business Key Performance Indicators when they exceed a low threshold: Having a low threshold alarm on your business KPIs help you know when your workload is unavailable or non functional. Alarm on events that invoke healing automation: You can directly invoke an SNS API to send notifications with any automation that you create.				
06. Send notifications when events impact availability	Automated	Medium	Open	
Recommendation: Send notifications when events impact availability				
07. Architect your product to meet availability targets and uptime service level agreements (SLAs)	Manual	Medium	Open	
Recommendation: Architect your product to meet availability targets and uptime service level agreements (SLAs)				

Pillar Name: Reliability

Question: 12. How do you test reliability?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Use playbooks to investigate failures	Manual	High	Open	
Recommendation:				

Pillar Name: Reliability

Question: 12. How do you test reliability?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Use playbooks to identify issues: Playbooks are documented processes to investigate issues. Enable consistent and prompt responses to failure scenarios by documenting processes in playbooks. Playbooks must contain the information and guidance necessary for an adequately skilled person to gather applicable information, identify potential sources of failure, isolate faults, and determine contributing factors (i.e. perform post incident analysis). Implement playbooks as code: Perform your operations as code by scripting your playbooks to ensure consistency and limit reduce errors caused by manual processes. Playbooks can be composed of multiple scripts representing the different steps that might be necessary to identify the contributing factors to an issue. Runbook activities can be triggered or performed as part of playbook activities, or may prompt for execution of a playbook in response to identified events.				
02. Perform post-incident analysis	Manual	High	Open	
Recommendation: Establish a standard for your post incident analysis: Good post incident analysis provides opportunities to propose common solutions for problems with architecture patterns that are used in other places in your systems. Ensure that the contributing factors are honest and blame free. If you do not document your problems, you cannot correct them. Ensure post incident analysis is blame free so you can be dispassionate about the proposed corrective actions and promote honest self assessment and collaboration on your application teams. Use a process to determine contributing factors: Have a process to identify and document the contributing factors of an event so that you can develop mitigations to limit or prevent recurrence and you can develop procedures for prompt and effective responses. Communicate contributing factors as appropriate, tailored to target audiences.				
03. Test functional requirements	Manual	High	Open	
Recommendation: Test functional requirements: These include unit tests and integration tests that validate required functionality.				
04. Test scaling and performance requirements	Manual	High	Open	
Recommendation: Test scaling and performance requirements: Perform load testing to validate that the workload meets scaling and performance requirements. Deploy your application in an environment identical to your production environment and execute a load test. Use infrastructure as code concepts to create an environment as similar to your production environment as possible.				
05. Test resiliency using chaos engineering	Manual	Medium	Open	
Recommendation: Test resiliency using chaos engineering: Run tests that inject failures regularly into pre production and production environments. Hypothesize how your workload will react to the failure, then compare your hypothesis to the testing results and iterate if they do not match. Ensure that production testing does not impact users. To inject fault into your workload use open source software Or use commercial software available through AWS Marketplace Or create your own failure injection code Test the failure of all components and external dependencies. Simulate conditions that can produce brownouts using extensions to common proxies to introduce latency and dropped messages. You can also create your own implementations to create brownout conditions.				
06. Conduct game days regularly	Manual	Medium	Open	
Recommendation: Schedule game days to regularly exercise your runbooks and playbooks.: Game days should involve everyone who would be involved in a production interruption: business owner, development staff, operational staff, and incident response teams. Execute your load or performance tests and then execute your failure injection. Look for anomalies in your runbooks and opportunities to exercise your playbooks. If you deviate from your runbooks, refine the runbook or correct the behavior. If you exercise your playbook, identify the runbook that should have been used, or create a new one.				

Pillar Name: Reliability

Question: 13. How do you plan for disaster recovery (DR)?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Define recovery objectives for downtime and data loss	Manual	High	Open	

Recommendation:
Establish categories of need for your workloads: Identify the primary business driver and enabler workloads. Identify the workloads that are internal only tools, and the workloads that are externally visible tools. Identify the business impact of down time for each workload. Create five or fewer categories and refine the range of your recovery time objective (RTO) and recovery point objective (RPO) requirements.
Identify the business mission critical workload typically the main revenue drivers and enablers
Identify the business important workload typically reporting and runtime workload modification tools (like content management systems)
Identify the non business driving workloads where data may be difficult to recreate (like test systems with cleansed data)
Identify the non business driving workloads where data is less difficult or easy to recreate (like development environments)
Identify other categories as needed

02. Use defined recovery strategies to meet the recovery objectives	Automated	High	Open	11

Recommendation:
Establish strategies to achieve the recovery time objective (RTO) and recovery point objective (RPO) for each category: If a multi region strategy is necessary for your workload, you should choose one of the following strategies. They are listed in increasing order of complexity, and decreasing order of RTO and RPO. Backup and restore to another AWS region can add another layer of assurance that data will be available when needed, but for the other strategies you should weigh their potential complexity and cost versus what you can achieve using multiple Availability Zones within an AWS Region.
Backup and restore (RPO in hours, RTO in 24 hours or less: Back up your data and applications into the DR Region. Restore this data when necessary to recover from a disaster.
Pilot light (RPO in minutes, RTO in hours): Maintain a minimal version of an environment always running the most critical core elements of your system in the DR Region. When the time comes for recovery, you can rapidly provision a full scale production environment around the critical core.
Warm standby (RPO in seconds, RTO in minutes): Maintain a scaled down version of a fully functional environment always running in the DR Region. Business critical systems are fully duplicated and are always on, but with a scaled down fleet. When the time comes for recovery, the system is scaled up quickly to handle the production load.
Multi region active active (RPO is none or possibly seconds, RTO in seconds): Your workload is deployed to, and actively serving traffic from, multiple AWS Regions. This strategy requires you to synchronize users and data across the Regions that you are using. When the time comes for recovery, use services like Amazon Route 53 or AWS Global Accelerator to route your user traffic to where your workload is healthy.

03. Test disaster recovery implementation to validate the implementation	Manual	High	Open	

Recommendation:
Engineer your workloads for recovery. Regularly test your recovery paths: Recovery Oriented Computing (ROC) identifies the characteristics in systems that enhance recovery. These characteristics are: isolation and redundancy, system wide ability to roll back changes, ability to monitor and determine health, ability to provide diagnostics, automated recovery, modular design, and ability to restart. Exercise the recovery path to ensure that you can accomplish the recovery in the specified time to the specified state. Use your runbooks during this recovery to document problems and find solutions for them before the next test.
Use CloudEndure Disaster Recovery to implement and test your DR strategy

04. Manage configuration drift at the DR site or Region	Automated	Medium	Verified	
05. Automate recovery	Automated	Medium	Open	11

Recommendation:
Automate recovery paths: For short recovery times, human judgment and action cannot be used for high availability scenarios. The system should automatically recover under every situation.
Use CloudEndure Disaster Recovery for automated Failover and Failback: CloudEndure Disaster Recovery continuously replicates your machines (including operating system, system state configuration, databases, applications, and files) into a low cost staging area in your target AWS account and preferred Region. In the case of a disaster, you can instruct CloudEndure Disaster Recovery to automatically launch thousands of your machines in their fully provisioned state in minutes.

Pillar Name: Security				
Question: 01. How do you securely operate your workload?				
Question Risk: High				
Best Practice	**Nature**	**Severity**	**Status**	**Violated Resources**
01. Separate workloads using accounts	Automated	High	Open	
Recommendation: Use AWS Organizations: Use AWS Organizations to centrally enforce policy based management for multiple AWS accounts. Consider AWS Control Tower: AWS Control Tower provides an easy way to set up and govern a new, secure, multi account AWS environment based on best practices.				
02. Secure account root user and properties	Automated	High	Verified	
03. Identify and validate control objectives	Manual	High	Open	
Recommendation: Identify compliance requirements: Discover the organizational, legal, and compliance requirements that your workload must comply with. Identify AWS compliance resources: Identify resources that AWS has available to assist you with compliance.				
04. Keep up-to-date with security threats	Automated	High	Open	
Recommendation: Subscribe to threat intelligence sources: Regularly review threat intelligence information from multiple sources that are relevant to the technologies used in your workload. Consider AWS Shield Advanced service: It provides near real time visibility into intelligence sources, if your workload is internet accessible.				
05. Keep up-to-date with security recommendations	Automated	High	Open	
Recommendation: Follow AWS updates: Subscribe or regularly check for new recommendations, tips and tricks. Subscribe to industry news: Regularly review news feeds from multiple sources that are relevant to the technologies that are used in your workload.				
06. Automate testing and validation of security controls in pipelines	Manual	Medium	Open	
Recommendation: Automate configuration management: Enforce and validate secure configurations automatically by using a configuration management service or tool.				
07. Identify threats and prioritize mitigations using a threat model	Manual	High	Open	
Recommendation: Create a threat model: A threat model can help you identify and address potential security threats.				
08. Evaluate and implement new security services and features regularly	Manual	Low	Open	
Recommendation: Plan regular reviews: Create a calendar of review activities that includes compliance requirements, evaluation of new AWS security features and services, and staying up to date with industry news. Discover AWS services and features: Discover the security features that are available for the services that you are using, and review new features as they are released. Define AWS service on boarding process: Define processes for on boarding of new AWS services. Include how you evaluate new AWS services for functionality, and the compliance requirements for your workload. Test new services and features: Test new services and features as they are released in a non production environment that closely replicates your production one. Implement other defense mechanisms: Implement automated mechanisms to defend your workload, explore the options available.				

Pillar Name: Security

Question: 02. How do you manage identities for people and machines?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Use strong sign-in mechanisms	Automated	High	Open	

Recommendation:
Create an IAM policy to enforce MFA sign in: Create a customer managed IAM policy that prohibits all IAM actions except for the ones that allow a user to assume roles, change their own credentials, and manage their MFA devices on the My Security Credentials page.
Enable MFA in your identity provider: Enable MFA in the identity provider or single sign on service, such as AWS Single Sign On (SSO), that you use.
Configure strong password policy: Configure a strong password policy in IAM and federated identity systems to help protect against brute force attacks.
Rotate credentials regularly: Ensure administrators of your workload change their passwords and access keys (if used) regularly.

02. Use temporary credentials	Automated	High	Open	1

Recommendation:
Implement least privilege policies: Assign access policies with least privilege to IAM groups and roles to reflect the user's role or function that you have defined.
Remove unnecessary permissions: Implement least privilege by removing permissions that are unnecessary.
Consider permissions boundaries: A permissions boundary is an advanced feature for using a managed policy that sets the maximum permissions that an identity based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity based policies and its permissions boundaries.
Consider resource tags for permissions: You can use tags to control access to your AWS resources that support tagging. You can also tag IAM users and roles to control what they can access.

03. Store and use secrets securely	Automated	High	Open	

Recommendation:
Use AWS Secrets Manager: AWS Secrets Manager is an AWS service that makes it easier for you to manage secrets. Secrets can be database credentials, passwords, third party API keys, and even arbitrary text.

04. Rely on a centralized identity provider	Automated	High	Verified	
05. Audit and rotate credentials periodically	Automated	Medium	Open	

Recommendation:
Regularly audit credentials: Use credential reports, and IAM Access Analyzer to audit IAM credentials and permissions.
Use Access Levels to Review IAM Permissions: To improve the security of your AWS account, regularly review and monitor each of your IAM policies. Make sure that your policies grant the least privilege that is needed to perform only the necessary actions.
Consider automating IAM resource creation and updates: AWS CloudFormation can be used to automate the deployment of IAM resources including roles and policies, to reduce human error, as the templates can be verified and version controlled.

06. Leverage user groups and attributes	Manual	Low	Open	

Recommendation:
If you are using AWS Single Sign On (SSO), configure groups: AWS SSO provides you with the ability to configure groups of users, and assign groups the desired level of permission.
Learn about attribute based access control (ABAC): Attribute based access control (ABAC) is an authorization strategy that defines permissions based on attributes.

Pillar Name: Security

Question: 03. How do you manage permissions for people and machines?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Define access requirements	Automated	High	Open	1
Recommendation: Define required privileges for job function and responsibilities: Based on the user's job function, role, or responsibilities, define what resources that they need access to and the conditions that may apply. Group the users with common requirements together to make delegation of policies easier.				
02. Grant least privilege access	Automated	High	Verified	
03. Establish emergency access process	Manual	Medium	Open	
Recommendation: Pre provision emergency access: Pre provisioning a role for emergency access from a trusted account, for example one that is used for security team, can help you gain access quickly.				
04. Reduce permissions continuously	Automated	Medium	Verified	
05. Define permission guardrails for your organization	Automated	Medium	Verified	
06. Manage access based on lifecycle	Automated	Low	Verified	
07. Analyze public and cross-account access	Automated	Low	Open	1
Recommendation: Configure IAM Access Analyzer: AWS IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, that are shared with an external entity.				
08. Share resources securely within your organization	Automated	Medium	Open	
Recommendation: Use AWS Resource Access Manager: AWS Resource Access Manager (RAM) is a service that enables you to easily and securely share AWS resources with any AWS account or within your AWS Organization.				
09. Share resources securely with a third party	Manual	Medium	Open	
Recommendation: Use AWS Resource Access Manager: AWS Resource Access Manager (RAM) is a service that enables you to easily and securely share AWS resources with any AWS account or within your AWS Organization.				

Pillar Name: Security

Question: 04. How do you detect and investigate security events?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Configure service and application logging	Automated	High	Open	34
Recommendation:				

Pillar Name: Security

Question: 04. How do you detect and investigate security events?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Enable logging of AWS services: Enable the logging of AWS services to meet your requirements. Logging capabilities include the following: VPC Flow Logs, ELB logs, S3 bucket logs, CloudFront access logs, Route 53 query logs, and Amazon RDS logs. Evaluate and enable logging of operating systems and application specific: Evaluate and enable logging of operating systems and application specific logs to detect suspicious behavior. Apply appropriate controls to the logs: Logs can contain sensitive information and only authorized users should have access. Consider restricting permissions to S3 buckets and CloudWatch Logs log groups. Configure Amazon GuardDuty: Amazon GuardDuty is a threat detection service that continuously looks for malicious activity and unauthorized behavior to protect your AWS accounts and workloads. Enable GuardDuty and configure automated alerts to email using the lab. Configure customized trail in CloudTrail: Configuring a trail enables you to store logs for longer than the default period, and analyze them later. Enable AWS Config: AWS Config provides a detailed view of the configuration of AWS resources in your AWS account. This view includes how the resources are related to one another and how they were previously configured so that you can see how the configurations and relationships change over time. Enable AWS Security Hub: AWS Security Hub provides you with a comprehensive view of your security state in AWS and helps you check your compliance with the security industry standards and best practices. Security Hub collects security data from across AWS accounts, services, and supported third party partner products and helps you analyze your security trends and identify the highest priority security issues.				
02. Analyze logs, findings, and metrics centrally	Automated	High	Open	
Recommendation: Evaluate log processing capabilities: Evaluate the options that are available for processing logs As a start for analyzing CloudTrail logs, test Amazon Athena Implement centralize logging in AWS: AWS example solution to centralize logging from multiple sources. Implement centralize logging with partner: APN Partners have solutions to help you analyze logs centrally.				
03. Automate response to events	Manual	Medium	Open	
Recommendation: Implement automated alerting with Amazon GuardDuty: Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts and workloads. Enable GuardDuty and configure automated alerts. Automate investigation processes: Develop automated processes that investigate an event and report information to an administrator to save time.				
04. Implement actionable security events	Automated	Low	Open	
Recommendation: Discover metrics available for AWS services: Discover the metrics that are available through CloudWatch for the services that you are using. Configure Amazon CloudWatch alarms: .				

Pillar Name: Security

Question: 05. How do you protect your network resources?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Create network layers	Manual	High	Open	
Recommendation: Create subnets in VPC: Create subnets for each layer (in groups that include multiple availability zones), and associate route tables to control routing.				
02. Control traffic at all layers	Automated	High	Open	24

Pillar Name: Security

Question: 05. How do you protect your network resources?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Recommendation: Control network traffic in a VPC: Implement VPC best practices to control traffic Control traffic at the edge: Implement edge services, such as Amazon CloudFront, to provide an additional layer of protection and other features. Control private network traffic: Implement services that protect your private traffic for your workload.				
03. Automate network protection	Automated	Medium	Open	
Recommendation: Automate protection for web based traffic: AWS offers a solution that uses AWS CloudFormation to automatically deploy a set of AWS WAF rules designed to filter common web based attacks. Users can select from preconfigured protective features that define the rules included in an AWS WAF web access control list (web ACL). Consider APN Partner solutions: APN Partners offer hundreds of industry leading products that are equivalent, identical to, or integrate with existing controls in your on premises environments. These products complement the existing AWS services to enable you to deploy a comprehensive security architecture and a more seamless experience across your cloud and on premises environments.				
04. Implement inspection and protection	Automated	Low	Open	6
Recommendation: Configure Amazon GuardDuty: Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts and workloads. Enable GuardDuty and configure automated alerts. Configure VPC Flow Logs: VPC Flow Logs is a feature that enables you to capture information about the IP traffic going to and from network interfaces in your VPC. Flow log data can be published to Amazon CloudWatch Logs and Amazon S3. After you've created a flow log, you can retrieve and view its data in the chosen destination. Consider VPC traffic mirroring: Traffic Mirroring is an Amazon VPC feature that you can use to copy network traffic from an elastic network interface of Amazon EC2 instances and then send it to out of band security and monitoring appliances for content inspection, threat monitoring, and troubleshooting.				

Pillar Name: Security

Question: 06. How do you protect your compute resources?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Perform vulnerability management	Automated	High	Verified	
02. Reduce attack surface	Automated	High	Open	
Recommendation: Harden operating system: Configure operating systems to meet best practices. Harden containerized resources: Configure containerized resources to meet security best practices. AWS Lambda best practices: Implement AWS Lambda best practices				
03. Implement managed services	Automated	Medium	Verified	
04. Automate compute protection	Automated	Medium	Open	
Recommendation: Automate configuration management: Enforce and validate secure configurations automatically by using a configuration management service or tool.				

Pillar Name: Security

Question: 06. How do you protect your compute resources?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Automate patching of EC2 instances: AWS Systems Manager Patch Manager automates the process of patching managed instances with both security related and other types of updates. You can use Patch Manager to apply patches for both operating systems and applications. Implement intrusion detection and prevention: Implement an intrusion detection and prevention tool to monitor and stop malicious activity on instances. Consider APN Partner solutions: APN Partners offer hundreds of industry leading products that are equivalent, identical to, or integrate with existing controls in your on premises environments. These products complement the existing AWS services to enable you to deploy a comprehensive security architecture and a more seamless experience across your cloud and on premises environments.				
05. Enable people to perform actions at a distance	Manual	Low	Open	
Recommendation: Replace console access: Replace console access (SSH or RDP) to instances with AWS Systems Manager Run Command to automate management tasks.				
06. Validate software integrity	Manual	Low	Open	
Recommendation: Investigate mechanisms: Code signing is one mechanism that can be used to validate software integrity.				

Pillar Name: Security

Question: 07. How do you classify your data?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Identify the data within your workload	Manual	High	Open	
Recommendation: Consider discovering data using Amazon Macie: Amazon Macie recognizes sensitive data such as personally identifiable information (PII) or intellectual property.				
02. Define data protection controls	Manual	High	Open	
Recommendation: Define your data identification and classification schema: Identification and classification of your data is performed to assess the potential impact and type of data you store, and who can access it. Discover available AWS controls: For the AWS services you are or plan to use, discover the security controls. Many services have a security section in their documentation Identify AWS compliance resources: Identify resources that AWS has available to assist.				
03. Automate identification and classification	Manual	Medium	Open	
Recommendation: Use Amazon S3 Inventory: Amazon S3 Inventory Amazon S3 inventory is one of the tools you can use to audit and report on the replication and encryption status of your objects. Consider Amazon Macie: Amazon Macie uses machine learning to automatically discover and classify data stored in Amazon S3.				
04. Define data lifecycle management	Manual	Low	Open	
Recommendation: Identify data types: Identify the types of data that you are storing or processing in your workload. That data could be text, images, binary databases, etc.				

Pillar Name: Security				
Question: 08. How do you protect your data at rest?				
Question Risk: High				
Best Practice	**Nature**	**Severity**	**Status**	**Violated Resources**
01. Implement secure key management	Automated	High	Open	
Recommendation: Implement AWS Key Management Service (AWS KMS): AWS Key Management Service (AWS KMS) makes it easy for you to create and manage keys and control the use of encryption across a wide range of AWS services and in your applications. AWS KMS is a secure and resilient service that uses FIPS 140 2 validated hardware security modules to protect your keys. Consider AWS Encryption SDK: Use the AWS Encryption SDK with AWS KMS integration when your application needs to encrypt data client side.				
02. Enforce encryption at rest	Automated	High	Open	
Recommendation: Enforce encryption at rest for Amazon S3: Implement S3 bucket default encryption. Use AWS Secrets Manager: AWS Secrets Manager is an AWS service that makes it easy for you to manage secrets. Secrets can be database credentials, passwords, third party API keys, and even arbitrary text. Configure default encryption for new EBS volumes: Specify that you want all newly created EBS volumes to be created in encrypted form, with the option of using the default key provided by AWS, or a key that you create. Configure encrypted Amazon Machine Images (AMIs): Copying an existing AMI with encryption enabled will automatically encrypt root volumes and snapshots. Configure Amazon RDS encryption: Configure encryption for your Amazon RDS DB clusters and snapshots at rest by enabling the encryption option. Configure encryption in additional AWS services: For the AWS services you use, determine the encryption capabilities.				
03. Automate data at rest protection	Automated	Medium	Open	
Recommendation: Implement AWS Key Management Service (AWS KMS): AWS Key Management Service (AWS KMS) makes it easy for you to create and manage keys and control the use of encryption across a wide range of AWS services and in your applications. AWS KMS is a secure and resilient service that uses FIPS 140 2 validated hardware security modules to protect your keys. Consider AWS Encryption SDK: Use the AWS Encryption SDK with AWS KMS integration when your application needs to encrypt data client side.				
04. Enforce access control	Automated	Low	Open	7
Recommendation: Enforce access control: Enforce access control with least privileges, including access to encryption keys. Separate data based on different classification levels: Use different AWS accounts for data classification levels managed by AWS Organizations. Review AWS KMS policies: Review the level of access granted in AWS KMS policies. Review S3 bucket and object permissions: Regularly review the level of access granted in Amazon S3 bucket policies. Best practice is to not have publicly readable or writeable buckets. Consider using AWS Config to detect buckets that are publicly available, and Amazon CloudFront to serve content from Amazon S3. Enable Amazon S3 versioning and object lock Use Amazon S3 Inventory: Amazon S3 Inventory Amazon S3 inventory is one of the tools you can use to audit and report on the replication and encryption status of your objects. Review Amazon EBS and AMI sharing permissions: Sharing permissions can allow images and volumes to be shared to AWS accounts external to your workload.				
05. Use mechanisms to keep people away from data	Automated	Low	Verified	

Pillar Name: Security

Question: 09. How do you protect your data in transit?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Implement secure key and certificate management	Automated	High	Open	

Recommendation:
Implement secure key and certificate management: Implement your defined secure key and certificate management solution.
Implement secure protocols: Use secure protocols that offer authentication and confidentiality, such as Transport Layer Security (TLS) or IPsec, to reduce the risk of data tampering or loss. Check the AWS documentation for the protocols and security relevant to the services that you are using.

02. Enforce encryption in transit	Automated	High	Open	

Recommendation:
Enforce encryption in transit: Your defined encryption requirements should be based on the latest standards and best practices and only allow secure protocols. For example, only configure a security group to allow HTTPS protocol to an Application Load Balancer or EC2 instance.
Configure secure protocols in edge services: Configure HTTPS with Amazon CloudFront and required ciphers.
Use a VPN for external connectivity: Consider using an IPsec VPN for securing point to point or network to network connections to provide both data privacy and integrity.
Configure secure protocols in load balancers: Enable HTTPS listener for securing connections to load balancers.
Configure secure protocols for instances: Consider configuring HTTPS encryption on instances.
Configure secure protocols in Amazon Relational Database Service (Amazon RDS): Use SSL/TLS to encrypt connection to database instances.
Configure secure protocols in Amazon Redshift: Configure your cluster to require an SSL/TLS connection.
Configure secure protocols in additional AWS services: For the AWS services you use, determine the encryption in transit capabilities.

03. Automate detection of unintended data access	Automated	Medium	Open	

Recommendation:
Automate detection of unintended data access: Use a tool or detection mechanism to automatically detect attempts to move data outside of defined boundaries, for example, to detect a database system that is copying data to an unrecognized host.
Consider Amazon Macie: Amazon Macie continuously monitors data access activity for anomalies, and generates detailed alerts when it detects risk of unauthorized access or inadvertent data leaks.

04. Authenticate network communications	Manual	Low	Open	

Recommendation:
Implement secure protocols: Use secure protocols that offer authentication and confidentiality, such as Transport Layer Security (TLS) or IPsec, to reduce the risk of data tampering or loss. Check the AWS documentation for the protocols and security relevant to the services you are using.

Pillar Name: Security

Question: 10. How do you anticipate, respond to, and recover from incidents?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Identify key personnel and external resources	Automated	High	Verified	
02. Develop incident management plans	Manual	High	Open	

Recommendation:

Pillar Name: Security

Question: 10. How do you anticipate, respond to, and recover from incidents?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Review available resources: AWS and industry resources are available for you to use. Develop incident response playbooks: Easy to follow playbooks should detail steps that you would take to respond and recover from an incident. Develop escalation and communications plans: Escalation and communications plans should include personnel in your organization, and external parties that you must notify at each stage during an incident. Develop external public relations plan: Develop a plan for public relations to release information about an incident.				
03. Prepare forensic capabilities	Automated	Medium	Open	34
Recommendation: Identify forensic capabilities: Research your organization's forensic investigation capabilities, available tools, and external specialists.				
04. Automate containment capability	Manual	Medium	Open	
Recommendation: Automate containment capability				
05. Pre-provision access	Manual	Medium	Open	
Recommendation: Pre provision access: Ensure that security personnel have the correct access pre provisioned in AWS so that an appropriate response can be made to an incident.				
06. Pre-deploy tools	Automated	Low	Verified	
07. Run game days	Manual	Low	Open	
Recommendation: Run game days: Run simulated incident response events (game days) for different threats that involve key staff and management. Capture lessons learned: Lessons learned from running game days should be part of a feedback look to improve your processes.				

Pillar Name: Security

Question: 11. How do you incorporate and validate the security properties of applications throughout the design, development, and deployment lifecycle?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
01. Train for application security	Manual	Medium	Open	
Recommendation: Provide training to the builders in your organization on common practices for the secure development and operation of applications. Adopting security focused development practices helps reduce the likelihood of issues that are only detected at the security review stage.				
02. Automate testing throughout the development and release lifecycle	Manual	Medium	Open	
Recommendation:				

Pillar Name: Security

Question: 11. How do you incorporate and validate the security properties of applications throughout the design, development, and deployment lifecycle?

Question Risk: High

Best Practice	Nature	Severity	Status	Violated Resources
Automate the testing for security properties throughout the development and release lifecycle. Automation makes it easier to consistently and repeatably identify potential issues in software prior to release, which reduces the risk of security issues in the software being provided.				
03. Perform regular penetration testing	Manual	High	Open	
Recommendation: Perform regular penetration testing of your software. This mechanism helps identify potential software issues that cannot be detected by automated testing or a manual code review. It can also help you understand the efficacy of your detective controls. Penetration testing should try to determine if the software can be made to perform in unexpected ways, such as exposing data that should be protected, or granting broader permissions than expected.				
04. Manual code reviews	Manual	Medium	Open	
Recommendation: Perform a manual code review of the software that you produce. This process helps verify that the person who wrote the code is not the only one checking the code quality.				
05. Deploy software programmatically	Manual	Medium	Open	
Recommendation: Provide centralized services for builder teams to obtain software packages and other dependencies. This allows the validation of packages before they are included in the software that you write, and provides a source of data for the analysis of the software being used in your organization.				
06. Deploy software programmatically	Manual	High	Open	
Recommendation: Perform software deployments programmatically where possible. This approach reduces the likelihood that a deployment fails or an unexpected issue is introduced due to human error.				
07. Regularly assess security properties of the pipelines	Manual	High	Open	
Recommendation: Apply the principles of the Well Architected Security Pillar to your pipelines, with particular attention to the separation of permissions. Regularly assess the security properties of your pipeline infrastructure. Effectively managing the security of the pipelines allows you to deliver the security of the software that passes through the pipelines.				
08. Build a program that embeds security ownership in workload teams	Manual	Low	Open	
Recommendation: Build a program or mechanism that empowers builder teams to make security decisions about the software that they create. Your security team still needs to validate these decisions during a review, but embedding security ownership in builder teams allows for faster, more secure workloads to be built. This mechanism also promotes a culture of ownership that positively impacts the operation of the systems you build.				

Pillar Name: Sustainability

Question: 01. How do you select Regions for your workload?

Question Risk: Medium

Best Practice	Nature	Severity	Status	Violated Resources
01. Choose Region based on both business requirements and sustainability goals	Manual	Medium	Open	

Recommendation:
Choose Regions near Amazon renewable energy projects and Regions where the grid has a published carbon intensity that is lower than other locations (or Regions).

Pillar Name: Sustainability

Question: 02. How do you align cloud resources to your demand?

Question Risk: Medium

Best Practice	Nature	Severity	Status	Violated Resources
01. Scale workload infrastructure dynamically	Automated	Medium	Open	11

Recommendation:
Analyze the effect of users on load and capacity utilization over time and respond to changes in demand by scaling in resources during periods of low utilization. Evaluate your workload for predictable patterns and proactively scale as you anticipate predicted and planned changes in demand.

02. Align SLAs with sustainability goals	Manual	Low	Open	

Recommendation:
Define SLAs that support your sustainability goals while meeting your business requirements. Redefine SLAs to meet business requirements, not exceed them. Make trade offs that significantly reduce sustainability impacts in exchange for acceptable decreases in service levels. Use circuit breakers and design patterns that prioritize business critical functions, and allow lower service levels (such as response time or recovery time objectives) for non critical functions.

03. Stop the creation and maintenance of unused assets	Automated	Low	Open	12

Recommendation:
Manage static assets and remove assets that are no longer required. Manage generated assets and stop generating and remove assets that are no longer required. Consolidate overlapping generated assets to remove redundant processing. Instruct third parties to stop producing and storing assets managed on your behalf that are no longer required. Instruct third parties to consolidate redundant assets produced on your behalf.

04. Optimize geographic placement of workloads based on their networking requirements	Manual	Medium	Open	

Recommendation:
Use local caching for frequently used resources. Use connection pooling to enable connection reuse and reduce required resources. Use edge caching to reduce the amount of data traversing the network from your origin server. Use distributed data stores that don't rely on persistent connections and synchronous updates for consistency to serve regional populations. Replace pre provisioned static network capacity with shared dynamic capacity, and share the sustainability impact of network capacity with other subscribers.

05. Optimize team member resources for activities performed	Manual	Low	Open	

Recommendation:
Provision workstations and other devices to align with how they are used. Use virtual desktops and application streaming to limit upgrade and device requirements. Move processor or memory intensive tasks to the cloud. Evaluate the impact of processes and systems on your device lifecycle, and select solutions that minimize the requirement for device replacement while satisfying business requirements. Implement remote management for devices to reduce required business travel.

Pillar Name: Sustainability

Question: 02. How do you align cloud resources to your demand?

Question Risk: Medium

Best Practice	Nature	Severity	Status	Violated Resources
06. Implement buffering or throttling to flatten the demand curve	Manual	Low	Open	

Recommendation:
Analyze the client requests to determine how to respond to them. Questions to consider include:
Can this request be processed asynchronously?
Does the client have retry capability?
If the client has retry capability, then you can implement throttling, which tells the source that if it cannot service the request at the current time, it should try again later.
You can use Amazon API Gateway to implement throttling.
For clients that cannot perform retries, a buffer needs to be implemented to flatten the demand curve. A buffer defers request processing, allowing applications that run at different rates to communicate effectively. A buffer based approach uses a queue or a stream to accept messages from producers. Messages are read by consumers and processed, allowing the messages to run at the rate that meets the consumers' business requirements.
Amazon Simple Queue Service (Amazon SQS) is a managed service that provides queues that allow a single consumer to read individual messages.
Amazon Kinesis provides a stream that allows many consumers to read the same messages.
Analyze the overall demand, rate of change, and required response time to right size the throttle or buffer required.

Pillar Name: Sustainability

Question: 03. How do you take advantage of software and architecture patterns to support your sustainability goals?

Question Risk: Medium

Best Practice	Nature	Severity	Status	Violated Resources
01. Optimize software and architecture for asynchronous and scheduled jobs	Automated	Medium	Open	1

Recommendation:
Queue requests that don't require immediate processing. Increase serialization to flatten utilization across your pipeline. Modify the capacity of individual components to prevent idling resources waiting for input. Create buffers and establish rate limiting to smooth the consumption of external services. Use the most efficient available hardware for your software optimizations. Use queue driven architectures, pipeline management, and On Demand Instance workers to maximize utilization for batch processing. Schedule tasks to avoid load spikes and resource contention from simultaneous execution. Schedule jobs during times of day where carbon intensity for power is lowest.

02. Remove or refactor workload components with low or no use	Automated	Medium	Open	12

Recommendation:
Analyze load (using indicators such as transaction flow and API calls) on functional components to identify unused and under utilized components. Retire components that are no longer needed. Refactor under utilized components. Consolidate under utilized components with other resources to improve utilization efficiency.

03. Optimize areas of code that consume the most time or resources	Manual	Medium	Open	

Recommendation:
Monitor performance as a function of resource usage to identify components with high resource requirements per unit of work as targets for optimization. Use a code profiler to identify the areas of code that use the most time or resources as targets for optimization. Replace algorithms with more efficient versions that produce the same result. Use hardware acceleration to improve the efficiency of blocks of code with long execution times. Use the most efficient operating system and programming language for the workload. Remove unnecessary sorting and formatting. Use data transfer patterns that minimize the resources used based on how frequently the data changes and how it is consumed. For example, push state change information to a client instead of having it consume resources to poll and receive valueless are no change messages.

04. Optimize impact on devices and equipment	Manual	Medium	Open	

Pillar Name: Sustainability

Question: 03. How do you take advantage of software and architecture patterns to support your sustainability goals?

Question Risk: Medium

Best Practice	Nature	Severity	Status	Violated Resources
Recommendation: Inventory the devices your customers use. Test using managed device farms with representative sets of hardware to understand the impact of your changes, and iterate development to maximize the devices supported. Account for network bandwidth and latency when building payloads, and implement capabilities that help your applications work well on low bandwidth, high latency links. Pre process data payloads to reduce local processing requirements and limit data transfer requirements. Perform computationally intense activities server side (such as image rendering), or use application streaming to improve the user experience on older devices. Segment and paginate output, especially for interactive sessions, to manage payloads and limit local storage requirements.				
05. Use software patterns and architectures that best support data access and storage patterns	Manual	Medium	Open	
Recommendation: Analyze your data access and storage patterns. Store data files in efficient file formats such as Parquet to prevent unnecessary processing (for example, when running analytics) and to reduce the total storage provisioned. Use technologies that work natively with compressed data. Use the database engine that best supports your dominant query pattern. Manage your database indexes to ensure index designs support efficient query execution. Select network protocols that reduce the amount of network capacity consumed.				

Pillar Name: Sustainability

Question: 04. How do you take advantage of data management policies and patterns to support your sustainability goals?

Question Risk: Medium

Best Practice	Nature	Severity	Status	Violated Resources
01. Implement a data classification policy	Manual	Medium	Open	
Recommendation: Determine requirements for the distribution, retention, and deletion of your data. Use tagging on volumes and objects to record the metadata thats used to determine how its managed, including data classification. Periodically audit your environment for untagged and unclassified data, and classify and tag the data appropriately.				
02. Use technologies that support data access and storage patterns	Automated	Low	Open	1
Recommendation: Monitor your data access patterns. Migrate data to the appropriate technology based on access pattern. Migrate archival data to storage designed for that purpose.				
03. Use policies to manage the lifecycle of your datasets	Automated	Medium	Verified	
04. Use elasticity and automation to expand block storage or file system	Automated	Medium	Open	1
Recommendation: Monitor the utilization of your data volumes. Use elastic volumes and managed block data services to automate allocation of additional storage as your persistent data grows. Set target levels of utilization for your data volumes, and resize volumes outside of expected ranges. Size read only volumes to fit the data. Migrate data to object stores to avoid provisioning the excess capacity from fixed volume sizes on block storage.				
05. Remove unneeded or redundant data	Automated	Medium	Verified	
06. Use shared file systems or storage to access common data	Manual	Medium	Open	
Recommendation:				

Pillar Name: Sustainability

Question: 04. How do you take advantage of data management policies and patterns to support your sustainability goals?

Question Risk: Medium

Best Practice	Nature	Severity	Status	Violated Resources
Migrate data to shared storage when the data has multiple consumers. Fetch data from shared storage only as needed. Delete data as appropriate for your usage patterns, and implement time to live functionality to manage cached data. Detach volumes from clients that are not actively using them.				
07. Minimize data movement across networks	Manual	Medium	Open	
Recommendation: Store data as close to the consumer as possible. Partition regionally consumed services so that their Region specific data is stored within the Region where it is consumed. Use block level duplication instead of file or object level duplication when copying changes across the network. Compress data before moving it over the network.				
08. Back up data only when difficult to recreate	Manual	Medium	Open	
Recommendation: Use your data classification to establish what data needs to be backed up. Exclude data that you can easily recreate. Exclude ephemeral data from your backups. Exclude local copies of data, unless the time required to restore that data from a common location exceeds your SLAs.				

Pillar Name: Sustainability

Question: 05. How do you select and use cloud hardware and services in your architecture to support your sustainability goals?

Question Risk: Medium

Best Practice	Nature	Severity	Status	Violated Resources
01. Use the minimum amount of hardware to meet your needs	Manual	Medium	Open	
Recommendation: Enable horizontal scaling, and use automation to scale out as loads increase and to scale in as loads decrease. Scale using small increments for variable workloads. Align scaling with cyclical utilization patterns (for example, a payroll system with intense bi weekly processing activities) as load varies over days, weeks, months, or years. Negotiate SLAs that allow for a temporary reduction in capacity while automation deploys replacement resources.				
02. Use instance types with the least impact	Manual	Medium	Open	
Recommendation: Use the most efficient instance type compatible with your workload. Modify your workload to work with different numbers of CPUs and different amounts of memory to maximize your choice of instance type. Migrate your workload to Regions that offer instances with the least sustainability impact and that meet your SLAs. Use burstable instances to support workloads with infrequent requirements for additional capacity. Use Spot Instances for stateless and fault tolerant workloads to increase overall utilization of the Cloud and reduce the sustainability impact of unused resources.				
03. Use managed services	Automated	Medium	Verified	
04. Optimize your use of hardware-based compute accelerators	Manual	Medium	Open	
Recommendation: Use GPUs only for tasks where they are more efficient than CPU based alternatives. Use automation to release GPU instances when not in use. Use flexible graphics acceleration rather than dedicated GPU instances. Take advantage of custom purpose hardware that is specific to your workload.				

Pillar Name: Sustainability

Question: 06. How do your organizational processes support your sustainability goals

Question Risk: Medium

Best Practice	Nature	Severity	Status	Violated Resources
01. Adopt methods that can rapidly introduce sustainability improvements	Manual	Medium	Open	
Recommendation: Add requirements for sustainability to your development process. Enable resources to work in parallel to develop, test, and deploy sustainability improvements. Test and validate potential sustainability impact improvements before deploying into production. Test potential improvements using the minimum viable representative components. Deploy tested sustainability improvements to production as they become available.				
02. Keep your workload up-to-date	Manual	Low	Open	
Recommendation: Update systems to gain performance efficiencies. Update systems to remove barriers for a planned improvement. Update systems to acquire software features to reduce sustainability impacts. Update systems to improve your ability to measure and manage sustainability impacts.				
03. Increase utilization of build environments	Manual	Low	Open	
Recommendation: Use automation to maximize utilization of your development and test environments. Use automation to manage the lifecycle of your development and test environments. Use minimum viable representative environments to develop and test potential improvements. Use On Demand Instances to supplement your developer devices. Use automation to maximize the efficiency of your build resources. Use instance types with burst capacity, Spot Instances, and other technologies to align build capacity with use. Adopt native cloud services for secure instance shell access rather than deploying fleets of bastion hosts.				
04. Use managed device farms for testing	Manual	Low	Open	
Recommendation: Test using managed device farms with representative sets of hardware to understand the impact of your changes, and iterate development to maximize the devices supported.				

Workload Assessment Details With Resources

Pillar Name	Cost Optimization	**Status**	Violation
Question	02. How do you govern usage?	**No. of violated resources**	11
Best Practice	06. Track project lifecycle	**Resource Category**	Storage
Severity	Low	**Resource**	Volumes
Policy	AWS EBS Volume Orphaned	**Resource Type**	EBS
Policy Description	This policy identifies the EBS Volumes which has not been used to reduce unnecessary charges and save cost. The policy becomes non-compliant when there is no instance attached to the EBS Volumes.		
Policy Recommendation	**Description** EBS volumes which are not attached (State as available) are considered as Orphaned.It is recommended to address the orphaned EBS Volumes by deleting them to avoid unnecessary usage cost being incurred on your bill. In most cases, we recommend taking a EBS snapshot before deleting a EBS Volumes so that it can be restored if required. It is the responsibility of the user to decide which EBS Volumes are required and take appropriate action based on the data available on these Volumes. **Remediation Actions** • Delete EBS Volumes		

Resource Id	Cloud Account	Region
vol-0e47a859161478f52	AWS_Bagchi	us-east-1
vol-0fb9b3484ca0fbc33	AWS_Bagchi	us-east-1
vol-0045b166860804a3c	AWS_Bagchi	us-east-1
vol-076d8124573051729	AWS_Bagchi	us-east-1
vol-0f7c20f3ecb1f5fe9	AWS_Bagchi	us-east-1
vol-0ab237df0b9671462	AWS_Bagchi	us-east-1
vol-0abf5b337987a1692	AWS_Bagchi	us-east-1
vol-011c568079be4b7a4	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
vol-00482b3d8144fc181	AWS_Bagchi	us-east-1
vol-03a32df71101c03df	AWS_Bagchi	us-east-1
vol-094195b5f3842d3e4	AWS_Bagchi	us-east-1

Pillar Name	Cost Optimization	**Status**	Violation
Question	02. How do you govern usage?	**No. of violated resources**	1
Best Practice	06. Track project lifecycle	**Resource Category**	Network
Severity	Low	**Resource**	Internet_Gateways
Policy	Aws Audit Vpc Unused Internet Gateway Policy CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether any unused VPC resources is available in the AWS account. For a better management of VPC resources, all unused (detached) Internet Gateways and Egress-Only Internet Gateways is recommended to be removed from AWS VPC environment. This rule becomes non-compliant if unused VPCs are not removed in the AWS account.		
Policy Recommendation	It is recommended to Identify and remove any unused VPC Internet Gateways (IGWs) and VPC Egress-Only Internet Gateways (EIGWs) in order to adhere to best practices and to avoid approaching the service limit (by default, you are limited to 5 IGWs and 5 EIGWs per AWS region). An Internet Gateway/Egress-Only Internet Gateway is evaluated as unused when is not attached anymore to an AWS Virtual Private Cloud (VPC).		

Resource Id	Cloud Account	Region
igw-070e4c8901108b0bb	AWS_Bagchi	us-east-1

Pillar Name	Cost Optimization	**Status**	Violation
Question	04. How do you decommission resources?	**No. of violated resources**	11
Best Practice	02. Implement a decommissioning process	**Resource Category**	Storage
Severity	High	**Resource**	Volumes
Policy	AWS EBS Volume Orphaned	**Resource Type**	EBS
Policy Description	This policy identifies the EBS Volumes which has not been used to reduce unnecessary charges and save cost. The policy becomes non-compliant when there is no instance attached to the EBS Volumes.		
Policy Recommendation	**Description** EBS volumes which are not attached (State as available) are considered as Orphaned.It is recommended to address the orphaned EBS Volumes by deleting them to avoid unnecessary usage cost being incurred on your bill. In most cases, we recommend taking a EBS snapshot before deleting a EBS Volumes so that it can be restored if required. It is the responsibility of the user to decide which EBS Volumes are required and take appropriate action based on the data available on these Volumes. **Remediation Actions** • Delete EBS Volumes		

Resource Id	Cloud Account	Region
vol-0e47a859161478f52	AWS_Bagchi	us-east-1
vol-0fb9b3484ca0fbc33	AWS_Bagchi	us-east-1
vol-0045b166860804a3c	AWS_Bagchi	us-east-1
vol-076d8124573051729	AWS_Bagchi	us-east-1
vol-0f7c20f3ecb1f5fe9	AWS_Bagchi	us-east-1
vol-0ab237df0b9671462	AWS_Bagchi	us-east-1
vol-0abf5b337987a1692	AWS_Bagchi	us-east-1
vol-011c568079be4b7a4	AWS_Bagchi	us-east-1
vol-00482b3d8144fc181	AWS_Bagchi	us-east-1
vol-03a32df71101c03df	AWS_Bagchi	us-east-1
vol-094195b5f3842d3e4	AWS_Bagchi	us-east-1

Pillar Name	Cost Optimization	**Status**	Violation
Question	04. How do you decommission resources?	**No. of violated resources**	1
Best Practice	02. Implement a decommissioning process	**Resource Category**	Network
Severity	High	**Resource**	Internet_Gateways
Policy	Aws Audit Vpc Unused Internet Gateway Policy CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether any unused VPC resources is available in the AWS account. For a better management of VPC resources, all unused (detached) Internet Gateways and Egress-Only Internet Gateways is recommended to be removed from AWS VPC environment. This rule becomes non-compliant if unused VPCs are not removed in the AWS account.		
Policy Recommendation	It is recommended to Identify and remove any unused VPC Internet Gateways (IGWs) and VPC Egress-Only Internet Gateways (EIGWs) in order to adhere to best practices and to avoid approaching the service limit (by default, you are limited to 5 IGWs and 5 EIGWs per AWS region). An Internet Gateway/Egress-Only Internet Gateway is evaluated as unused when is not attached anymore to an AWS Virtual Private Cloud (VPC).		

Resource Id	Cloud Account	Region
igw-070e4c8901108b0bb	AWS_Bagchi	us-east-1

Pillar Name	Cost Optimization	**Status**	Violation
Question	04. How do you decommission resources?	**No. of violated resources**	11
Best Practice	03. Decommission resources	**Resource Category**	Storage
Severity	Medium	**Resource**	Volumes
Policy	AWS EBS Volume Orphaned	**Resource Type**	EBS
Policy Description	This policy identifies the EBS Volumes which has not been used to reduce unnecessary charges and save cost. The policy becomes non-compliant when there is no instance attached to the EBS Volumes.		
Policy Recommendation	**Description** EBS volumes which are not attached (State as available) are considered as Orphaned.It is recommended to address the orphaned EBS Volumes by deleting them to avoid unnecessary usage cost being incurred on your bill. In most cases, we recommend taking a EBS snapshot before deleting a EBS Volumes so that it can be restored if required. It is the responsibility of the user to decide which EBS Volumes are required and take appropriate action based on the data available on these Volumes. **Remediation Actions** • Delete EBS Volumes		

Resource Id	Cloud Account	Region
vol-0e47a859161478f52	AWS_Bagchi	us-east-1
vol-0fb9b3484ca0fbc33	AWS_Bagchi	us-east-1
vol-0045b166860804a3c	AWS_Bagchi	us-east-1
vol-076d8124573051729	AWS_Bagchi	us-east-1
vol-0f7c20f3ecb1f5fe9	AWS_Bagchi	us-east-1
vol-0ab237df0b9671462	AWS_Bagchi	us-east-1
vol-0abf5b337987a1692	AWS_Bagchi	us-east-1
vol-011c568079be4b7a4	AWS_Bagchi	us-east-1
vol-00482b3d8144fc181	AWS_Bagchi	us-east-1
vol-03a32df71101c03df	AWS_Bagchi	us-east-1
vol-094195b5f3842d3e4	AWS_Bagchi	us-east-1

Pillar Name	Cost Optimization	**Status**	Violation
Question	04. How do you decommission resources?	**No. of violated resources**	1
Best Practice	03. Decommission resources	**Resource Category**	Network
Severity	Medium	**Resource**	Internet_Gateways
Policy	Aws Audit Vpc Unused Internet Gateway Policy CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether any unused VPC resources is available in the AWS account. For a better management of VPC resources, all unused (detached) Internet Gateways and Egress-Only Internet Gateways is recommended to be removed from AWS VPC environment. This rule becomes non-compliant if unused VPCs are not removed in the AWS account.		
Policy Recommendation	It is recommended to Identify and remove any unused VPC Internet Gateways (IGWs) and VPC Egress-Only Internet Gateways (EIGWs) in order to adhere to best practices and to avoid approaching the service limit (by default, you are limited to 5 IGWs and 5 EIGWs per AWS region). An Internet Gateway/Egress-Only Internet Gateway is evaluated as unused when is not attached anymore to an AWS Virtual Private Cloud (VPC).		

Resource Id	Cloud Account	Region
igw-070e4c8901108b0bb	AWS_Bagchi	us-east-1

Pillar Name	Cost Optimization	**Status**	Violation
Question	04. How do you decommission resources?	**No. of violated resources**	11
Best Practice	04. Decommission resources automatically	**Resource Category**	Storage
Severity	Low	**Resource**	Volumes
Policy	AWS EBS Volume Orphaned	**Resource Type**	EBS
Policy Description	This policy identifies the EBS Volumes which has not been used to reduce unnecessary charges and save cost. The policy becomes non-compliant when there is no instance attached to the EBS Volumes.		
Policy Recommendation	**Description** EBS volumes which are not attached (State as available) are considered as Orphaned.It is recommended to address the orphaned EBS Volumes by deleting them to avoid unnecessary usage cost being incurred on your bill. In most cases, we recommend taking a EBS snapshot before deleting a EBS Volumes so that it can be restored if required. It is the responsibility of the user to decide which EBS Volumes are required and take appropriate action based on the data available on these Volumes. **Remediation Actions** • Delete EBS Volumes		

Resource Id	Cloud Account	Region
vol-0e47a859161478f52	AWS_Bagchi	us-east-1
vol-0fb9b3484ca0fbc33	AWS_Bagchi	us-east-1
vol-0045b166860804a3c	AWS_Bagchi	us-east-1
vol-076d8124573051729	AWS_Bagchi	us-east-1
vol-0f7c20f3ecb1f5fe9	AWS_Bagchi	us-east-1
vol-0ab237df0b9671462	AWS_Bagchi	us-east-1
vol-0abf5b337987a1692	AWS_Bagchi	us-east-1
vol-011c568079be4b7a4	AWS_Bagchi	us-east-1
vol-00482b3d8144fc181	AWS_Bagchi	us-east-1
vol-03a32df71101c03df	AWS_Bagchi	us-east-1
vol-094195b5f3842d3e4	AWS_Bagchi	us-east-1

Pillar Name	Cost Optimization	**Status**	Violation
Question	04. How do you decommission resources?	**No. of violated resources**	1
Best Practice	04. Decommission resources automatically	**Resource Category**	Network
Severity	Low	**Resource**	Internet_Gateways
Policy	Aws Audit Vpc Unused Internet Gateway Policy CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether any unused VPC resources is available in the AWS account. For a better management of VPC resources, all unused (detached) Internet Gateways and Egress-Only Internet Gateways is recommended to be removed from AWS VPC environment. This rule becomes non-compliant if unused VPCs are not removed in the AWS account.		
Policy Recommendation	It is recommended to Identify and remove any unused VPC Internet Gateways (IGWs) and VPC Egress-Only Internet Gateways (EIGWs) in order to adhere to best practices and to avoid approaching the service limit (by default, you are limited to 5 IGWs and 5 EIGWs per AWS region). An Internet Gateway/Egress-Only Internet Gateway is evaluated as unused when is not attached anymore to an AWS Virtual Private Cloud (VPC).		

Resource Id	Cloud Account	Region
igw-070e4c8901108b0bb	AWS_Bagchi	us-east-1

Pillar Name	Cost Optimization	**Status**	Violation
Question	05. How do you evaluate cost when you select services?	**No. of violated resources**	11
Best Practice	02. Analyze all components of the workload	**Resource Category**	Storage
Severity	High	**Resource**	Volumes
Policy	AWS EBS Volume Orphaned	**Resource Type**	EBS
Policy Description	This policy identifies the EBS Volumes which has not been used to reduce unnecessary charges and save cost. The policy becomes non-compliant when there is no instance attached to the EBS Volumes.		
Policy Recommendation	**Description** EBS volumes which are not attached (State as available) are considered as Orphaned.It is recommended to address the orphaned EBS Volumes by deleting them to avoid unnecessary usage cost being incurred on your bill. In most cases, we recommend taking a EBS snapshot before deleting a EBS Volumes so that it can be restored if required. It is the responsibility of the user to decide which EBS Volumes are required and take appropriate action based on the data available on these Volumes. **Remediation Actions** • Delete EBS Volumes		

Resource Id	Cloud Account	Region
vol-0e47a859161478f52	AWS_Bagchi	us-east-1
vol-0fb9b3484ca0fbc33	AWS_Bagchi	us-east-1
vol-0045b166860804a3c	AWS_Bagchi	us-east-1
vol-076d8124573051729	AWS_Bagchi	us-east-1
vol-0f7c20f3ecb1f5fe9	AWS_Bagchi	us-east-1
vol-0ab237df0b9671462	AWS_Bagchi	us-east-1
vol-0abf5b337987a1692	AWS_Bagchi	us-east-1
vol-011c568079be4b7a4	AWS_Bagchi	us-east-1
vol-00482b3d8144fc181	AWS_Bagchi	us-east-1
vol-03a32df71101c03df	AWS_Bagchi	us-east-1
vol-094195b5f3842d3e4	AWS_Bagchi	us-east-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	02. How do you structure your organization to support your business outcomes?	**No. of violated resources**	1599
Best Practice	01. Resources have identified owners	**Resource Category**	Compute
Severity	High	**Resource**	Key Pairs
Policy	AWS Audit Required Owner Tags CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether your resources have the owner tags specified. The tag structure could differ from organization to organization. However a minimum set of required tags needs to be maintained for tracking the usage of each of the resources within the AWS infrastructure. This rule will be non-compliant if the owner tags and values do not exist for each of the resources within the AWS account.		
Policy Recommendation	Resources that do not have specified Owner tags are recommended to define Owners at resource level for better visibility of resource usage. Tags are mandatory to keep track of the costs and define budgets.		

Note : The below table displays only 100 resource violations out of 1599 resource violations. To view the complete resource violations, please go to the policy violations under for the current assessment in the platform UI

Resource Id	Cloud Account	Region
key-002c7129dbb643e8c	AWS_Bagchi	us-east-1
vol-0a7bcd7fe14ab9958	AWS_Bagchi	us-east-1
vol-0321fb4cab10dd421	AWS_Bagchi	us-east-1
key-0dc2196c5fac7434b	AWS_Bagchi	us-east-1
key-079fb09f81a049d50	AWS_Bagchi	us-east-1
vol-094176eec3c11c16e	AWS_Bagchi	us-east-1
eni-040c8268dadb62a8b	AWS_Bagchi	us-east-1
vol-088d600517109f4a3	AWS_Bagchi	us-east-1
key-06f71271a1b8bc897	AWS_Bagchi	us-east-1
vol-0e47a859161478f52	AWS_Bagchi	us-east-1
key-09e17afe882917766	AWS_Bagchi	us-east-1
eni-0a08698cc8eb9e20e	AWS_Bagchi	us-east-1
key-01478eb308c256a35	AWS_Bagchi	us-east-1
vol-0fb9b3484ca0fbc33	AWS_Bagchi	us-east-1
eni-0f265e8799ee906d0	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
eni-00a792f323fb0e2f7	AWS_Bagchi	us-east-1
eni-0d7f174eabd45ef58	AWS_Bagchi	us-east-1
vol-0a0b7bf7687dc55f2	AWS_Bagchi	us-east-1
vol-005658743b6efb900	AWS_Bagchi	us-east-1
eni-0241fa4013e32d470	AWS_Bagchi	us-east-1
vol-0611b706e1adddc02	AWS_Bagchi	us-east-1
eni-0fc2d9f862aadf65d	AWS_Bagchi	us-east-1
eni-0dd1745847ecda4d8	AWS_Bagchi	us-east-1
eni-075e208940bdd775b	AWS_Bagchi	us-east-1
eni-09787aebf5dfaebd5	AWS_Bagchi	us-east-1
vol-08cc085215cb7212c	AWS_Bagchi	us-east-1
vol-07c6661b47a956bf7	AWS_Bagchi	us-east-1
eni-0d5df1812d124ddc9	AWS_Bagchi	us-east-1
vol-0ab237df0b9671462	AWS_Bagchi	us-east-1
eni-076d2153a53d345e0	AWS_Bagchi	us-east-1
vol-0abf5b337987a1692	AWS_Bagchi	us-east-1
eni-0a09afc3790e57e00	AWS_Bagchi	us-east-1
eni-074fea6e1482aa926	AWS_Bagchi	us-east-1
vol-011c568079be4b7a4	AWS_Bagchi	us-east-1
eni-058c6d0af148baee4	AWS_Bagchi	us-east-1
vol-00482b3d8144fc181	AWS_Bagchi	us-east-1
eni-0cd5c667f238e3d7f	AWS_Bagchi	us-east-1
vol-03a32df71101c03df	AWS_Bagchi	us-east-1
eni-0d1f9fbb5c83cd9a3	AWS_Bagchi	us-east-1
vol-094195b5f3842d3e4	AWS_Bagchi	us-east-1
snap-0b1e854018b33bcfd	AWS_Bagchi	us-east-1
snap-0d91cc78a32973f56	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
snap-03a0557bc4e0119e2	AWS_Bagchi	us-east-1
subnet-0a26c4cf6183911d5	AWS_Bagchi	us-east-1
subnet-09ea0a37274c42302	AWS_Bagchi	us-east-1
subnet-95e017d9	AWS_Bagchi	us-east-1
subnet-acc2e08d	AWS_Bagchi	us-east-1
subnet-573c1c08	AWS_Bagchi	us-east-1
subnet-410f6470	AWS_Bagchi	us-east-1
subnet-0d45555af01e7748c	AWS_Bagchi	us-east-1
subnet-599bbf3f	AWS_Bagchi	us-east-1
subnet-02aed858349185ef7	AWS_Bagchi	us-east-1
subnet-c0263fce	AWS_Bagchi	us-east-1
subnet-2207fa49	AWS_Bagchi	ap-south-1
subnet-fd661286	AWS_Bagchi	ap-south-1
subnet-e6fbd9aa	AWS_Bagchi	ap-south-1
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1
rtb-0b27ea4933a478a68	AWS_Bagchi	us-east-1
rtb-029b9474d9236c4b9	AWS_Bagchi	us-east-1
rtb-0d209864fc44d0b3f	AWS_Bagchi	us-east-1
rtb-03078cf8a619ce59c	AWS_Bagchi	us-east-1
rtb-0b54935567c7c737f	AWS_Bagchi	us-east-1
rtb-37679846	AWS_Bagchi	us-east-1
rtb-0a1fc464bc5d26b79	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
rtb-0b601f59648f1f149	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
default	AWS_Bagchi	us-east-1
rtb-dfbd28b4	AWS_Bagchi	ap-south-1
default-vpc-f6ed9d8b	AWS_Bagchi	us-east-1
igw-01bd7078297c638e5	AWS_Bagchi	us-east-1
igw-03ae7fe010cfc1d4e	AWS_Bagchi	us-east-1
igw-044f286afa839dc73	AWS_Bagchi	us-east-1
igw-070e4c8901108b0bb	AWS_Bagchi	us-east-1
igw-1097306a	AWS_Bagchi	us-east-1
sg-0bec9689805cc2af4	AWS_Bagchi	us-east-1
sg-06cf8f6c379f72972	AWS_Bagchi	us-east-1
sg-00c6642262b469e78	AWS_Bagchi	us-east-1
sg-0ae0afc704d6f4a42	AWS_Bagchi	us-east-1
sg-0a92d97292ef1e035	AWS_Bagchi	us-east-1
sg-0da9b07b3d85448fb	AWS_Bagchi	us-east-1
sg-044cebb2f306fa722	AWS_Bagchi	us-east-1
sg-02462c0c85e68a867	AWS_Bagchi	us-east-1
sg-07dc3a3d238a30dcd	AWS_Bagchi	us-east-1
sg-0932dceb37352180b	AWS_Bagchi	us-east-1
sg-634fdb7d	AWS_Bagchi	us-east-1
sg-0dcb42941edc7dc45	AWS_Bagchi	us-east-1
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1
sg-0009bad022b207bb4	AWS_Bagchi	us-east-1
sg-08930526ca35700f6	AWS_Bagchi	us-east-1
sg-08a870c6344de8357	AWS_Bagchi	us-east-1
sg-0f2117e9cc6bc46cc	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
sg-0ca41ccae1d763064	AWS_Bagchi	us-east-1
sg-0a0fe7fa081d4dd46	AWS_Bagchi	us-east-1
igw-321e7f5a	AWS_Bagchi	ap-south-1
default.mysql8.0	AWS_Bagchi	us-east-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	02. How do you structure your organization to support your business outcomes?	**No. of violated resources**	1599
Best Practice	03. Operations activities have identified owners responsible for their performance	**Resource Category**	Compute
Severity	High	**Resource**	Key Pairs
Policy	AWS Audit Required Owner Tags CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether your resources have the owner tags specified. The tag structure could differ from organization to organization. However a minimum set of required tags needs to be maintained for tracking the usage of each of the resources within the AWS infrastructure. This rule will be non-compliant if the owner tags and values do not exist for each of the resources within the AWS account.		
Policy Recommendation	Resources that do not have specified Owner tags are recommended to define Owners at resource level for better visibility of resource usage. Tags are mandatory to keep track of the costs and define budgets.		

Note : The below table displays only 100 resource violations out of 1599 resource violations. To view the complete resource violations, please go to the policy violations under for the current assessment in the platform UI

Resource Id	Cloud Account	Region
key-002c7129dbb643e8c	AWS_Bagchi	us-east-1
vol-0a7bcd7fe14ab9958	AWS_Bagchi	us-east-1
vol-0321fb4cab10dd421	AWS_Bagchi	us-east-1
key-0dc2196c5fac7434b	AWS_Bagchi	us-east-1
key-079fb09f81a049d50	AWS_Bagchi	us-east-1
vol-094176eec3c11c16e	AWS_Bagchi	us-east-1
eni-040c8268dadb62a8b	AWS_Bagchi	us-east-1
vol-088d600517109f4a3	AWS_Bagchi	us-east-1
key-06f71271a1b8bc897	AWS_Bagchi	us-east-1
vol-0e47a859161478f52	AWS_Bagchi	us-east-1
key-09e17afe882917766	AWS_Bagchi	us-east-1
eni-0a08698cc8eb9e20e	AWS_Bagchi	us-east-1
key-01478eb308c256a35	AWS_Bagchi	us-east-1
vol-0fb9b3484ca0fbc33	AWS_Bagchi	us-east-1
eni-0f265e8799ee906d0	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
eni-00a792f323fb0e2f7	AWS_Bagchi	us-east-1
eni-0d7f174eabd45ef58	AWS_Bagchi	us-east-1
vol-0a0b7bf7687dc55f2	AWS_Bagchi	us-east-1
vol-005658743b6efb900	AWS_Bagchi	us-east-1
eni-0241fa4013e32d470	AWS_Bagchi	us-east-1
vol-0611b706e1adddc02	AWS_Bagchi	us-east-1
eni-0fc2d9f862aadf65d	AWS_Bagchi	us-east-1
eni-0dd1745847ecda4d8	AWS_Bagchi	us-east-1
eni-075e208940bdd775b	AWS_Bagchi	us-east-1
eni-09787aebf5dfaebd5	AWS_Bagchi	us-east-1
vol-08cc085215cb7212c	AWS_Bagchi	us-east-1
vol-07c6661b47a956bf7	AWS_Bagchi	us-east-1
eni-0d5df1812d124ddc9	AWS_Bagchi	us-east-1
vol-0ab237df0b9671462	AWS_Bagchi	us-east-1
eni-076d2153a53d345e0	AWS_Bagchi	us-east-1
vol-0abf5b337987a1692	AWS_Bagchi	us-east-1
eni-0a09afc3790e57e00	AWS_Bagchi	us-east-1
eni-074fea6e1482aa926	AWS_Bagchi	us-east-1
vol-011c568079be4b7a4	AWS_Bagchi	us-east-1
eni-058c6d0af148baee4	AWS_Bagchi	us-east-1
vol-00482b3d8144fc181	AWS_Bagchi	us-east-1
eni-0cd5c667f238e3d7f	AWS_Bagchi	us-east-1
vol-03a32df71101c03df	AWS_Bagchi	us-east-1
eni-0d1f9fbb5c83cd9a3	AWS_Bagchi	us-east-1
vol-094195b5f3842d3e4	AWS_Bagchi	us-east-1
snap-0b1e854018b33bcfd	AWS_Bagchi	us-east-1
snap-0d91cc78a32973f56	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
snap-03a0557bc4e0119e2	AWS_Bagchi	us-east-1
subnet-0a26c4cf6183911d5	AWS_Bagchi	us-east-1
subnet-09ea0a37274c42302	AWS_Bagchi	us-east-1
subnet-95e017d9	AWS_Bagchi	us-east-1
subnet-acc2e08d	AWS_Bagchi	us-east-1
subnet-573c1c08	AWS_Bagchi	us-east-1
subnet-410f6470	AWS_Bagchi	us-east-1
subnet-0d45555af01e7748c	AWS_Bagchi	us-east-1
subnet-599bbf3f	AWS_Bagchi	us-east-1
subnet-02aed858349185ef7	AWS_Bagchi	us-east-1
subnet-c0263fce	AWS_Bagchi	us-east-1
subnet-2207fa49	AWS_Bagchi	ap-south-1
subnet-fd661286	AWS_Bagchi	ap-south-1
subnet-e6fbd9aa	AWS_Bagchi	ap-south-1
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1
rtb-0b27ea4933a478a68	AWS_Bagchi	us-east-1
rtb-029b9474d9236c4b9	AWS_Bagchi	us-east-1
rtb-0d209864fc44d0b3f	AWS_Bagchi	us-east-1
rtb-03078cf8a619ce59c	AWS_Bagchi	us-east-1
rtb-0b54935567c7c737f	AWS_Bagchi	us-east-1
rtb-37679846	AWS_Bagchi	us-east-1
rtb-0a1fc464bc5d26b79	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
rtb-0b601f59648f1f149	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
default	AWS_Bagchi	us-east-1
rtb-dfbd28b4	AWS_Bagchi	ap-south-1
default-vpc-f6ed9d8b	AWS_Bagchi	us-east-1
igw-01bd7078297c638e5	AWS_Bagchi	us-east-1
igw-03ae7fe010cfc1d4e	AWS_Bagchi	us-east-1
igw-044f286afa839dc73	AWS_Bagchi	us-east-1
igw-070e4c8901108b0bb	AWS_Bagchi	us-east-1
igw-1097306a	AWS_Bagchi	us-east-1
sg-0bec9689805cc2af4	AWS_Bagchi	us-east-1
sg-06cf8f6c379f72972	AWS_Bagchi	us-east-1
sg-00c6642262b469e78	AWS_Bagchi	us-east-1
sg-0ae0afc704d6f4a42	AWS_Bagchi	us-east-1
sg-0a92d97292ef1e035	AWS_Bagchi	us-east-1
sg-0da9b07b3d85448fb	AWS_Bagchi	us-east-1
sg-044cebb2f306fa722	AWS_Bagchi	us-east-1
sg-02462c0c85e68a867	AWS_Bagchi	us-east-1
sg-07dc3a3d238a30dcd	AWS_Bagchi	us-east-1
sg-0932dceb37352180b	AWS_Bagchi	us-east-1
sg-634fdb7d	AWS_Bagchi	us-east-1
sg-0dcb42941edc7dc45	AWS_Bagchi	us-east-1
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1
sg-0009bad022b207bb4	AWS_Bagchi	us-east-1
sg-08930526ca35700f6	AWS_Bagchi	us-east-1
sg-08a870c6344de8357	AWS_Bagchi	us-east-1
sg-0f2117e9cc6bc46cc	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
sg-0ca41ccae1d763064	AWS_Bagchi	us-east-1
sg-0a0fe7fa081d4dd46	AWS_Bagchi	us-east-1
igw-321e7f5a	AWS_Bagchi	ap-south-1
default.mysql8.0	AWS_Bagchi	us-east-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	02. How do you structure your organization to support your business outcomes?	**No. of violated resources**	1
Best Practice	04. Team members know what they are responsible for	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit EC2 Instance Using IAM Roles CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether IAM roles are used to grant permission to access AWS resource from EC2 instances. AWS IAM roles reduce the risks associated with sharing and rotating credentials that can be used outside of AWS itself. If credentials are compromised, they can be used from outside of the AWS account. This rule is non-compliant if IAM roles are not used to access AWS resources by the applications hosted in EC2 instances.		
Policy Recommendation	It is recommended to use Instance Profiles/IAM Roles to appropriately grant permissions to applications running on amazon EC2 instances. AWS IAM roles reduce the risks associated with sharing and rotating credentials that can be used outside of AWS itself. IAM roles can only be associated at the launch of an instance.		

Resource Id	Cloud Account	Region
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	02. How do you structure your organization to support your business outcomes?	**No. of violated resources**	1599
Best Practice	04. Team members know what they are responsible for	**Resource Category**	Compute
Severity	High	**Resource**	Key Pairs
Policy	AWS Audit Required Owner Tags CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether your resources have the owner tags specified. The tag structure could differ from organization to organization. However a minimum set of required tags needs to be maintained for tracking the usage of each of the resources within the AWS infrastructure. This rule will be non-compliant if the owner tags and values do not exist for each of the resources within the AWS account.		
Policy Recommendation	Resources that do not have specified Owner tags are recommended to define Owners at resource level for better visibility of resource usage. Tags are mandatory to keep track of the costs and define budgets.		

Note : The below table displays only 100 resource violations out of 1599 resource violations. To view the complete resource violations, please go to the policy violations under for the current assessment in the platform UI

Resource Id	Cloud Account	Region
key-002c7129dbb643e8c	AWS_Bagchi	us-east-1
vol-0a7bcd7fe14ab9958	AWS_Bagchi	us-east-1
vol-0321fb4cab10dd421	AWS_Bagchi	us-east-1
key-0dc2196c5fac7434b	AWS_Bagchi	us-east-1
key-079fb09f81a049d50	AWS_Bagchi	us-east-1
vol-094176eec3c11c16e	AWS_Bagchi	us-east-1
eni-040c8268dadb62a8b	AWS_Bagchi	us-east-1
vol-088d600517109f4a3	AWS_Bagchi	us-east-1
key-06f71271a1b8bc897	AWS_Bagchi	us-east-1
vol-0e47a859161478f52	AWS_Bagchi	us-east-1
key-09e17afe882917766	AWS_Bagchi	us-east-1
eni-0a08698cc8eb9e20e	AWS_Bagchi	us-east-1
key-01478eb308c256a35	AWS_Bagchi	us-east-1
vol-0fb9b3484ca0fbc33	AWS_Bagchi	us-east-1
eni-0f265e8799ee906d0	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
eni-00a792f323fb0e2f7	AWS_Bagchi	us-east-1
eni-0d7f174eabd45ef58	AWS_Bagchi	us-east-1
vol-0a0b7bf7687dc55f2	AWS_Bagchi	us-east-1
vol-005658743b6efb900	AWS_Bagchi	us-east-1
eni-0241fa4013e32d470	AWS_Bagchi	us-east-1
vol-0611b706e1adddc02	AWS_Bagchi	us-east-1
eni-0fc2d9f862aadf65d	AWS_Bagchi	us-east-1
eni-0dd1745847ecda4d8	AWS_Bagchi	us-east-1
eni-075e208940bdd775b	AWS_Bagchi	us-east-1
eni-09787aebf5dfaebd5	AWS_Bagchi	us-east-1
vol-08cc085215cb7212c	AWS_Bagchi	us-east-1
vol-07c6661b47a956bf7	AWS_Bagchi	us-east-1
eni-0d5df1812d124ddc9	AWS_Bagchi	us-east-1
vol-0ab237df0b9671462	AWS_Bagchi	us-east-1
eni-076d2153a53d345e0	AWS_Bagchi	us-east-1
vol-0abf5b337987a1692	AWS_Bagchi	us-east-1
eni-0a09afc3790e57e00	AWS_Bagchi	us-east-1
eni-074fea6e1482aa926	AWS_Bagchi	us-east-1
vol-011c568079be4b7a4	AWS_Bagchi	us-east-1
eni-058c6d0af148baee4	AWS_Bagchi	us-east-1
vol-00482b3d8144fc181	AWS_Bagchi	us-east-1
eni-0cd5c667f238e3d7f	AWS_Bagchi	us-east-1
vol-03a32df71101c03df	AWS_Bagchi	us-east-1
eni-0d1f9fbb5c83cd9a3	AWS_Bagchi	us-east-1
vol-094195b5f3842d3e4	AWS_Bagchi	us-east-1
snap-0b1e854018b33bcfd	AWS_Bagchi	us-east-1
snap-0d91cc78a32973f56	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
snap-03a0557bc4e0119e2	AWS_Bagchi	us-east-1
subnet-0a26c4cf6183911d5	AWS_Bagchi	us-east-1
subnet-09ea0a37274c42302	AWS_Bagchi	us-east-1
subnet-95e017d9	AWS_Bagchi	us-east-1
subnet-acc2e08d	AWS_Bagchi	us-east-1
subnet-573c1c08	AWS_Bagchi	us-east-1
subnet-410f6470	AWS_Bagchi	us-east-1
subnet-0d45555af01e7748c	AWS_Bagchi	us-east-1
subnet-599bbf3f	AWS_Bagchi	us-east-1
subnet-02aed858349185ef7	AWS_Bagchi	us-east-1
subnet-c0263fce	AWS_Bagchi	us-east-1
subnet-2207fa49	AWS_Bagchi	ap-south-1
subnet-fd661286	AWS_Bagchi	ap-south-1
subnet-e6fbd9aa	AWS_Bagchi	ap-south-1
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1
rtb-0b27ea4933a478a68	AWS_Bagchi	us-east-1
rtb-029b9474d9236c4b9	AWS_Bagchi	us-east-1
rtb-0d209864fc44d0b3f	AWS_Bagchi	us-east-1
rtb-03078cf8a619ce59c	AWS_Bagchi	us-east-1
rtb-0b54935567c7c737f	AWS_Bagchi	us-east-1
rtb-37679846	AWS_Bagchi	us-east-1
rtb-0a1fc464bc5d26b79	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
rtb-0b601f59648f1f149	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
default	AWS_Bagchi	us-east-1
rtb-dfbd28b4	AWS_Bagchi	ap-south-1
default-vpc-f6ed9d8b	AWS_Bagchi	us-east-1
igw-01bd7078297c638e5	AWS_Bagchi	us-east-1
igw-03ae7fe010cfc1d4e	AWS_Bagchi	us-east-1
igw-044f286afa839dc73	AWS_Bagchi	us-east-1
igw-070e4c8901108b0bb	AWS_Bagchi	us-east-1
igw-1097306a	AWS_Bagchi	us-east-1
sg-0bec9689805cc2af4	AWS_Bagchi	us-east-1
sg-06cf8f6c379f72972	AWS_Bagchi	us-east-1
sg-00c6642262b469e78	AWS_Bagchi	us-east-1
sg-0ae0afc704d6f4a42	AWS_Bagchi	us-east-1
sg-0a92d97292ef1e035	AWS_Bagchi	us-east-1
sg-0da9b07b3d85448fb	AWS_Bagchi	us-east-1
sg-044cebb2f306fa722	AWS_Bagchi	us-east-1
sg-02462c0c85e68a867	AWS_Bagchi	us-east-1
sg-07dc3a3d238a30dcd	AWS_Bagchi	us-east-1
sg-0932dceb37352180b	AWS_Bagchi	us-east-1
sg-634fdb7d	AWS_Bagchi	us-east-1
sg-0dcb42941edc7dc45	AWS_Bagchi	us-east-1
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1
sg-0009bad022b207bb4	AWS_Bagchi	us-east-1
sg-08930526ca35700f6	AWS_Bagchi	us-east-1
sg-08a870c6344de8357	AWS_Bagchi	us-east-1
sg-0f2117e9cc6bc46cc	AWS_Bagchi	us-east-1

Resource Id	Cloud Account	Region
sg-0ca41ccae1d763064	AWS_Bagchi	us-east-1
sg-0a0fe7fa081d4dd46	AWS_Bagchi	us-east-1
igw-321e7f5a	AWS_Bagchi	ap-south-1
default.mysql8.0	AWS_Bagchi	us-east-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	02. How do you structure your organization to support your business outcomes?	**No. of violated resources**	1
Best Practice	05. Mechanisms exist to identify responsibility and ownership	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit EC2 Instance Using IAM Roles CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether IAM roles are used to grant permission to access AWS resource from EC2 instances. AWS IAM roles reduce the risks associated with sharing and rotating credentials that can be used outside of AWS itself. If credentials are compromised, they can be used from outside of the AWS account. This rule is non-compliant if IAM roles are not used to access AWS resources by the applications hosted in EC2 instances.		
Policy Recommendation	It is recommended to use Instance Profiles/IAM Roles to appropriately grant permissions to applications running on amazon EC2 instances. AWS IAM roles reduce the risks associated with sharing and rotating credentials that can be used outside of AWS itself. IAM roles can only be associated at the launch of an instance.		

Resource Id	Cloud Account	Region
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	04. How do you design your workload so that you can understand its state?	**No. of violated resources**	7
Best Practice	01. Implement application telemetry	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit EC2 Instance Detailed Monitoring Enabled CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether detailed monitoring is enabled for the EC2 instances. Enabling detailed monitoring allows the capture of metric data every minute, providing valuable insights for effective decision-making in architecting and managing compute resources within the account. This rule is Non-compliant when the EC2 instances are not configured with detailed monitoring.		
Policy Recommendation	It is recommended to enable detailed monitoring on the EC2 instances in order to have more granular control over the utilization and help track the performance of the instance. It also helps to get aggregated data across groups of similar instances.		

Resource Id	Cloud Account	Region
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	04. How do you design your workload so that you can understand its state?	**No. of violated resources**	6
Best Practice	02. Implement and configure workload telemetry	**Resource Category**	Network
Severity	High	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	04. How do you design your workload so that you can understand its state?	**No. of violated resources**	28
Best Practice	02. Implement and configure workload telemetry	**Resource Category**	Storage
Severity	High	**Resource**	Buckets
Policy	AWS Audit S3 Bucket Logging Enabled CS Policy	**Resource Type**	S3
Policy Description	This policy audits whether Server access logging is enabled for Amazon S3 buckets. Server access logging provides detailed records for the requests that are made to an Amazon S3 bucket, which might be useful in security and access audits. This rule becomes non-compliant if Server access logging is not enabled for Amazon S3 buckets.		
Policy Recommendation	It is recommended to enable server access logging for S3 buckets. Server access logging provides detailed records for the requests that are made to a bucket. This helps in access tracking and trend analysis.		

Resource Id	Cloud Account	Region
abagchi-cloudtrail-bucket-us-east-1	AWS_Bagchi	global
bagchi-billing	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-1	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-2	AWS_Bagchi	global
bagchi-guardduty-ca-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-north-1	AWS_Bagchi	global
bagchi-guardduty-eu-west-2	AWS_Bagchi	global
bagchi-guardduty-eu-west-3	AWS_Bagchi	global
bagchi-guardduty-us-east-1	AWS_Bagchi	global
bagchi-guardduty-us-east-2	AWS_Bagchi	global
bagchi-guardduty-us-west-1	AWS_Bagchi	global
bagchi-guardduty-us-west-2	AWS_Bagchi	global
bagchi-guardduty2-ap-south-1	AWS_Bagchi	global
bagchi-guardduty2-eu-west-2	AWS_Bagchi	global
bagchi-guardduty2-us-east-1	AWS_Bagchi	global
bagchi-guardduty2-us-east-2	AWS_Bagchi	global

Resource Id	Cloud Account	Region
bagchi-guardduty2-us-west-1	AWS_Bagchi	global
bagchi-guardduty2-us-west-2	AWS_Bagchi	global
cf-templates-fg4y3e089bnm-us-east-1	AWS_Bagchi	global
cloud-trail-sales2-ap-south-1	AWS_Bagchi	global
cloud-trail-sales2-eu-west-2	AWS_Bagchi	global
cloud-trail-sales2-us-east-1	AWS_Bagchi	global
config-bucket-967852078706	AWS_Bagchi	global
demoess	AWS_Bagchi	global
feb14test001	AWS_Bagchi	global
bucket05061234	AWS_Bagchi	global
bagchi-billing2	AWS_Bagchi	global

Pillar Name	Operational Excellence	**Status**	Violation
Question	04. How do you design your workload so that you can understand its state?	**No. of violated resources**	7
Best Practice	02. Implement and configure workload telemetry	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit EC2 Instance Detailed Monitoring Enabled CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether detailed monitoring is enabled for the EC2 instances. Enabling detailed monitoring allows the capture of metric data every minute, providing valuable insights for effective decision-making in architecting and managing compute resources within the account. This rule is Non-compliant when the EC2 instances are not configured with detailed monitoring.		
Policy Recommendation	It is recommended to enable detailed monitoring on the EC2 instances in order to have more granular control over the utilization and help track the performance of the instance. It also helps to get aggregated data across groups of similar instances.		

Resource Id	Cloud Account	Region
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	04. How do you design your workload so that you can understand its state?	**No. of violated resources**	7
Best Practice	05. Implement transaction traceability	**Resource Category**	Compute
Severity	Low	**Resource**	Instances
Policy	AWS Audit EC2 Instance Detailed Monitoring Enabled CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether detailed monitoring is enabled for the EC2 instances. Enabling detailed monitoring allows the capture of metric data every minute, providing valuable insights for effective decision-making in architecting and managing compute resources within the account. This rule is Non-compliant when the EC2 instances are not configured with detailed monitoring.		
Policy Recommendation	It is recommended to enable detailed monitoring on the EC2 instances in order to have more granular control over the utilization and help track the performance of the instance. It also helps to get aggregated data across groups of similar instances.		

Resource Id	Cloud Account	Region
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	08. How do you understand the health of your workload?	**No. of violated resources**	6
Best Practice	03. Collect and analyze workload metrics	**Resource Category**	Network
Severity	High	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	08. How do you understand the health of your workload?	**No. of violated resources**	28
Best Practice	03. Collect and analyze workload metrics	**Resource Category**	Storage
Severity	High	**Resource**	Buckets
Policy	AWS Audit S3 Bucket Logging Enabled CS Policy	**Resource Type**	S3
Policy Description	This policy audits whether Server access logging is enabled for Amazon S3 buckets. Server access logging provides detailed records for the requests that are made to an Amazon S3 bucket, which might be useful in security and access audits. This rule becomes non-compliant if Server access logging is not enabled for Amazon S3 buckets.		
Policy Recommendation	It is recommended to enable server access logging for S3 buckets. Server access logging provides detailed records for the requests that are made to a bucket. This helps in access tracking and trend analysis.		

Resource Id	Cloud Account	Region
abagchi-cloudtrail-bucket-us-east-1	AWS_Bagchi	global
bagchi-billing	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-1	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-2	AWS_Bagchi	global
bagchi-guardduty-ca-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-north-1	AWS_Bagchi	global
bagchi-guardduty-eu-west-2	AWS_Bagchi	global
bagchi-guardduty-eu-west-3	AWS_Bagchi	global
bagchi-guardduty-us-east-1	AWS_Bagchi	global
bagchi-guardduty-us-east-2	AWS_Bagchi	global
bagchi-guardduty-us-west-1	AWS_Bagchi	global
bagchi-guardduty-us-west-2	AWS_Bagchi	global
bagchi-guardduty2-ap-south-1	AWS_Bagchi	global
bagchi-guardduty2-eu-west-2	AWS_Bagchi	global
bagchi-guardduty2-us-east-1	AWS_Bagchi	global
bagchi-guardduty2-us-east-2	AWS_Bagchi	global

Resource Id	Cloud Account	Region
bagchi-guardduty2-us-west-1	AWS_Bagchi	global
bagchi-guardduty2-us-west-2	AWS_Bagchi	global
cf-templates-fg4y3e089bnm-us-east-1	AWS_Bagchi	global
cloud-trail-sales2-ap-south-1	AWS_Bagchi	global
cloud-trail-sales2-eu-west-2	AWS_Bagchi	global
cloud-trail-sales2-us-east-1	AWS_Bagchi	global
config-bucket-967852078706	AWS_Bagchi	global
demoess	AWS_Bagchi	global
feb14test001	AWS_Bagchi	global
bucket05061234	AWS_Bagchi	global
bagchi-billing2	AWS_Bagchi	global

Pillar Name	Operational Excellence	**Status**	Violation
Question	09. How do you understand the health of your operations?	**No. of violated resources**	6
Best Practice	03. Collect and analyze operations metrics	**Resource Category**	Network
Severity	High	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	09. How do you understand the health of your operations?	**No. of violated resources**	28
Best Practice	03. Collect and analyze operations metrics	**Resource Category**	Storage
Severity	High	**Resource**	Buckets
Policy	AWS Audit S3 Bucket Logging Enabled CS Policy	**Resource Type**	S3
Policy Description	This policy audits whether Server access logging is enabled for Amazon S3 buckets. Server access logging provides detailed records for the requests that are made to an Amazon S3 bucket, which might be useful in security and access audits. This rule becomes non-compliant if Server access logging is not enabled for Amazon S3 buckets.		
Policy Recommendation	It is recommended to enable server access logging for S3 buckets. Server access logging provides detailed records for the requests that are made to a bucket. This helps in access tracking and trend analysis.		

Resource Id	Cloud Account	Region
abagchi-cloudtrail-bucket-us-east-1	AWS_Bagchi	global
bagchi-billing	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-1	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-2	AWS_Bagchi	global
bagchi-guardduty-ca-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-north-1	AWS_Bagchi	global
bagchi-guardduty-eu-west-2	AWS_Bagchi	global
bagchi-guardduty-eu-west-3	AWS_Bagchi	global
bagchi-guardduty-us-east-1	AWS_Bagchi	global
bagchi-guardduty-us-east-2	AWS_Bagchi	global
bagchi-guardduty-us-west-1	AWS_Bagchi	global
bagchi-guardduty-us-west-2	AWS_Bagchi	global
bagchi-guardduty2-ap-south-1	AWS_Bagchi	global
bagchi-guardduty2-eu-west-2	AWS_Bagchi	global
bagchi-guardduty2-us-east-1	AWS_Bagchi	global
bagchi-guardduty2-us-east-2	AWS_Bagchi	global

Resource Id	Cloud Account	Region
bagchi-guardduty2-us-west-1	AWS_Bagchi	global
bagchi-guardduty2-us-west-2	AWS_Bagchi	global
cf-templates-fg4y3e089bnm-us-east-1	AWS_Bagchi	global
cloud-trail-sales2-ap-south-1	AWS_Bagchi	global
cloud-trail-sales2-eu-west-2	AWS_Bagchi	global
cloud-trail-sales2-us-east-1	AWS_Bagchi	global
config-bucket-967852078706	AWS_Bagchi	global
demoess	AWS_Bagchi	global
feb14test001	AWS_Bagchi	global
bucket05061234	AWS_Bagchi	global
bagchi-billing2	AWS_Bagchi	global

Pillar Name	Operational Excellence	**Status**	Violation
Question	09. How do you understand the health of your operations?	**No. of violated resources**	6
Best Practice	06. Alert when operations outcomes are at risk	**Resource Category**	Network
Severity	Medium	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	09. How do you understand the health of your operations?	**No. of violated resources**	28
Best Practice	06. Alert when operations outcomes are at risk	**Resource Category**	Storage
Severity	Medium	**Resource**	Buckets
Policy	AWS Audit S3 Bucket Logging Enabled CS Policy	**Resource Type**	S3
Policy Description	This policy audits whether Server access logging is enabled for Amazon S3 buckets. Server access logging provides detailed records for the requests that are made to an Amazon S3 bucket, which might be useful in security and access audits. This rule becomes non-compliant if Server access logging is not enabled for Amazon S3 buckets.		
Policy Recommendation	It is recommended to enable server access logging for S3 buckets. Server access logging provides detailed records for the requests that are made to a bucket. This helps in access tracking and trend analysis.		

Resource Id	Cloud Account	Region
abagchi-cloudtrail-bucket-us-east-1	AWS_Bagchi	global
bagchi-billing	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-1	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-2	AWS_Bagchi	global
bagchi-guardduty-ca-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-north-1	AWS_Bagchi	global
bagchi-guardduty-eu-west-2	AWS_Bagchi	global
bagchi-guardduty-eu-west-3	AWS_Bagchi	global
bagchi-guardduty-us-east-1	AWS_Bagchi	global
bagchi-guardduty-us-east-2	AWS_Bagchi	global
bagchi-guardduty-us-west-1	AWS_Bagchi	global
bagchi-guardduty-us-west-2	AWS_Bagchi	global
bagchi-guardduty2-ap-south-1	AWS_Bagchi	global
bagchi-guardduty2-eu-west-2	AWS_Bagchi	global
bagchi-guardduty2-us-east-1	AWS_Bagchi	global
bagchi-guardduty2-us-east-2	AWS_Bagchi	global

Resource Id	Cloud Account	Region
bagchi-guardduty2-us-west-1	AWS_Bagchi	global
bagchi-guardduty2-us-west-2	AWS_Bagchi	global
cf-templates-fg4y3e089bnm-us-east-1	AWS_Bagchi	global
cloud-trail-sales2-ap-south-1	AWS_Bagchi	global
cloud-trail-sales2-eu-west-2	AWS_Bagchi	global
cloud-trail-sales2-us-east-1	AWS_Bagchi	global
config-bucket-967852078706	AWS_Bagchi	global
demoess	AWS_Bagchi	global
feb14test001	AWS_Bagchi	global
bucket05061234	AWS_Bagchi	global
bagchi-billing2	AWS_Bagchi	global

Pillar Name	Operational Excellence	**Status**	Violation
Question	09. How do you understand the health of your operations?	**No. of violated resources**	6
Best Practice	07. Alert when operations anomalies are detected	**Resource Category**	Network
Severity	Low	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	09. How do you understand the health of your operations?	**No. of violated resources**	28
Best Practice	07. Alert when operations anomalies are detected	**Resource Category**	Storage
Severity	Low	**Resource**	Buckets
Policy	AWS Audit S3 Bucket Logging Enabled CS Policy	**Resource Type**	S3
Policy Description	This policy audits whether Server access logging is enabled for Amazon S3 buckets. Server access logging provides detailed records for the requests that are made to an Amazon S3 bucket, which might be useful in security and access audits. This rule becomes non-compliant if Server access logging is not enabled for Amazon S3 buckets.		
Policy Recommendation	It is recommended to enable server access logging for S3 buckets. Server access logging provides detailed records for the requests that are made to a bucket. This helps in access tracking and trend analysis.		

Resource Id	Cloud Account	Region
abagchi-cloudtrail-bucket-us-east-1	AWS_Bagchi	global
bagchi-billing	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-1	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-2	AWS_Bagchi	global
bagchi-guardduty-ca-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-north-1	AWS_Bagchi	global
bagchi-guardduty-eu-west-2	AWS_Bagchi	global
bagchi-guardduty-eu-west-3	AWS_Bagchi	global
bagchi-guardduty-us-east-1	AWS_Bagchi	global
bagchi-guardduty-us-east-2	AWS_Bagchi	global
bagchi-guardduty-us-west-1	AWS_Bagchi	global
bagchi-guardduty-us-west-2	AWS_Bagchi	global
bagchi-guardduty2-ap-south-1	AWS_Bagchi	global
bagchi-guardduty2-eu-west-2	AWS_Bagchi	global
bagchi-guardduty2-us-east-1	AWS_Bagchi	global
bagchi-guardduty2-us-east-2	AWS_Bagchi	global

Resource Id	Cloud Account	Region
bagchi-guardduty2-us-west-1	AWS_Bagchi	global
bagchi-guardduty2-us-west-2	AWS_Bagchi	global
cf-templates-fg4y3e089bnm-us-east-1	AWS_Bagchi	global
cloud-trail-sales2-ap-south-1	AWS_Bagchi	global
cloud-trail-sales2-eu-west-2	AWS_Bagchi	global
cloud-trail-sales2-us-east-1	AWS_Bagchi	global
config-bucket-967852078706	AWS_Bagchi	global
demoess	AWS_Bagchi	global
feb14test001	AWS_Bagchi	global
bucket05061234	AWS_Bagchi	global
bagchi-billing2	AWS_Bagchi	global

Pillar Name	Operational Excellence	**Status**	Violation
Question	09. How do you understand the health of your operations?	**No. of violated resources**	6
Best Practice	08. KPIs and metrics	**Resource Category**	Network
Severity	Low	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Operational Excellence	**Status**	Violation
Question	09. How do you understand the health of your operations?	**No. of violated resources**	28
Best Practice	08. KPIs and metrics	**Resource Category**	Storage
Severity	Low	**Resource**	Buckets
Policy	AWS Audit S3 Bucket Logging Enabled CS Policy	**Resource Type**	S3
Policy Description	This policy audits whether Server access logging is enabled for Amazon S3 buckets. Server access logging provides detailed records for the requests that are made to an Amazon S3 bucket, which might be useful in security and access audits. This rule becomes non-compliant if Server access logging is not enabled for Amazon S3 buckets.		
Policy Recommendation	It is recommended to enable server access logging for S3 buckets. Server access logging provides detailed records for the requests that are made to a bucket. This helps in access tracking and trend analysis.		

Resource Id	Cloud Account	Region
abagchi-cloudtrail-bucket-us-east-1	AWS_Bagchi	global
bagchi-billing	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-1	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-2	AWS_Bagchi	global
bagchi-guardduty-ca-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-north-1	AWS_Bagchi	global
bagchi-guardduty-eu-west-2	AWS_Bagchi	global
bagchi-guardduty-eu-west-3	AWS_Bagchi	global
bagchi-guardduty-us-east-1	AWS_Bagchi	global
bagchi-guardduty-us-east-2	AWS_Bagchi	global
bagchi-guardduty-us-west-1	AWS_Bagchi	global
bagchi-guardduty-us-west-2	AWS_Bagchi	global
bagchi-guardduty2-ap-south-1	AWS_Bagchi	global
bagchi-guardduty2-eu-west-2	AWS_Bagchi	global
bagchi-guardduty2-us-east-1	AWS_Bagchi	global
bagchi-guardduty2-us-east-2	AWS_Bagchi	global

Resource Id	Cloud Account	Region
bagchi-guardduty2-us-west-1	AWS_Bagchi	global
bagchi-guardduty2-us-west-2	AWS_Bagchi	global
cf-templates-fg4y3e089bnm-us-east-1	AWS_Bagchi	global
cloud-trail-sales2-ap-south-1	AWS_Bagchi	global
cloud-trail-sales2-eu-west-2	AWS_Bagchi	global
cloud-trail-sales2-us-east-1	AWS_Bagchi	global
config-bucket-967852078706	AWS_Bagchi	global
demoess	AWS_Bagchi	global
feb14test001	AWS_Bagchi	global
bucket05061234	AWS_Bagchi	global
bagchi-billing2	AWS_Bagchi	global

Pillar Name	Performance Efficiency	**Status**	Violation
Question	02. How do you select your compute solution?	**No. of violated resources**	7
Best Practice	03. Collect compute-related metrics	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit EC2 Instance Detailed Monitoring Enabled CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether detailed monitoring is enabled for the EC2 instances. Enabling detailed monitoring allows the capture of metric data every minute, providing valuable insights for effective decision-making in architecting and managing compute resources within the account. This rule is Non-compliant when the EC2 instances are not configured with detailed monitoring.		
Policy Recommendation	It is recommended to enable detailed monitoring on the EC2 instances in order to have more granular control over the utilization and help track the performance of the instance. It also helps to get aggregated data across groups of similar instances.		

Resource Id	Cloud Account	Region
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Performance Efficiency	**Status**	Violation
Question	02. How do you select your compute solution?	**No. of violated resources**	11
Best Practice	05. Use the available elasticity of resources	**Resource Category**	Compute
Severity	Medium	**Resource**	Instances
Policy	AWS Audit Instance In Auto Scaling Group CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether the EC2 instances are in a AutoScaling Group. This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances. This rule will be non-compliant if the EC2 instances are not in a AutoScaling Group.		
Policy Recommendation	It is recommended that every EC2 instance is launched inside an Auto Scaling Group (ASG). This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances.		

Resource Id	Cloud Account	Region
i-0b3b9be418dc822d3	AWS_Bagchi	us-east-1
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0981155c9a8265683	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-048d69c6517f7a34a	AWS_Bagchi	us-east-1
i-0019b50a9ac031efe	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Performance Efficiency	**Status**	Violation
Question	05. How do you configure your networking solution?	**No. of violated resources**	1
Best Practice	07. Optimize network configuration based on metrics	**Resource Category**	Compute
Severity	Medium	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted NetBIOS Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow NetBIOS traffic (TCP port 139) from all addresses (0.0.0.0/0). This means the inbound access for NetBIOS is unrestrictedly allowed. Restricting access to only those IP addresses and NetBIOS port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if NetBIOS port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to limit access to TCP port 139 in EC2 Security groups only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). TCP Port 139 is used by NetBIOS. This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1

Pillar Name	Performance Efficiency	**Status**	Violation
Question	05. How do you configure your networking solution?	**No. of violated resources**	2
Best Practice	07. Optimize network configuration based on metrics	**Resource Category**	Compute
Severity	Medium	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted RPC Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow RPC access (TCP port 135) from all addresses (0.0.0.0/0). This means the inbound access for RPC is unrestrictedly allowed. Restricting access to only those IP addresses and RPC port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if RPC port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to limit access to TCP port 135 in EC2 Security groups only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). TCP Port 135 is used by RPC. This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1
sg-0a0fe7fa081d4dd46	AWS_Bagchi	us-east-1

Pillar Name	Performance Efficiency	**Status**	Violation
Question	05. How do you configure your networking solution?	**No. of violated resources**	6
Best Practice	07. Optimize network configuration based on metrics	**Resource Category**	Network
Severity	Medium	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Performance Efficiency	**Status**	Violation
Question	05. How do you configure your networking solution?	**No. of violated resources**	4
Best Practice	07. Optimize network configuration based on metrics	**Resource Category**	Network
Severity	Medium	**Resource**	Network_Acls
Policy	AWS Audit NetworkACL With Inbound Rule Allow All Traffic	**Resource Type**	VPC
Policy Description	This policy audits whether Network ACL inbound rules that allow all addresses (0.0.0.0/0). This means the inbound access is unrestrictedly allowed. Restricting access to only those IP addresses enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if the Network ACL inbound rules allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to configure Network ACL inbound rules to limit access only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
acl-225e855e	AWS_Bagchi	us-east-1
acl-025e1a3940e1d15af	AWS_Bagchi	us-east-1
acl-08ae090929b76abbf	AWS_Bagchi	us-east-1
acl-0681def15a2582b69	AWS_Bagchi	us-east-1

Pillar Name	Performance Efficiency	**Status**	Violation
Question	05. How do you configure your networking solution?	**No. of violated resources**	4
Best Practice	07. Optimize network configuration based on metrics	**Resource Category**	Network
Severity	Medium	**Resource**	Network_Acls
Policy	AWS Audit NetworkACL With Outbound Rule Allow All Traffic	**Resource Type**	VPC
Policy Description	This policy audits whether Network ACL outbound rules that allow all addresses (0.0.0.0/0). This means the outbound access is unrestrictedly allowed. Restricting access to only those IP addresses enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if the Network ACL outbound rules allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to configure Network ACL outbound rules to limit access only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
acl-225e855e	AWS_Bagchi	us-east-1
acl-025e1a3940e1d15af	AWS_Bagchi	us-east-1
acl-08ae090929b76abbf	AWS_Bagchi	us-east-1
acl-0681def15a2582b69	AWS_Bagchi	us-east-1

Pillar Name	Performance Efficiency	**Status**	Violation
Question	07. How do you monitor your resources to ensure they are performing?	**No. of violated resources**	7
Best Practice	01. Record performance-related metrics	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit EC2 Instance Detailed Monitoring Enabled CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether detailed monitoring is enabled for the EC2 instances. Enabling detailed monitoring allows the capture of metric data every minute, providing valuable insights for effective decision-making in architecting and managing compute resources within the account. This rule is Non-compliant when the EC2 instances are not configured with detailed monitoring.		
Policy Recommendation	It is recommended to enable detailed monitoring on the EC2 instances in order to have more granular control over the utilization and help track the performance of the instance. It also helps to get aggregated data across groups of similar instances.		

Resource Id	Cloud Account	Region
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Performance Efficiency	**Status**	Violation
Question	07. How do you monitor your resources to ensure they are performing?	**No. of violated resources**	7
Best Practice	02. Analyze metrics when events or incidents occur	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit EC2 Instance Detailed Monitoring Enabled CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether detailed monitoring is enabled for the EC2 instances. Enabling detailed monitoring allows the capture of metric data every minute, providing valuable insights for effective decision-making in architecting and managing compute resources within the account. This rule is Non-compliant when the EC2 instances are not configured with detailed monitoring.		
Policy Recommendation	It is recommended to enable detailed monitoring on the EC2 instances in order to have more granular control over the utilization and help track the performance of the instance. It also helps to get aggregated data across groups of similar instances.		

Resource Id	Cloud Account	Region
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Performance Efficiency	**Status**	Violation
Question	07. How do you monitor your resources to ensure they are performing?	**No. of violated resources**	7
Best Practice	04. Use monitoring to generate alarm-based notifications	**Resource Category**	Compute
Severity	Medium	**Resource**	Instances
Policy	AWS Audit EC2 Instance Detailed Monitoring Enabled CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether detailed monitoring is enabled for the EC2 instances. Enabling detailed monitoring allows the capture of metric data every minute, providing valuable insights for effective decision-making in architecting and managing compute resources within the account. This rule is Non-compliant when the EC2 instances are not configured with detailed monitoring.		
Policy Recommendation	It is recommended to enable detailed monitoring on the EC2 instances in order to have more granular control over the utilization and help track the performance of the instance. It also helps to get aggregated data across groups of similar instances.		

Resource Id	Cloud Account	Region
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Performance Efficiency	**Status**	Violation
Question	07. How do you monitor your resources to ensure they are performing?	**No. of violated resources**	7
Best Practice	05. Review metrics at regular intervals	**Resource Category**	Compute
Severity	Medium	**Resource**	Instances
Policy	AWS Audit EC2 Instance Detailed Monitoring Enabled CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether detailed monitoring is enabled for the EC2 instances. Enabling detailed monitoring allows the capture of metric data every minute, providing valuable insights for effective decision-making in architecting and managing compute resources within the account. This rule is Non-compliant when the EC2 instances are not configured with detailed monitoring.		
Policy Recommendation	It is recommended to enable detailed monitoring on the EC2 instances in order to have more granular control over the utilization and help track the performance of the instance. It also helps to get aggregated data across groups of similar instances.		

Resource Id	Cloud Account	Region
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Performance Efficiency	**Status**	Violation
Question	07. How do you monitor your resources to ensure they are performing?	**No. of violated resources**	7
Best Practice	06. Monitor and alarm proactively	**Resource Category**	Compute
Severity	Low	**Resource**	Instances
Policy	AWS Audit EC2 Instance Detailed Monitoring Enabled CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether detailed monitoring is enabled for the EC2 instances. Enabling detailed monitoring allows the capture of metric data every minute, providing valuable insights for effective decision-making in architecting and managing compute resources within the account. This rule is Non-compliant when the EC2 instances are not configured with detailed monitoring.		
Policy Recommendation	It is recommended to enable detailed monitoring on the EC2 instances in order to have more granular control over the utilization and help track the performance of the instance. It also helps to get aggregated data across groups of similar instances.		

Resource Id	Cloud Account	Region
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Reliability	**Status**	Violation
Question	06. How do you monitor workload resources?	**No. of violated resources**	6
Best Practice	02. Define and calculate metrics (Aggregation)	**Resource Category**	Network
Severity	High	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Reliability	**Status**	Violation
Question	06. How do you monitor workload resources?	**No. of violated resources**	28
Best Practice	02. Define and calculate metrics (Aggregation)	**Resource Category**	Storage
Severity	High	**Resource**	Buckets
Policy	AWS Audit S3 Bucket Logging Enabled CS Policy	**Resource Type**	S3
Policy Description	This policy audits whether Server access logging is enabled for Amazon S3 buckets. Server access logging provides detailed records for the requests that are made to an Amazon S3 bucket, which might be useful in security and access audits. This rule becomes non-compliant if Server access logging is not enabled for Amazon S3 buckets.		
Policy Recommendation	It is recommended to enable server access logging for S3 buckets. Server access logging provides detailed records for the requests that are made to a bucket. This helps in access tracking and trend analysis.		

Resource Id	Cloud Account	Region
abagchi-cloudtrail-bucket-us-east-1	AWS_Bagchi	global
bagchi-billing	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-1	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-2	AWS_Bagchi	global
bagchi-guardduty-ca-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-north-1	AWS_Bagchi	global
bagchi-guardduty-eu-west-2	AWS_Bagchi	global
bagchi-guardduty-eu-west-3	AWS_Bagchi	global
bagchi-guardduty-us-east-1	AWS_Bagchi	global
bagchi-guardduty-us-east-2	AWS_Bagchi	global
bagchi-guardduty-us-west-1	AWS_Bagchi	global
bagchi-guardduty-us-west-2	AWS_Bagchi	global
bagchi-guardduty2-ap-south-1	AWS_Bagchi	global
bagchi-guardduty2-eu-west-2	AWS_Bagchi	global
bagchi-guardduty2-us-east-1	AWS_Bagchi	global
bagchi-guardduty2-us-east-2	AWS_Bagchi	global

Resource Id	Cloud Account	Region
bagchi-guardduty2-us-west-1	AWS_Bagchi	global
bagchi-guardduty2-us-west-2	AWS_Bagchi	global
cf-templates-fg4y3e089bnm-us-east-1	AWS_Bagchi	global
cloud-trail-sales2-ap-south-1	AWS_Bagchi	global
cloud-trail-sales2-eu-west-2	AWS_Bagchi	global
cloud-trail-sales2-us-east-1	AWS_Bagchi	global
config-bucket-967852078706	AWS_Bagchi	global
demoess	AWS_Bagchi	global
feb14test001	AWS_Bagchi	global
bucket05061234	AWS_Bagchi	global
bagchi-billing2	AWS_Bagchi	global

Pillar Name	Reliability	**Status**	Violation
Question	06. How do you monitor workload resources?	**No. of violated resources**	6
Best Practice	05. Analytics	**Resource Category**	Network
Severity	Medium	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Reliability	**Status**	Violation
Question	06. How do you monitor workload resources?	**No. of violated resources**	6
Best Practice	06. Conduct reviews regularly	**Resource Category**	Network
Severity	Medium	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Reliability	**Status**	Violation
Question	06. How do you monitor workload resources?	**No. of violated resources**	28
Best Practice	06. Conduct reviews regularly	**Resource Category**	Storage
Severity	Medium	**Resource**	Buckets
Policy	AWS Audit S3 Bucket Logging Enabled CS Policy	**Resource Type**	S3
Policy Description	This policy audits whether Server access logging is enabled for Amazon S3 buckets. Server access logging provides detailed records for the requests that are made to an Amazon S3 bucket, which might be useful in security and access audits. This rule becomes non-compliant if Server access logging is not enabled for Amazon S3 buckets.		
Policy Recommendation	It is recommended to enable server access logging for S3 buckets. Server access logging provides detailed records for the requests that are made to a bucket. This helps in access tracking and trend analysis.		

Resource Id	Cloud Account	Region
abagchi-cloudtrail-bucket-us-east-1	AWS_Bagchi	global
bagchi-billing	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-1	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-2	AWS_Bagchi	global
bagchi-guardduty-ca-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-north-1	AWS_Bagchi	global
bagchi-guardduty-eu-west-2	AWS_Bagchi	global
bagchi-guardduty-eu-west-3	AWS_Bagchi	global
bagchi-guardduty-us-east-1	AWS_Bagchi	global
bagchi-guardduty-us-east-2	AWS_Bagchi	global
bagchi-guardduty-us-west-1	AWS_Bagchi	global
bagchi-guardduty-us-west-2	AWS_Bagchi	global
bagchi-guardduty2-ap-south-1	AWS_Bagchi	global
bagchi-guardduty2-eu-west-2	AWS_Bagchi	global
bagchi-guardduty2-us-east-1	AWS_Bagchi	global
bagchi-guardduty2-us-east-2	AWS_Bagchi	global

Resource Id	Cloud Account	Region
bagchi-guardduty2-us-west-1	AWS_Bagchi	global
bagchi-guardduty2-us-west-2	AWS_Bagchi	global
cf-templates-fg4y3e089bnm-us-east-1	AWS_Bagchi	global
cloud-trail-sales2-ap-south-1	AWS_Bagchi	global
cloud-trail-sales2-eu-west-2	AWS_Bagchi	global
cloud-trail-sales2-us-east-1	AWS_Bagchi	global
config-bucket-967852078706	AWS_Bagchi	global
demoess	AWS_Bagchi	global
feb14test001	AWS_Bagchi	global
bucket05061234	AWS_Bagchi	global
bagchi-billing2	AWS_Bagchi	global

Pillar Name	Reliability	**Status**	Violation
Question	07. How do you design your workload to adapt to changes in demand?	**No. of violated resources**	11
Best Practice	01. Use automation when obtaining or scaling resources	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit Instance In Auto Scaling Group CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether the EC2 instances are in a AutoScaling Group. This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances. This rule will be non-compliant if the EC2 instances are not in a AutoScaling Group.		
Policy Recommendation	It is recommended that every EC2 instance is launched inside an Auto Scaling Group (ASG). This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances.		

Resource Id	Cloud Account	Region
i-0b3b9be418dc822d3	AWS_Bagchi	us-east-1
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0981155c9a8265683	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-048d69c6517f7a34a	AWS_Bagchi	us-east-1
i-0019b50a9ac031efe	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Reliability	**Status**	Violation
Question	07. How do you design your workload to adapt to changes in demand?	**No. of violated resources**	11
Best Practice	02. Obtain resources upon detection of impairment to a workload	**Resource Category**	Compute
Severity	Medium	**Resource**	Instances
Policy	AWS Audit Instance In Auto Scaling Group CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether the EC2 instances are in a AutoScaling Group. This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances. This rule will be non-compliant if the EC2 instances are not in a AutoScaling Group.		
Policy Recommendation	It is recommended that every EC2 instance is launched inside an Auto Scaling Group (ASG). This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances.		

Resource Id	Cloud Account	Region
i-0b3b9be418dc822d3	AWS_Bagchi	us-east-1
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0981155c9a8265683	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-048d69c6517f7a34a	AWS_Bagchi	us-east-1
i-0019b50a9ac031efe	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Reliability	**Status**	Violation
Question	07. How do you design your workload to adapt to changes in demand?	**No. of violated resources**	11
Best Practice	03. Obtain resources upon detection that more resources are needed for a workload	**Resource Category**	Compute
Severity	Medium	**Resource**	Instances
Policy	AWS Audit Instance In Auto Scaling Group CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether the EC2 instances are in a AutoScaling Group. This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances. This rule will be non-compliant if the EC2 instances are not in a AutoScaling Group.		
Policy Recommendation	It is recommended that every EC2 instance is launched inside an Auto Scaling Group (ASG). This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances.		

Resource Id	Cloud Account	Region
i-0b3b9be418dc822d3	AWS_Bagchi	us-east-1
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0981155c9a8265683	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-048d69c6517f7a34a	AWS_Bagchi	us-east-1
i-0019b50a9ac031efe	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Reliability	**Status**	Violation
Question	09. How do you back up data?	**No. of violated resources**	9
Best Practice	01. Identify and back up all data that needs to be backed up, or reproduce the data from sources	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit Ec2 Instances Without Backup CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether the backups are available for the EC2 instances. This can be achieved by creating a full backup of an EC2 instance by taking a snapshot of an individual volumes. This rule becomes non-compliant if the EC2 instances does not have any snapshots created for the volumes used by the EC2 instances.		
Policy Recommendation	IIt is recommended to take a snapshot of all the required volumes which are in-use by the AWS EC2 instances. Amazon EBS snapshots enable point-in-time backups of the data, which are durable and highly available and can be used for disaster recovery (DR) when information is lost due to human error.		

Resource Id	Cloud Account	Region
i-0b3b9be418dc822d3	AWS_Bagchi	us-east-1
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0981155c9a8265683	AWS_Bagchi	us-east-1
i-048d69c6517f7a34a	AWS_Bagchi	us-east-1
i-0019b50a9ac031efe	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Reliability	**Status**	Violation
Question	09. How do you back up data?	**No. of violated resources**	9
Best Practice	03. Perform data backup automatically	**Resource Category**	Compute
Severity	Medium	**Resource**	Instances
Policy	AWS Audit Ec2 Instances Without Backup CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether the backups are available for the EC2 instances. This can be achieved by creating a full backup of an EC2 instance by taking a snapshot of an individual volumes. This rule becomes non-compliant if the EC2 instances does not have any snapshots created for the volumes used by the EC2 instances.		
Policy Recommendation	IIt is recommended to take a snapshot of all the required volumes which are in-use by the AWS EC2 instances. Amazon EBS snapshots enable point-in-time backups of the data, which are durable and highly available and can be used for disaster recovery (DR) when information is lost due to human error.		

Resource Id	Cloud Account	Region
i-0b3b9be418dc822d3	AWS_Bagchi	us-east-1
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0981155c9a8265683	AWS_Bagchi	us-east-1
i-048d69c6517f7a34a	AWS_Bagchi	us-east-1
i-0019b50a9ac031efe	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Reliability	**Status**	Violation
Question	11. How do you design your workload to withstand component failures?	**No. of violated resources**	11
Best Practice	02. Fail over to healthy resources	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit Instance In Auto Scaling Group CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether the EC2 instances are in a AutoScaling Group. This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances. This rule will be non-compliant if the EC2 instances are not in a AutoScaling Group.		
Policy Recommendation	It is recommended that every EC2 instance is launched inside an Auto Scaling Group (ASG). This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances.		

Resource Id	Cloud Account	Region
i-0b3b9be418dc822d3	AWS_Bagchi	us-east-1
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0981155c9a8265683	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-048d69c6517f7a34a	AWS_Bagchi	us-east-1
i-0019b50a9ac031efe	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Reliability	**Status**	Violation
Question	13. How do you plan for disaster recovery (DR)?	**No. of violated resources**	11
Best Practice	02. Use defined recovery strategies to meet the recovery objectives	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit Instance In Auto Scaling Group CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether the EC2 instances are in a AutoScaling Group. This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances. This rule will be non-compliant if the EC2 instances are not in a AutoScaling Group.		
Policy Recommendation	It is recommended that every EC2 instance is launched inside an Auto Scaling Group (ASG). This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances.		

Resource Id	Cloud Account	Region
i-0b3b9be418dc822d3	AWS_Bagchi	us-east-1
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0981155c9a8265683	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-048d69c6517f7a34a	AWS_Bagchi	us-east-1
i-0019b50a9ac031efe	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Reliability	**Status**	Violation
Question	13. How do you plan for disaster recovery (DR)?	**No. of violated resources**	11
Best Practice	05. Automate recovery	**Resource Category**	Compute
Severity	Medium	**Resource**	Instances
Policy	AWS Audit Instance In Auto Scaling Group CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether the EC2 instances are in a AutoScaling Group. This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances. This rule will be non-compliant if the EC2 instances are not in a AutoScaling Group.		
Policy Recommendation	It is recommended that every EC2 instance is launched inside an Auto Scaling Group (ASG). This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances.		

Resource Id	Cloud Account	Region
i-0b3b9be418dc822d3	AWS_Bagchi	us-east-1
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0981155c9a8265683	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-048d69c6517f7a34a	AWS_Bagchi	us-east-1
i-0019b50a9ac031efe	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	02. How do you manage identities for people and machines?	**No. of violated resources**	1
Best Practice	02. Use temporary credentials	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit EC2 Instance Using IAM Roles CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether IAM roles are used to grant permission to access AWS resource from EC2 instances. AWS IAM roles reduce the risks associated with sharing and rotating credentials that can be used outside of AWS itself. If credentials are compromised, they can be used from outside of the AWS account. This rule is non-compliant if IAM roles are not used to access AWS resources by the applications hosted in EC2 instances.		
Policy Recommendation	It is recommended to use Instance Profiles/IAM Roles to appropriately grant permissions to applications running on amazon EC2 instances. AWS IAM roles reduce the risks associated with sharing and rotating credentials that can be used outside of AWS itself. IAM roles can only be associated at the launch of an instance.		

Resource Id	Cloud Account	Region
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	03. How do you manage permissions for people and machines?	**No. of violated resources**	1
Best Practice	01. Define access requirements	**Resource Category**	Compute
Severity	High	**Resource**	Instances
Policy	AWS Audit EC2 Instance Using IAM Roles CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether IAM roles are used to grant permission to access AWS resource from EC2 instances. AWS IAM roles reduce the risks associated with sharing and rotating credentials that can be used outside of AWS itself. If credentials are compromised, they can be used from outside of the AWS account. This rule is non-compliant if IAM roles are not used to access AWS resources by the applications hosted in EC2 instances.		
Policy Recommendation	It is recommended to use Instance Profiles/IAM Roles to appropriately grant permissions to applications running on amazon EC2 instances. AWS IAM roles reduce the risks associated with sharing and rotating credentials that can be used outside of AWS itself. IAM roles can only be associated at the launch of an instance.		

Resource Id	Cloud Account	Region
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	03. How do you manage permissions for people and machines?	**No. of violated resources**	1
Best Practice	07. Analyze public and cross-account access	**Resource Category**	Compute
Severity	Low	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted RDP Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow RDP access (TCP port 3389) from all addresses (0.0.0.0/0). This means the inbound access for RDP is unrestrictedly allowed. Restricting access to only those IP addresses and RDP port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if RDP port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to limit access to TCP port 3389 in EC2 Security groups only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). TCP Port 3389 is used by RDP. This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	04. How do you detect and investigate security events?	**No. of violated resources**	6
Best Practice	01. Configure service and application logging	**Resource Category**	Network
Severity	High	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Security	**Status**	Violation
Question	04. How do you detect and investigate security events?	**No. of violated resources**	28
Best Practice	01. Configure service and application logging	**Resource Category**	Storage
Severity	High	**Resource**	Buckets
Policy	AWS Audit S3 Bucket Logging Enabled CS Policy	**Resource Type**	S3
Policy Description	This policy audits whether Server access logging is enabled for Amazon S3 buckets. Server access logging provides detailed records for the requests that are made to an Amazon S3 bucket, which might be useful in security and access audits. This rule becomes non-compliant if Server access logging is not enabled for Amazon S3 buckets.		
Policy Recommendation	It is recommended to enable server access logging for S3 buckets. Server access logging provides detailed records for the requests that are made to a bucket. This helps in access tracking and trend analysis.		

Resource Id	Cloud Account	Region
abagchi-cloudtrail-bucket-us-east-1	AWS_Bagchi	global
bagchi-billing	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-1	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-2	AWS_Bagchi	global
bagchi-guardduty-ca-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-north-1	AWS_Bagchi	global
bagchi-guardduty-eu-west-2	AWS_Bagchi	global
bagchi-guardduty-eu-west-3	AWS_Bagchi	global
bagchi-guardduty-us-east-1	AWS_Bagchi	global
bagchi-guardduty-us-east-2	AWS_Bagchi	global
bagchi-guardduty-us-west-1	AWS_Bagchi	global
bagchi-guardduty-us-west-2	AWS_Bagchi	global
bagchi-guardduty2-ap-south-1	AWS_Bagchi	global
bagchi-guardduty2-eu-west-2	AWS_Bagchi	global
bagchi-guardduty2-us-east-1	AWS_Bagchi	global
bagchi-guardduty2-us-east-2	AWS_Bagchi	global

Resource Id	Cloud Account	Region
bagchi-guardduty2-us-west-1	AWS_Bagchi	global
bagchi-guardduty2-us-west-2	AWS_Bagchi	global
cf-templates-fg4y3e089bnm-us-east-1	AWS_Bagchi	global
cloud-trail-sales2-ap-south-1	AWS_Bagchi	global
cloud-trail-sales2-eu-west-2	AWS_Bagchi	global
cloud-trail-sales2-us-east-1	AWS_Bagchi	global
config-bucket-967852078706	AWS_Bagchi	global
demoess	AWS_Bagchi	global
feb14test001	AWS_Bagchi	global
bucket05061234	AWS_Bagchi	global
bagchi-billing2	AWS_Bagchi	global

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	1
Best Practice	02. Control traffic at all layers	**Resource Category**	Compute
Severity	High	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted NetBIOS Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow NetBIOS traffic (TCP port 139) from all addresses (0.0.0.0/0). This means the inbound access for NetBIOS is unrestrictedly allowed. Restricting access to only those IP addresses and NetBIOS port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if NetBIOS port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to limit access to TCP port 139 in EC2 Security groups only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). TCP Port 139 is used by NetBIOS. This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	1
Best Practice	02. Control traffic at all layers	**Resource Category**	Compute
Severity	High	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted RDP Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow RDP access (TCP port 3389) from all addresses (0.0.0.0/0). This means the inbound access for RDP is unrestrictedly allowed. Restricting access to only those IP addresses and RDP port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if RDP port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to limit access to TCP port 3389 in EC2 Security groups only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). TCP Port 3389 is used by RDP. This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	2
Best Practice	02. Control traffic at all layers	**Resource Category**	Compute
Severity	High	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted RPC Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow RPC access (TCP port 135) from all addresses (0.0.0.0/0). This means the inbound access for RPC is unrestrictedly allowed. Restricting access to only those IP addresses and RPC port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if RPC port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to limit access to TCP port 135 in EC2 Security groups only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). TCP Port 135 is used by RPC. This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1
sg-0a0fe7fa081d4dd46	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	1
Best Practice	02. Control traffic at all layers	**Resource Category**	Compute
Severity	High	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted Oracle Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow Oracle access (TCP port 1521) from all addresses (0.0.0.0/0). This means the inbound access for Oracle is unrestrictedly allowed. Restricting access to only those IP addresses and Oracle port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if Oracle port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to verify if EC2 security groups for inbound rules that allow unrestricted access to TCP port 1521 and restrict access to only those IP addresses that require it in order to implement the principle of least privilege and reduce the possibility of a breach. TCP port 1521 is used by the Oracle Database Server which is an object-relational database management system (RDBMS) server developed by Oracle Corporation:ÃŠ https://en.wikipedia.org/wiki/Oracle_Database		

Resource Id	Cloud Account	Region
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	7
Best Practice	02. Control traffic at all layers	**Resource Category**	Compute
Severity	High	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted HTTP Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow HTTP traffic (TCP port 80) from all addresses (0.0.0.0/0). This means the inbound access for HTTP is unrestrictedly allowed. Restricting access to only those IP addresses and HTTP port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if HTTP port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to limit access to TCP port 80 in EC2 Security groups only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). TCP Port 80 is used by HTTP. This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
sg-06cf8f6c379f72972	AWS_Bagchi	us-east-1
sg-0a92d97292ef1e035	AWS_Bagchi	us-east-1
sg-044cebb2f306fa722	AWS_Bagchi	us-east-1
sg-07dc3a3d238a30dcd	AWS_Bagchi	us-east-1
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1
sg-0009bad022b207bb4	AWS_Bagchi	us-east-1
sg-08a870c6344de8357	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	1
Best Practice	02. Control traffic at all layers	**Resource Category**	Compute
Severity	High	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted MongoDB Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow MongoDB access (TCP port 27017) from all addresses (0.0.0.0/0). This means the inbound access for MongoDB is unrestrictedly allowed. Restricting access to only those IP addresses and MongoDB port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if MongoDB port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to limit access to TCP port 27017 in EC2 Security groups only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). TCP Port 27017 is used by MongoDB. This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	1
Best Practice	02. Control traffic at all layers	**Resource Category**	Compute
Severity	High	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted MySQL Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow MySQL access (TCP port 3306) from all addresses (0.0.0.0/0). This means the inbound access for MySQL is unrestrictedly allowed. Restricting access to only those IP addresses and MySQL port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if MySQL port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to limit access to TCP port 3306 in EC2 Security groups only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). TCP Port 3306 is used by MySQL. This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	1
Best Practice	02. Control traffic at all layers	**Resource Category**	Compute
Severity	High	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted HTTPS Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow HTTPS traffic (TCP port 443) from all addresses (0.0.0.0/0). This means the inbound access for HTTPS is unrestrictedly allowed. Restricting access to only those IP addresses and HTTPS port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if HTTPS port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to limit access to TCP port 443 in EC2 Security groups only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). TCP Port 443 is used by HTTPS. This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	1
Best Practice	02. Control traffic at all layers	**Resource Category**	Compute
Severity	High	**Resource**	Security_Groups
Policy	AWS Audit Unrestricted ElasticSearch Access CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether EC2 security groups for inbound rules that allow Elasticsearch access (TCP port 9200) from all addresses (0.0.0.0/0). This means the inbound access for Elasticsearch is unrestrictedly allowed. Restricting access to only those IP addresses and Elasticsearch port enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if Elasticsearch port allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to limit access to TCP port 9200 in EC2 Security groups only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). TCP Port 9200 is used by Elasticsearch. This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
sg-0b9f8be0fe1bfa144	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	4
Best Practice	02. Control traffic at all layers	**Resource Category**	Network
Severity	High	**Resource**	Network_Acls
Policy	AWS Audit NetworkACL With Inbound Rule Allow All Traffic	**Resource Type**	VPC
Policy Description	This policy audits whether Network ACL inbound rules that allow all addresses (0.0.0.0/0). This means the inbound access is unrestrictedly allowed. Restricting access to only those IP addresses enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if the Network ACL inbound rules allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to configure Network ACL inbound rules to limit access only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
acl-225e855e	AWS_Bagchi	us-east-1
acl-025e1a3940e1d15af	AWS_Bagchi	us-east-1
acl-08ae090929b76abbf	AWS_Bagchi	us-east-1
acl-0681def15a2582b69	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	4
Best Practice	02. Control traffic at all layers	**Resource Category**	Network
Severity	High	**Resource**	Network_Acls
Policy	AWS Audit NetworkACL With Outbound Rule Allow All Traffic	**Resource Type**	VPC
Policy Description	This policy audits whether Network ACL outbound rules that allow all addresses (0.0.0.0/0). This means the outbound access is unrestrictedly allowed. Restricting access to only those IP addresses enables the principle of least privilege and reduce the possibility of a breach. This rule becomes non-compliant if the Network ACL outbound rules allows all addresses (i.e. 0.0.0.0/0 or ::/0).		
Policy Recommendation	It is recommended to configure Network ACL outbound rules to limit access only to the required IP addresses instead of allowing unrestricted access (i.e., 0.0.0.0/0). This enables the principle of least privilege and reduce the possibility of a breach.		

Resource Id	Cloud Account	Region
acl-225e855e	AWS_Bagchi	us-east-1
acl-025e1a3940e1d15af	AWS_Bagchi	us-east-1
acl-08ae090929b76abbf	AWS_Bagchi	us-east-1
acl-0681def15a2582b69	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	05. How do you protect your network resources?	**No. of violated resources**	6
Best Practice	04. Implement inspection and protection	**Resource Category**	Network
Severity	Low	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Security	**Status**	Violation
Question	08. How do you protect your data at rest?	**No. of violated resources**	1
Best Practice	04. Enforce access control	**Resource Category**	Compute
Severity	Low	**Resource**	Instances
Policy	AWS Audit EC2 Instance Using IAM Roles CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether IAM roles are used to grant permission to access AWS resource from EC2 instances. AWS IAM roles reduce the risks associated with sharing and rotating credentials that can be used outside of AWS itself. If credentials are compromised, they can be used from outside of the AWS account. This rule is non-compliant if IAM roles are not used to access AWS resources by the applications hosted in EC2 instances.		
Policy Recommendation	It is recommended to use Instance Profiles/IAM Roles to appropriately grant permissions to applications running on amazon EC2 instances. AWS IAM roles reduce the risks associated with sharing and rotating credentials that can be used outside of AWS itself. IAM roles can only be associated at the launch of an instance.		

Resource Id	Cloud Account	Region
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1

Pillar Name	Security	**Status**	Violation
Question	08. How do you protect your data at rest?	**No. of violated resources**	6
Best Practice	04. Enforce access control	**Resource Category**	Compute
Severity	Low	**Resource**	Security_Groups
Policy	AWS Audit Default SecurityGroup Closed CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether the default security group of any Amazon Virtual Private Cloud (Amazon VPC) does not allow inbound or outbound traffic. The rule becomes non-compliant if the default security group has one or more inbound or outbound traffic rules.		
Policy Recommendation	It is recommended to restrict Amazon EC2 default security groups all inbound public traffic in order to enforce AWS users (administrators, resource managers, etc.) to create custom security groups that exercise the Principle of Least Privilege (POLP) instead of using the default security groups.		

Resource Id	Cloud Account	Region
sg-0ae0afc704d6f4a42	AWS_Bagchi	us-east-1
sg-0da9b07b3d85448fb	AWS_Bagchi	us-east-1
sg-0932dceb37352180b	AWS_Bagchi	us-east-1
sg-634fdb7d	AWS_Bagchi	us-east-1
sg-0f2117e9cc6bc46cc	AWS_Bagchi	us-east-1
sg-c3034cbc	AWS_Bagchi	ap-south-1

Pillar Name	Security	**Status**	Violation
Question	10. How do you anticipate, respond to, and recover from incidents?	**No. of violated resources**	6
Best Practice	03. Prepare forensic capabilities	**Resource Category**	Network
Severity	Medium	**Resource**	VPC
Policy	AWS Audit VPC Flow Logs Enabled CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether VPC flow logs is enabled. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues. This rule becomes non-compliant if VPC flow logs is not enabled.		
Policy Recommendation	It is recommended to enable VPC flow logs. Enabling VPC flow logs will start collecting network traffic data to and from your VPC network, data that can be useful to detect and troubleshoot security issues.		

Resource Id	Cloud Account	Region
vpc-0a121c8358ac1b75f	AWS_Bagchi	us-east-1
vpc-04b17dc3f05481449	AWS_Bagchi	us-east-1
vpc-f6ed9d8b	AWS_Bagchi	us-east-1
vpc-04e3274156b6a9c39	AWS_Bagchi	us-east-1
vpc-0ff157a4c653078a4	AWS_Bagchi	us-east-1
vpc-ee904785	AWS_Bagchi	ap-south-1

Pillar Name	Security	**Status**	Violation
Question	10. How do you anticipate, respond to, and recover from incidents?	**No. of violated resources**	28
Best Practice	03. Prepare forensic capabilities	**Resource Category**	Storage
Severity	Medium	**Resource**	Buckets
Policy	AWS Audit S3 Bucket Logging Enabled CS Policy	**Resource Type**	S3
Policy Description	This policy audits whether Server access logging is enabled for Amazon S3 buckets. Server access logging provides detailed records for the requests that are made to an Amazon S3 bucket, which might be useful in security and access audits. This rule becomes non-compliant if Server access logging is not enabled for Amazon S3 buckets.		
Policy Recommendation	It is recommended to enable server access logging for S3 buckets. Server access logging provides detailed records for the requests that are made to a bucket. This helps in access tracking and trend analysis.		

Resource Id	Cloud Account	Region
abagchi-cloudtrail-bucket-us-east-1	AWS_Bagchi	global
bagchi-billing	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-1	AWS_Bagchi	global
bagchi-guardduty-ap-southeast-2	AWS_Bagchi	global
bagchi-guardduty-ca-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-central-1	AWS_Bagchi	global
bagchi-guardduty-eu-north-1	AWS_Bagchi	global
bagchi-guardduty-eu-west-2	AWS_Bagchi	global
bagchi-guardduty-eu-west-3	AWS_Bagchi	global
bagchi-guardduty-us-east-1	AWS_Bagchi	global
bagchi-guardduty-us-east-2	AWS_Bagchi	global
bagchi-guardduty-us-west-1	AWS_Bagchi	global
bagchi-guardduty-us-west-2	AWS_Bagchi	global
bagchi-guardduty2-ap-south-1	AWS_Bagchi	global
bagchi-guardduty2-eu-west-2	AWS_Bagchi	global
bagchi-guardduty2-us-east-1	AWS_Bagchi	global
bagchi-guardduty2-us-east-2	AWS_Bagchi	global

Resource Id	Cloud Account	Region
bagchi-guardduty2-us-west-1	AWS_Bagchi	global
bagchi-guardduty2-us-west-2	AWS_Bagchi	global
cf-templates-fg4y3e089bnm-us-east-1	AWS_Bagchi	global
cloud-trail-sales2-ap-south-1	AWS_Bagchi	global
cloud-trail-sales2-eu-west-2	AWS_Bagchi	global
cloud-trail-sales2-us-east-1	AWS_Bagchi	global
config-bucket-967852078706	AWS_Bagchi	global
demoess	AWS_Bagchi	global
feb14test001	AWS_Bagchi	global
bucket05061234	AWS_Bagchi	global
bagchi-billing2	AWS_Bagchi	global

Pillar Name	Sustainability	**Status**	Violation
Question	02. How do you align cloud resources to your demand?	**No. of violated resources**	11
Best Practice	01. Scale workload infrastructure dynamically	**Resource Category**	Compute
Severity	Medium	**Resource**	Instances
Policy	AWS Audit Instance In Auto Scaling Group CS Policy	**Resource Type**	EC2
Policy Description	This policy audits whether the EC2 instances are in a AutoScaling Group. This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances. This rule will be non-compliant if the EC2 instances are not in a AutoScaling Group.		
Policy Recommendation	It is recommended that every EC2 instance is launched inside an Auto Scaling Group (ASG). This maintains the availability of the EC2 resources in the event of a failure and improve scaling for the instances.		

Resource Id	Cloud Account	Region
i-0b3b9be418dc822d3	AWS_Bagchi	us-east-1
i-09ad198bfc7874934	AWS_Bagchi	us-east-1
i-061557dee14ac995a	AWS_Bagchi	us-east-1
i-0981155c9a8265683	AWS_Bagchi	us-east-1
i-0271fc0d0a890a9a3	AWS_Bagchi	us-east-1
i-0b36e1b8517d3a7ae	AWS_Bagchi	us-east-1
i-048d69c6517f7a34a	AWS_Bagchi	us-east-1
i-0019b50a9ac031efe	AWS_Bagchi	us-east-1
i-0bcecc40e1b3fa897	AWS_Bagchi	us-east-1
i-09bd67df221c21764	AWS_Bagchi	us-east-1
i-07dd9dac9de6c7730	AWS_Bagchi	us-east-1

Pillar Name	Sustainability	**Status**	Violation
Question	02. How do you align cloud resources to your demand?	**No. of violated resources**	11
Best Practice	03. Stop the creation and maintenance of unused assets	**Resource Category**	Storage
Severity	Low	**Resource**	Volumes
Policy	AWS EBS Volume Orphaned	**Resource Type**	EBS
Policy Description	This policy identifies the EBS Volumes which has not been used to reduce unnecessary charges and save cost. The policy becomes non-compliant when there is no instance attached to the EBS Volumes.		
Policy Recommendation	**Description** EBS volumes which are not attached (State as available) are considered as Orphaned.It is recommended to address the orphaned EBS Volumes by deleting them to avoid unnecessary usage cost being incurred on your bill. In most cases, we recommend taking a EBS snapshot before deleting a EBS Volumes so that it can be restored if required. It is the responsibility of the user to decide which EBS Volumes are required and take appropriate action based on the data available on these Volumes. **Remediation Actions** • Delete EBS Volumes		

Resource Id	Cloud Account	Region
vol-0e47a859161478f52	AWS_Bagchi	us-east-1
vol-0fb9b3484ca0fbc33	AWS_Bagchi	us-east-1
vol-0045b166860804a3c	AWS_Bagchi	us-east-1
vol-076d8124573051729	AWS_Bagchi	us-east-1
vol-0f7c20f3ecb1f5fe9	AWS_Bagchi	us-east-1
vol-0ab237df0b9671462	AWS_Bagchi	us-east-1
vol-0abf5b337987a1692	AWS_Bagchi	us-east-1
vol-011c568079be4b7a4	AWS_Bagchi	us-east-1
vol-00482b3d8144fc181	AWS_Bagchi	us-east-1
vol-03a32df71101c03df	AWS_Bagchi	us-east-1
vol-094195b5f3842d3e4	AWS_Bagchi	us-east-1

Pillar Name	Sustainability	**Status**	Violation
Question	02. How do you align cloud resources to your demand?	**No. of violated resources**	1
Best Practice	03. Stop the creation and maintenance of unused assets	**Resource Category**	Network
Severity	Low	**Resource**	Internet_Gateways
Policy	Aws Audit Vpc Unused Internet Gateway Policy CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether any unused VPC resources is available in the AWS account. For a better management of VPC resources, all unused (detached) Internet Gateways and Egress-Only Internet Gateways is recommended to be removed from AWS VPC environment. This rule becomes non-compliant if unused VPCs are not removed in the AWS account.		
Policy Recommendation	It is recommended to Identify and remove any unused VPC Internet Gateways (IGWs) and VPC Egress-Only Internet Gateways (EIGWs) in order to adhere to best practices and to avoid approaching the service limit (by default, you are limited to 5 IGWs and 5 EIGWs per AWS region). An Internet Gateway/Egress-Only Internet Gateway is evaluated as unused when is not attached anymore to an AWS Virtual Private Cloud (VPC).		

Resource Id	Cloud Account	Region
igw-070e4c8901108b0bb	AWS_Bagchi	us-east-1

Pillar Name	Sustainability	**Status**	Violation
Question	03. How do you take advantage of software and architecture patterns to support your sustainability goals?	**No. of violated resources**	1
Best Practice	01. Optimize software and architecture for asynchronous and scheduled jobs	**Resource Category**	Network
Severity	Medium	**Resource**	Internet_Gateways
Policy	Aws Audit Vpc Unused Internet Gateway Policy CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether any unused VPC resources is available in the AWS account. For a better management of VPC resources, all unused (detached) Internet Gateways and Egress-Only Internet Gateways is recommended to be removed from AWS VPC environment. This rule becomes non-compliant if unused VPCs are not removed in the AWS account.		
Policy Recommendation	It is recommended to Identify and remove any unused VPC Internet Gateways (IGWs) and VPC Egress-Only Internet Gateways (EIGWs) in order to adhere to best practices and to avoid approaching the service limit (by default, you are limited to 5 IGWs and 5 EIGWs per AWS region). An Internet Gateway/Egress-Only Internet Gateway is evaluated as unused when is not attached anymore to an AWS Virtual Private Cloud (VPC).		

Resource Id	Cloud Account	Region
igw-070e4c8901108b0bb	AWS_Bagchi	us-east-1

Pillar Name	Sustainability	**Status**	Violation
Question	03. How do you take advantage of software and architecture patterns to support your sustainability goals?	**No. of violated resources**	11
Best Practice	02. Remove or refactor workload components with low or no use	**Resource Category**	Storage
Severity	Medium	**Resource**	Volumes
Policy	AWS EBS Volume Orphaned	**Resource Type**	EBS
Policy Description	This policy identifies the EBS Volumes which has not been used to reduce unnecessary charges and save cost. The policy becomes non-compliant when there is no instance attached to the EBS Volumes.		
Policy Recommendation	**Description** EBS volumes which are not attached (State as available) are considered as Orphaned.It is recommended to address the orphaned EBS Volumes by deleting them to avoid unnecessary usage cost being incurred on your bill. In most cases, we recommend taking a EBS snapshot before deleting a EBS Volumes so that it can be restored if required. It is the responsibility of the user to decide which EBS Volumes are required and take appropriate action based on the data available on these Volumes. **Remediation Actions** • Delete EBS Volumes		

Resource Id	Cloud Account	Region
vol-0e47a859161478f52	AWS_Bagchi	us-east-1
vol-0fb9b3484ca0fbc33	AWS_Bagchi	us-east-1
vol-0045b166860804a3c	AWS_Bagchi	us-east-1
vol-076d8124573051729	AWS_Bagchi	us-east-1
vol-0f7c20f3ecb1f5fe9	AWS_Bagchi	us-east-1
vol-0ab237df0b9671462	AWS_Bagchi	us-east-1
vol-0abf5b337987a1692	AWS_Bagchi	us-east-1
vol-011c568079be4b7a4	AWS_Bagchi	us-east-1
vol-00482b3d8144fc181	AWS_Bagchi	us-east-1
vol-03a32df71101c03df	AWS_Bagchi	us-east-1
vol-094195b5f3842d3e4	AWS_Bagchi	us-east-1

Pillar Name	Sustainability	**Status**	Violation
Question	03. How do you take advantage of software and architecture patterns to support your sustainability goals?	**No. of violated resources**	1
Best Practice	02. Remove or refactor workload components with low or no use	**Resource Category**	Network
Severity	Medium	**Resource**	Internet_Gateways
Policy	Aws Audit Vpc Unused Internet Gateway Policy CS Policy	**Resource Type**	VPC
Policy Description	This policy audits whether any unused VPC resources is available in the AWS account. For a better management of VPC resources, all unused (detached) Internet Gateways and Egress-Only Internet Gateways is recommended to be removed from AWS VPC environment. This rule becomes non-compliant if unused VPCs are not removed in the AWS account.		
Policy Recommendation	It is recommended to Identify and remove any unused VPC Internet Gateways (IGWs) and VPC Egress-Only Internet Gateways (EIGWs) in order to adhere to best practices and to avoid approaching the service limit (by default, you are limited to 5 IGWs and 5 EIGWs per AWS region). An Internet Gateway/Egress-Only Internet Gateway is evaluated as unused when is not attached anymore to an AWS Virtual Private Cloud (VPC).		

Resource Id	Cloud Account	Region
igw-070e4c8901108b0bb	AWS_Bagchi	us-east-1

Pillar Name	Sustainability	**Status**	Violation
Question	04. How do you take advantage of data management policies and patterns to support your sustainability goals?	**No. of violated resources**	1
Best Practice	02. Use technologies that support data access and storage patterns	**Resource Category**	Storage
Severity	Low	**Resource**	Volumes
Policy	AWS EBS Provisioned IOPS-(io2) SSD Recommend	**Resource Type**	EBS
Policy Description	This policy identifies workloads specific EBS Provisioned IOPS (io2) Volumes to save unnecessary cost spent due to undesirable Volume Type. The policy becomes non-compliant when the EBS volumes have utilization suitable for EBS Provisioned IOPS (io2) disks.		
Policy Recommendation	**Description** It is highly recommended to change the EBS Current Volume Type io1 to Provisioned (io2) as it is the latest generation Volume Type with the same cost and better performance than previous generation (io1). It is the responsibility of the user to decide whether making this change will have an impact based on the environment and business criticality of the data stored in EBS Volumes. It is also to be noted that Multi-Attach feature of io1 is not yet supported in io		

Resource Id	Cloud Account	Region
vol-07c6661b47a956bf7	AWS_Bagchi	us-east-1

Pillar Name	Sustainability	**Status**	Violation
Question	04. How do you take advantage of data management policies and patterns to support your sustainability goals?	**No. of violated resources**	1
Best Practice	04. Use elasticity and automation to expand block storage or file system	**Resource Category**	Storage
Severity	Medium	**Resource**	Volumes
Policy	AWS EBS Provisioned IOPS-(io2) SSD Recommend	**Resource Type**	EBS
Policy Description	This policy identifies workloads specific EBS Provisioned IOPS (io2) Volumes to save unnecessary cost spent due to undesirable Volume Type. The policy becomes non-compliant when the EBS volumes have utilization suitable for EBS Provisioned IOPS (io2) disks.		
Policy Recommendation	**Description** It is highly recommended to change the EBS Current Volume Type io1 to Provisioned (io2) as it is the latest generation Volume Type with the same cost and better performance than previous generation (io1). It is the responsibility of the user to decide whether making this change will have an impact based on the environment and business criticality of the data stored in EBS Volumes. It is also to be noted that Multi-Attach feature of io1 is not yet supported in io		

Resource Id	Cloud Account	Region
vol-07c6661b47a956bf7	AWS_Bagchi	us-east-1