



The Server Labs G-Cloud 10 Service Description

G-Cloud 10 Service Definition - Amazon Web Services

Ref: TSL/GCLOUD10/SERVICE_DESC_AWS

Issue: 1.0

Date: May 2018

For: **G-Cloud**

Table of Contents

| | |
|---|-----------|
| 1.0 INTRODUCTION..... | 18 |
| 2.0 SERVICE DEFINITION INTRODUCTION | 18 |
| 3.0 SERVICE DEFINITION – AMAZON ELASTIC COMPUTE CLOUD (AMAZON EC2) . | 18 |
| 3.1 Service Overview | 18 |
| 3.2 Backup/Restore and Disaster Recovery | 29 |
| 3.3 Service Constraints | 30 |
| 3.4 Technical Requirements | 30 |
| 4.0 SERVICE DEFINITION – AMAZON VIRTUAL PRIVATE CLOUD (AMAZON VPC)... | 31 |
| 4.1 Service Overview | 31 |
| 4.2 Backup/Restore and Disaster Recovery | 32 |
| 4.3 Service Constraints | 32 |
| 4.4 Technical Requirements | 32 |
| 5.0 SERVICE DEFINITION – AWS DIRECT CONNECT..... | 32 |
| 5.1 Service Overview | 32 |
| 5.2 Backup/Restore and Disaster Recovery | 33 |
| 5.3 Service Constraints | 34 |
| 5.4 Technical Requirements | 34 |
| 6.0 SERVICE DEFINITION – AMAZON SIMPLE STORAGE SERVICE (AMAZON S3)... | 34 |
| 6.1 Service Overview | 34 |
| 6.2 Backup/Restore and Disaster Recovery | 35 |
| 6.3 Service Constraints | 35 |
| 6.4 Technical Requirements | 35 |
| 7.0 SERVICE DEFINITION – AMAZON GLACIER..... | 36 |
| 7.1 Service Overview | 36 |
| 7.2 Backup/Restore and Disaster Recovery | 37 |
| 7.3 Service Constraints | 37 |
| 7.4 Technical Requirements | 37 |
| 8.0 SERVICE DEFINITION – AMAZON ELASTIC BLOCK STORE (AMAZON EBS) | 37 |
| 8.1 Service Overview | 37 |

| | | |
|-------------|--|-----------|
| 8.2 | Backup/Restore and Disaster Recovery | 38 |
| 8.3 | Service Constraints | 39 |
| 8.4 | Technical Requirements | 39 |
| 9.0 | SERVICE DEFINITION – AMAZON RELATIONAL DATABASE SERVICE (AMAZON RDS)..... | 39 |
| 9.1 | Service Overview | 40 |
| 9.2 | Backup/Restore and Disaster Recovery | 40 |
| 9.3 | Service Constraints | 43 |
| 9.4 | Technical Requirements | 43 |
| 9.5 | Service Overview | 44 |
| 9.6 | Backup/Restore and Disaster Recovery | 44 |
| 9.7 | Service Constraints | 44 |
| 9.8 | Technical Requirements | 45 |
| 10.0 | SERVICE DEFINITION – AMAZON CLOUD DIRECTORY..... | 45 |
| 10.1 | Service Overview | 45 |
| 10.2 | Backup/Restore and Disaster Recovery | 47 |
| 10.3 | Service Constraints | 47 |
| 10.4 | Technical Requirements | 47 |
| 11.0 | SERVICE DEFINITION – AMAZON CLOUDWATCH | 47 |
| 11.1 | Service Overview | 47 |
| 11.2 | Backup/Restore and Disaster Recovery | 48 |
| 11.3 | Service Constraints | 48 |
| 11.4 | Technical Requirements | 48 |
| 12.0 | SERVICE DEFINITION – AMAZON EMR | 49 |
| 12.1 | Service Overview | 49 |
| 12.2 | Backup/Restore and Disaster Recovery | 50 |
| 12.3 | Service Constraints | 50 |
| 12.4 | Technical Requirements | 50 |
| 13.0 | SERVICE DEFINITION – AWS CLOUDHSM | 50 |
| 13.1 | Service Overview | 50 |
| 13.2 | Backup/Restore and Disaster Recovery | 52 |
| 13.3 | Service Constraints | 52 |
| 13.4 | Technical Requirements | 52 |

| | |
|--|-----------|
| 14.0 SERVICE DEFINITION – AWS CLOUDTRAIL | 52 |
| 14.1 Service Overview | 52 |
| 14.2 Backup/Restore and Disaster Recovery | 53 |
| 14.3 Service Constraints | 54 |
| 14.4 Technical Requirements | 54 |
| 15.0 SERVICE DEFINITION – AWS CONFIG | 54 |
| 15.1 Service Overview | 54 |
| 15.2 Backup/Restore and Disaster Recovery | 55 |
| 15.3 Service Constraints | 55 |
| 15.4 Technical Requirements | 55 |
| 16.0 SERVICE DEFINITION – AMAZON MQ | 55 |
| 16.1 Service Overview | 55 |
| 16.2 Backup/Restore and Disaster Recovery | 56 |
| 16.3 Service Constraints | 57 |
| 16.4 Technical Requirements | 57 |
| 17.0 SERVICE DEFINITION – AMAZON SIMPLE NOTIFICATION SERVICE (AMAZON SNS) | |
| 57 | |
| 17.1 Service Overview | 57 |
| 17.2 Backup/Restore and Disaster Recovery | 58 |
| 17.3 Service Constraints | 58 |
| 17.4 Technical Requirements | 58 |
| 18.0 SERVICE DEFINITION – AWS KEY MANAGEMENT SERVICE (AWS KMS) | 59 |
| 18.1 Service Overview | 59 |
| 18.2 Backup/Restore and Disaster Recovery | 60 |
| 18.3 Service Constraints | 60 |
| 18.4 Technical Requirements | 60 |
| 19.0 SERVICE DEFINITION – AMAZON CLOUDSEARCH | 60 |
| 19.1 Service Overview | 60 |
| 19.2 Backup/Restore and Disaster Recovery | 62 |
| 19.3 Service Constraints | 62 |
| 19.4 Technical Requirements | 62 |
| 20.0 SERVICE DEFINITION – AMAZON REDSHIFT | 62 |
| 20.1 Service Overview | 62 |

| | | |
|-------------|---|-----------|
| 20.2 | Backup/Restore and Disaster Recovery | 64 |
| 20.3 | Service Constraints | 65 |
| 20.4 | Technical Requirements | 65 |
| 21.0 | SERVICE DEFINITION – AMAZON EC2 CONTAINER REGISTRY (AMAZON ECR) 65 | |
| 21.1 | Service Overview | 65 |
| 21.2 | Backup/Restore and Disaster Recovery | 66 |
| 21.3 | Service Constraints | 66 |
| 21.4 | Technical Requirements | 66 |
| 22.0 | SERVICE DEFINITION – AMAZON EC2 CONTAINER SERVICE (AMAZON ECS)... 67 | |
| 22.1 | Service Overview | 67 |
| 22.2 | Backup/Restore and Disaster Recovery | 68 |
| 22.3 | Service Constraints | 68 |
| 22.4 | Technical Requirements | 68 |
| 23.0 | SERVICE DEFINITION – AMAZON EKS (IN PREVIEW)..... 68 | |
| 23.1 | Service Overview | 68 |
| 23.2 | Backup/Restore and Disaster Recovery | 69 |
| 23.3 | Service Constraints | 69 |
| 23.4 | Technical Requirements | 69 |
| 24.0 | SERVICE DEFINITION – AWS FARGATE | 69 |
| 24.1 | Service Overview | 69 |
| 24.2 | Backup/Restore and Disaster Recovery | 71 |
| 24.3 | Service Constraints | 71 |
| 24.4 | Technical Requirements | 71 |
| 25.0 | SERVICE DEFINITION – AWS ELASTIC BEANSTALK | 71 |
| 25.1 | Service Overview | 71 |
| 25.2 | Backup/Restore and Disaster Recovery | 72 |
| 25.3 | Service Constraints | 72 |
| 25.4 | Technical Requirements | 72 |
| 26.0 | SERVICE DEFINITION – AWS LAMBDA..... 73 | |
| 26.1 | Service Overview | 73 |
| 26.2 | Backup/Restore and Disaster Recovery | 74 |
| 26.3 | Service Constraints | 74 |
| 26.4 | Technical Requirements | 74 |

| | |
|---|-----------|
| 27.0 SERVICE DEFINITION – AWS AUTO SCALING | 74 |
| 27.1 Service Overview | 74 |
| 27.2 Backup/Restore and Disaster Recovery | 75 |
| 27.3 Service Constraints | 75 |
| 27.4 Technical Requirements | 75 |
| 28.0 SERVICE DEFINITION – AMAZON API GATEWAY | 75 |
| 28.1 Service Overview | 76 |
| 28.2 Backup/Restore and Disaster Recovery | 77 |
| 28.3 Service Constraints | 77 |
| 28.4 Technical Requirements | 77 |
| 29.0 SERVICE DEFINITION – AMAZON ROUTE 53..... | 77 |
| 29.1 Service Overview | 77 |
| 29.2 Backup/Restore and Disaster Recovery | 79 |
| 29.3 Service Constraints | 79 |
| 29.4 Technical Requirements | 79 |
| 30.0 SERVICE DEFINITION – AWS STORAGE GATEWAY | 80 |
| 30.1 Service Overview | 80 |
| 30.2 Backup/Restore and Disaster Recovery | 81 |
| 30.3 Service Constraints | 81 |
| 30.4 Technical Requirements | 81 |
| 31.0 SERVICE DEFINITION – AMAZON CLOUDFRONT | 82 |
| 31.1 Service Overview | 82 |
| 31.2 Backup/Restore and Disaster Recovery | 83 |
| 31.3 Service Constraints | 83 |
| 31.4 Technical Requirements | 83 |
| 32.0 SERVICE DEFINITION – AWS WEB APPLICATION FIREWALL (AWS WAF) | 84 |
| 32.1 Service Overview | 84 |
| 32.2 Backup/Restore and Disaster Recovery | 85 |
| 32.3 Service Constraints | 85 |
| 32.4 Technical Requirements | 85 |
| 33.0 SERVICE DEFINITION – AWS ARTIFACT..... | 85 |
| 33.1 Service Overview | 85 |
| 33.2 Backup/Restore and Disaster Recovery | 86 |

| | | |
|-------------|---|-----------|
| 33.3 | Service Constraints | 86 |
| 33.4 | Technical Requirements | 86 |
| 34.0 | SERVICE DEFINITION – AWS MOBILE HUB | 86 |
| 34.1 | Service Overview | 86 |
| 34.2 | Backup/Restore and Disaster Recovery | 87 |
| 34.3 | Service Constraints | 87 |
| 34.4 | Technical Requirements | 87 |
| 35.0 | SERVICE DEFINITION – AWS IMPORT/ EXPORT DISK | 88 |
| 35.1 | Service Overview | 88 |
| 35.2 | Backup/Restore and Disaster Recovery | 88 |
| 35.3 | Service Constraints | 88 |
| 35.4 | Technical Requirements | 88 |
| 36.0 | SERVICE DEFINITION – AMAZON ELASTIC FILE SYSTEM (AMAZON EFS) | 89 |
| 36.1 | Service Overview | 89 |
| 36.2 | Backup/Restore and Disaster Recovery | 90 |
| 36.3 | Service Constraints | 90 |
| 36.4 | Technical Requirements | 90 |
| 37.0 | SERVICE DEFINITION – AMAZON AURORA | 91 |
| 37.1 | Service Overview | 91 |
| 37.2 | Backup/Restore and Disaster Recovery | 92 |
| 37.3 | Service Constraints | 92 |
| 37.4 | Technical Requirements | 92 |
| 38.0 | SERVICE DEFINITION – AWS DATABASE MIGRATION SERVICE | 92 |
| 38.1 | Service Overview | 93 |
| 38.2 | Backup/Restore and Disaster Recovery | 94 |
| 38.3 | Service Constraints | 94 |
| 38.4 | Technical Requirements | 94 |
| 39.0 | SERVICE DEFINITION – AWS MIGRATION HUB | 95 |
| 39.1 | Service Overview | 95 |
| 39.2 | Backup/Restore and Disaster Recovery | 96 |
| 39.3 | Service Constraints | 96 |
| 39.4 | Technical Requirements | 96 |
| 40.0 | SERVICE DEFINITION – AWS SERVER MIGRATION SERVICE (SMS) | 96 |

| | | |
|-------------|---|------------|
| 40.1 | Service Overview | 97 |
| 40.2 | Backup/Restore and Disaster Recovery | 97 |
| 40.3 | Service Constraints | 97 |
| 40.4 | Technical Requirements | 97 |
| 41.0 | SERVICE DEFINITION – AWS SNOWBALL | 98 |
| 41.1 | Service Overview | 98 |
| 41.2 | Backup/Restore and Disaster Recovery | 99 |
| 41.3 | Service Constraints | 100 |
| 41.4 | Technical Requirements | 100 |
| 42.0 | SERVICE DEFINITION – AWS SNOWBALL EDGE | 100 |
| 42.1 | Service Overview | 100 |
| 42.2 | Backup/Restore and Disaster Recovery | 101 |
| 42.3 | Technical Requirements | 101 |
| 43.0 | SERVICE DEFINITION – AWS SNOWMOBILE | 102 |
| 43.1 | Service Overview | 102 |
| 43.2 | Backup/Restore and Disaster Recovery | 103 |
| 43.3 | Service Constraints | 103 |
| 43.4 | Technical Requirements | 103 |
| 44.0 | SERVICE DEFINITION – AMAZON DYNAMODB | 104 |
| 44.1 | Service Overview | 104 |
| 44.2 | Backup/Restore and Disaster Recovery | 104 |
| 44.3 | Service Constraints | 105 |
| 44.4 | Technical Requirements | 105 |
| 45.0 | SERVICE DEFINITION – AMAZON ELASTICACHE | 105 |
| 45.1 | Service Overview | 106 |
| 45.2 | Backup/Restore and Disaster Recovery | 107 |
| 45.3 | Service Constraints | 107 |
| 45.4 | Technical Requirements | 107 |
| 46.0 | SERVICE DEFINITION – AWS DIRECTORY SERVICE | 108 |
| 46.1 | Service Overview | 108 |
| 46.2 | Backup/Restore and Disaster Recovery | 109 |
| 46.3 | Service Constraints | 109 |
| 46.4 | Technical Requirements | 109 |

| | |
|---|------------|
| 47.0 SERVICE DEFINITION – AMAZON INSPECTOR | 110 |
| 47.1 Service Overview | 110 |
| 47.2 Backup/Restore and Disaster Recovery | 111 |
| 47.3 Service Constraints | 111 |
| 47.4 Technical Requirements | 111 |
| 48.0 SERVICE DEFINITION – AMAZON GUARDDUTY | 112 |
| 48.1 Service Overview | 112 |
| 48.2 Backup/Restore and Disaster Recovery | 113 |
| 48.3 Service Constraints | 113 |
| 48.4 Technical Requirements | 113 |
| 49.0 SERVICE DEFINITION – AMAZON MACIE..... | 114 |
| 49.1 Service Overview | 114 |
| 49.2 Backup/Restore and Disaster Recovery | 115 |
| 49.3 Service Constraints | 115 |
| 49.4 Technical Requirements | 115 |
| 50.0 SERVICE DEFINITION – AWS CERTIFICATE MANAGER..... | 115 |
| 50.1 Service Overview | 115 |
| 50.2 Backup/Restore and Disaster Recovery | 117 |
| 50.3 Service Constraints | 117 |
| 50.4 Technical Requirements | 117 |
| 51.0 SERVICE DEFINITION – AWS OPSWORKS | 118 |
| 51.1 Service Overview | 118 |
| 51.2 Backup/Restore and Disaster Recovery | 119 |
| 51.3 Service Constraints | 119 |
| 51.4 Technical Requirements | 119 |
| 52.0 SERVICE DEFINITION – AWS SYSTEMS MANAGER | 120 |
| 52.1 Service Overview | 120 |
| 52.2 Backup/Restore and Disaster Recovery | 122 |
| 52.3 Service Constraints | 122 |
| 52.4 Technical Requirements | 122 |
| 53.0 SERVICE DEFINITION – AWS CODECOMMIT | 122 |
| 53.1 Service Overview | 123 |
| 53.2 Backup/Restore and Disaster Recovery | 123 |

| | | |
|-------------|---|------------|
| 53.3 | Service Constraints | 123 |
| 53.4 | Technical Requirements | 123 |
| 54.0 | SERVICE DEFINITION – AWS CODEDEPLOY | 124 |
| 54.1 | Service Overview | 124 |
| 54.2 | Backup/Restore and Disaster Recovery | 125 |
| 54.3 | Service Constraints | 125 |
| 54.4 | Technical Requirements | 125 |
| 55.0 | SERVICE DEFINITION – AWS CODEPIPELINE | 125 |
| 55.1 | Service Overview | 126 |
| 55.2 | Backup/Restore and Disaster Recovery | 126 |
| 55.3 | Service Constraints | 126 |
| 55.4 | Technical Requirements | 127 |
| 56.0 | SERVICE DEFINITION – AWS X-RAY | 127 |
| 56.1 | Service Overview | 127 |
| 56.2 | Backup/Restore and Disaster Recovery | 128 |
| 56.3 | Service Constraints | 128 |
| 56.4 | Technical Requirements | 128 |
| 57.0 | SERVICE DEFINITION – AWS SERVICE CATALOG | 128 |
| 57.1 | Service Overview | 129 |
| 57.2 | Backup/Restore and Disaster Recovery | 129 |
| 57.3 | Service Constraints | 129 |
| 57.4 | Technical Requirements | 129 |
| 58.0 | CLOUDFORMATION | 130 |
| 58.1 | Service Overview | 130 |
| 58.2 | Backup/Restore and Disaster Recovery | 131 |
| 58.3 | Service Constraints | 132 |
| 58.4 | Technical Requirements | 132 |
| 59.0 | SERVICE DEFINITION – AMAZON QUICKSIGHT | 132 |
| 59.1 | Service Overview | 132 |
| 59.2 | Backup/Restore and Disaster Recovery | 133 |
| 59.3 | Service Constraints | 133 |
| 59.4 | Technical Requirements | 133 |
| 60.0 | SERVICE DEFINITION – AMAZON MACHINE LEARNING | 134 |

| | | |
|------|---|-----|
| 60.1 | Service Overview | 134 |
| 60.2 | Backup/Restore and Disaster Recovery | 135 |
| 60.3 | Service Constraints | 135 |
| 60.4 | Technical Requirements | 135 |
| 61.0 | SERVICE DEFINITION – AMAZON TRANSLATE | 136 |
| 61.1 | Service Overview | 136 |
| 61.2 | Backup/Restore and Disaster Recovery | 136 |
| 61.3 | Service Constraints | 136 |
| 61.4 | Technical Requirements | 137 |
| 62.0 | SERVICE DEFINITION – AMAZON TRANSCRIBE | 137 |
| 62.1 | Service Overview | 137 |
| 62.2 | Backup/Restore and Disaster Recovery | 138 |
| 62.3 | Service Constraints | 138 |
| 62.4 | Technical Requirements | 138 |
| 63.0 | SERVICE DEFINITION – AWS DEEPLENS | 138 |
| 63.1 | Service Overview | 139 |
| 63.2 | Backup/Restore and Disaster Recovery | 140 |
| 63.3 | Service Constraints | 140 |
| 63.4 | Technical Requirements | 140 |
| 64.0 | SERVICE DEFINITION – AMAZON TRANSCRIBE | 140 |
| 64.1 | Service Overview | 140 |
| 64.2 | Backup/Restore and Disaster Recovery | 141 |
| 64.3 | Service Constraints | 141 |
| 64.4 | Technical Requirements | 142 |
| 65.0 | SERVICE DEFINITION – AMAZON KINESIS | 142 |
| 65.1 | Service Overview | 142 |
| 65.2 | Backup/Restore and Disaster Recovery | 143 |
| 65.3 | Service Constraints | 143 |
| 65.4 | Technical Requirements | 143 |
| 66.0 | SERVICE DEFINITION – AMAZON ELASTIC TRANSCODER | 144 |
| 66.1 | Service Overview | 144 |
| 66.2 | Backup/Restore and Disaster Recovery | 145 |
| 66.3 | Service Constraints | 145 |

| | | |
|-------------|--|------------|
| 66.4 | Technical Requirements | 146 |
| 67.0 | SERVICE DEFINITION – AWS ELEMENTAL MEDIA CONVERT | 146 |
| 67.1 | Service Overview | 146 |
| 67.2 | Backup/Restore and Disaster Recovery | 147 |
| 67.3 | Service Constraints | 147 |
| 67.4 | Technical Requirements | 147 |
| 68.0 | SERVICE DEFINITION – AWS ELEMENTAL MEDIA LIVE | 147 |
| 68.1 | Service Overview | 147 |
| 68.2 | Backup/Restore and Disaster Recovery | 148 |
| 68.3 | Service Constraints | 149 |
| 68.4 | Technical Requirements | 149 |
| 69.0 | SERVICE DEFINITION – AWS ELEMENTAL MEDIA PACKAGE | 149 |
| 69.1 | Service Overview | 149 |
| 69.2 | Backup/Restore and Disaster Recovery | 150 |
| 69.3 | Service Constraints | 150 |
| 69.4 | Technical Requirements | 150 |
| 70.0 | SERVICE DEFINITION – AWS ELEMENTAL MEDIA STORE | 150 |
| 70.1 | Service Overview | 151 |
| 70.2 | Backup/Restore and Disaster Recovery | 151 |
| 70.3 | Service Constraints | 151 |
| 70.4 | Technical Requirements | 151 |
| 71.0 | SERVICE DEFINITION – AWS ELEMENTAL MEDIA TAILOR | 152 |
| 71.1 | Service Overview | 152 |
| 71.2 | Backup/Restore and Disaster Recovery | 153 |
| 71.3 | Service Constraints | 153 |
| 71.4 | Technical Requirements | 153 |
| 72.0 | SERVICE DEFINITION – AMAZON ELASTICSEARCH SERVICE | 153 |
| 72.1 | Service Overview | 153 |
| 72.2 | Backup/Restore and Disaster Recovery | 154 |
| 72.3 | Service Constraints | 155 |
| 72.4 | Technical Requirements | 155 |
| 73.0 | SERVICE DEFINITION – AWS DATA PIPELINE | 155 |
| 73.1 | Service Overview | 155 |

| | | |
|------|---|------------|
| 73.2 | Backup/Restore and Disaster Recovery | 156 |
| 73.3 | Service Constraints | 156 |
| 73.4 | Technical Requirements | 156 |
| 74.0 | SERVICE DEFINITION – AWS GLUE | 157 |
| 74.1 | Service Overview | 157 |
| 74.2 | Backup/Restore and Disaster Recovery | 158 |
| 74.3 | Service Constraints | 158 |
| 74.4 | Technical Requirements | 158 |
| 75.0 | SERVICE DEFINITION – AMAZON SIMPLE EMAIL SERVICE (AMAZON SES)..... | 158 |
| 75.1 | Service Overview | 159 |
| 75.2 | Backup/Restore and Disaster Recovery | 159 |
| 75.3 | Service Constraints | 159 |
| 75.4 | Technical Requirements | 159 |
| 76.0 | SERVICE DEFINITION – AMAZON SIMPLE QUEUE SERVICE (AMAZON SQS)... | 160 |
| 76.1 | Service Overview | 160 |
| 76.2 | Backup/Restore and Disaster Recovery | 161 |
| 76.3 | Service Constraints | 161 |
| 76.4 | Technical Requirements | 161 |
| 77.0 | SERVICE DEFINITION – AWS APPSYNC | 161 |
| 77.1 | Service Overview | 162 |
| 77.2 | Backup/Restore and Disaster Recovery | 163 |
| 77.3 | Service Constraints | 163 |
| 77.4 | Technical Requirements | 163 |
| 78.0 | SERVICE DEFINITION – AMAZON SIMPLE WORKFLOW SERVICE (AMAZON SWF) | |
| | 163 | |
| 78.1 | Service Overview | 163 |
| 78.2 | Backup/Restore and Disaster Recovery | 164 |
| 78.3 | Service Constraints | 164 |
| 78.4 | Technical Requirements | 164 |
| 79.0 | SERVICE DEFINITION – ALEXA FOR BUSINESS | 165 |
| 79.1 | Service Overview | 165 |
| 79.2 | Backup/Restore and Disaster Recovery | 166 |
| 79.3 | Service Constraints | 166 |

| | | |
|-------------|--|------------|
| 79.4 | Technical Requirements | 166 |
| 80.0 | SERVICE DEFINITION – AMAZON APPSTREAM 2.0 | 166 |
| 80.1 | Service Overview | 167 |
| 80.2 | Backup/Restore and Disaster Recovery | 168 |
| 80.3 | Service Constraints | 168 |
| 80.4 | Technical Requirements | 168 |
| 81.0 | SERVICE DEFINITION – AWS GREENGRASS | 168 |
| 81.1 | Service Overview | 168 |
| 81.2 | Backup/Restore and Disaster Recovery | 170 |
| 81.3 | Service Constraints | 170 |
| 81.4 | Technical Requirements | 170 |
| 82.0 | SERVICE DEFINITION – AMAZON COGNITO | 170 |
| 82.1 | Service Overview | 170 |
| 82.2 | Backup/Restore and Disaster Recovery | 171 |
| 82.3 | Service Constraints | 171 |
| 82.4 | Technical Requirements | 172 |
| 83.0 | SERVICE DEFINITION – AWS DEVICE FARM | 172 |
| 83.1 | Service Overview | 172 |
| 83.2 | Backup/Restore and Disaster Recovery | 173 |
| 83.3 | Service Constraints | 173 |
| 83.4 | Technical Requirements | 173 |
| 84.0 | SERVICE DEFINITION – AWS IOT | 173 |
| 84.1 | Service Overview | 173 |
| 84.2 | Backup/Restore and Disaster Recovery | 174 |
| 84.3 | Service Constraints | 174 |
| 84.4 | Technical Requirements | 174 |
| 85.0 | SERVICE DEFINITION – AWS IOT CORE | 175 |
| 85.1 | Service Overview | 175 |
| 85.2 | Backup/Restore and Disaster Recovery | 176 |
| 85.3 | Service Constraints | 176 |
| 85.4 | Technical Requirements | 176 |
| 86.0 | SERVICE DEFINITION – AMAZON FREERTOS | 176 |
| 86.1 | Service Overview | 176 |

| | | |
|-------------|---|------------|
| 86.2 | Backup/Restore and Disaster Recovery | 178 |
| 86.3 | Service Constraints | 178 |
| 86.4 | Technical Requirements | 178 |
| 87.0 | SERVICE DEFINITION – AWS IOT 1-CLICK | 178 |
| 87.1 | Service Overview | 178 |
| 87.2 | Backup/Restore and Disaster Recovery | 179 |
| 87.3 | Service Constraints | 179 |
| 87.4 | Technical Requirements | 180 |
| 88.0 | SERVICE DEFINITION – AWS IOT ANALYTICS | 180 |
| 88.1 | Service Overview | 180 |
| 88.2 | Backup/Restore and Disaster Recovery | 182 |
| 88.3 | Service Constraints | 182 |
| 88.4 | Technical Requirements | 182 |
| 89.0 | SERVICE DEFINITION – AWS IOT BUTTON..... | 182 |
| 89.1 | Service Overview | 183 |
| 90.0 | SERVICE DEFINITION – AWS IOT DEVICE DEFENDER (IN PREVIEW)..... | 183 |
| 90.1 | Service Overview | 183 |
| 91.0 | SERVICE DEFINITION – AWS IOT DEVICE MANAGEMENT..... | 185 |
| 91.1 | Service Overview | 185 |
| 91.2 | Backup/Restore and Disaster Recovery | 186 |
| 91.3 | Service Constraints | 186 |
| 91.4 | Technical Requirements | 186 |
| 92.0 | SERVICE DEFINITION – AMAZON WORKSPACES | 186 |
| 92.1 | Service Overview | 186 |
| 92.2 | Backup/Restore and Disaster Recovery | 187 |
| 92.3 | Service Constraints | 187 |
| 92.4 | Technical Requirements | 187 |
| 93.0 | SERVICE DEFINITION – AMAZON WORKMAIL | 188 |
| 93.1 | Service Overview | 188 |
| 93.2 | Backup/Restore and Disaster Recovery | 189 |
| 93.3 | Service Constraints | 189 |
| 93.4 | Technical Requirements | 189 |
| 94.0 | SERVICE DEFINITION – AMAZON WORKDOCS..... | 189 |

| | | |
|--------------|---|------------|
| 94.1 | Service Overview | 189 |
| 94.2 | Backup/Restore and Disaster Recovery | 190 |
| 94.3 | Service Constraints | 190 |
| 94.4 | Technical Requirements | 191 |
| 95.0 | SERVICE DEFINITION – AWS SERVER MIGRATION SERVICE | 191 |
| 95.1 | Service Overview | 191 |
| 95.2 | Backup/Restore and Disaster Recovery | 191 |
| 95.3 | Service Constraints | 192 |
| 95.4 | Technical Requirements | 192 |
| 96.0 | SERVICE DEFINITION – AMAZON LIGHTSAIL | 192 |
| 96.1 | Service Overview | 192 |
| 96.2 | Backup/Restore and Disaster Recovery | 193 |
| 96.3 | Service Constraints | 193 |
| 96.4 | Technical Requirements | 193 |
| 97.0 | SERVICE DEFINITION – AWS BATCH | 193 |
| 97.1 | Service Overview | 194 |
| 97.2 | Backup/Restore and Disaster Recovery | 194 |
| 97.3 | Service Constraints | 194 |
| 97.4 | Technical Requirements | 194 |
| 98.0 | SERVICE DEFINITION – AWS SERVERLESS APPLICATION REPOSITORY | 195 |
| 98.1 | Service Overview | 195 |
| 98.2 | Backup/Restore and Disaster Recovery | 196 |
| 98.3 | Service Constraints | 196 |
| 98.4 | Technical Requirements | 196 |
| 99.0 | SERVICE DEFINITION – AWS ORGANIZATIONS | 196 |
| 99.1 | Service Overview | 196 |
| 99.2 | Backup/Restore and Disaster Recovery | 197 |
| 99.3 | Service Constraints | 197 |
| 99.4 | Technical Requirements | 197 |
| 100.0 | SERVICE DEFINITION – AWS SHIELD | 198 |
| 100.1 | Service Overview | 198 |
| 100.2 | Backup/Restore and Disaster Recovery | 199 |
| 100.3 | Service Constraints | 199 |

| | | |
|--------------|---|------------|
| 100.4 | Technical Requirements..... | 200 |
| 101.0 | SERVICE DEFINITION – AWS APPLICATION DISCOVERY SERVICE | 200 |
| 101.1 | Service Overview | 200 |
| 101.2 | Backup/Restore and Disaster Recovery | 201 |
| 101.3 | Service Constraints | 201 |
| 101.4 | Technical Requirements..... | 201 |
| 102.0 | SERVICE DEFINITION – AWS CODEBUILD..... | 202 |
| 102.1 | Service Overview | 202 |
| 102.2 | Backup/Restore and Disaster Recovery | 202 |
| 102.3 | Service Constraints | 203 |
| 102.4 | Technical Requirements..... | 203 |
| 103.0 | SERVICE DEFINITION – AWS CODESTAR..... | 203 |
| 103.1 | Service Overview | 203 |
| 103.2 | Backup/Restore and Disaster Recovery | 204 |
| 103.3 | Service Constraints | 204 |
| 103.4 | Technical Requirements..... | 205 |
| 104.0 | SERVICE DEFINITION – AWS STEP FUNCTIONS..... | 205 |
| 104.1 | Service Overview | 205 |
| 104.2 | Backup/Restore and Disaster Recovery | 206 |
| 104.3 | Service Constraints | 206 |
| 104.4 | Technical Requirements..... | 206 |
| 105.0 | SERVICE DEFINITION – AMAZON ATHENA..... | 206 |
| 105.1 | Service Overview | 207 |
| 105.2 | Backup/Restore and Disaster Recovery | 207 |
| 105.3 | Service Constraints | 207 |
| 105.4 | Technical Requirements..... | 208 |
| 106.0 | SERVICE DEFINITION – AMAZON PINPOINT..... | 208 |
| 106.1 | Service Overview | 208 |
| 106.2 | Backup/Restore and Disaster Recovery | 209 |
| 106.3 | Service Constraints | 209 |
| 106.4 | Technical Requirements..... | 209 |
| 107.0 | SERVICE DEFINITION – AMAZON SAGEMAKER | 210 |
| 107.1 | Service Overview | 210 |

| | | |
|--------------|--|------------|
| 107.2 | Backup/Restore and Disaster Recovery | 211 |
| 107.3 | Service Constraints | 211 |
| 107.4 | Technical Requirements..... | 211 |
| 108.0 | SERVICE DEFINITION – AMAZON COMPREHEND..... | 212 |
| 108.1 | Service Overview | 212 |
| 108.2 | Backup/Restore and Disaster Recovery | 213 |
| 108.3 | Service Constraints | 213 |
| 108.4 | Technical Requirements..... | 213 |
| 109.0 | SERVICE DEFINITION – AMAZON LEX | 213 |
| 109.1 | Service Overview | 213 |
| 109.2 | Backup/Restore and Disaster Recovery | 215 |
| 109.3 | Service Constraints | 215 |
| 109.4 | Technical Requirements..... | 215 |
| 110.0 | SERVICE DEFINITION – AMAZON POLLY | 215 |
| 110.1 | Service Overview | 215 |
| 110.2 | Backup/Restore and Disaster Recovery | 217 |
| 110.3 | Service Constraints | 217 |
| 110.4 | Technical Requirements..... | 217 |
| 111.0 | SERVICE DEFINITION – AMAZON REKOGNITION | 217 |
| 111.1 | Service Overview | 217 |
| 111.2 | Backup/Restore and Disaster Recovery | 218 |
| 111.3 | Service Constraints | 219 |
| 111.4 | Technical Requirements..... | 219 |
| 112.0 | SERVICE DEFINITION – AMAZON CHIME..... | 219 |
| 112.1 | Service Overview | 219 |
| 112.2 | Backup/Restore and Disaster Recovery | 219 |
| 112.3 | Service Constraints | 220 |
| 112.4 | Technical Requirements..... | 220 |
| 113.0 | SERVICE DEFINITION – AMAZON CONNECT..... | 220 |
| 113.1 | Service Overview | 220 |
| 113.2 | Backup/Restore and Disaster Recovery | 221 |
| 113.3 | Service Constraints | 221 |
| 113.4 | Technical Requirements..... | 221 |

| | | |
|--------------|---|------------|
| 114.0 | SERVICE DEFINITION – AMAZON LUMBERYARD | 222 |
| 114.1 | Service Overview | 222 |
| 114.2 | Backup/Restore and Disaster Recovery | 222 |
| 114.3 | Service Constraints | 222 |
| 114.4 | Technical Requirements..... | 223 |
| 115.0 | SERVICE DEFINITION – AMAZON SUMERIAN (IN PREVIEW) | 223 |
| 115.1 | Service Overview | 223 |
| 115.2 | Backup/Restore and Disaster Recovery | 224 |
| 115.3 | Service Constraints | 224 |
| 115.4 | Technical Requirements..... | 224 |
| 116.0 | SERVICE DEFINITION – AMAZON GAMELIFT | 225 |
| 116.1 | Service Overview | 225 |
| 116.2 | Backup/Restore and Disaster Recovery | 225 |
| 116.3 | Service Constraints | 226 |
| 116.4 | Technical Requirements..... | 226 |
| 117.0 | CROSS SERVICES..... | 227 |
| 117.1 | Availability | 227 |
| 117.2 | On-Boarding/Off-Boarding Processes and Service Migration | 236 |
| 117.3 | Service Management Details | 237 |
| 117.4 | Service Levels and Service Credits | 238 |
| 117.1 | Trial Service Details | 238 |
| 117.2 | Termination Terms..... | 238 |
| 117.3 | Data Restoration/Service Migration | 238 |
| 117.4 | Customer Responsibilities | 245 |
| 118.0 | SELECTED CUSTOMER REFERENCES | 249 |

1.0 Introduction

This document provides you with a description of each of The Server Labs' AWS Cloud Hosting Services.

If you wish to receive further information please contact sales@theserverlabs.com

The Server Labs are an AWS Consulting Partner and Channel Reseller

2.0 Service Definition Introduction

This document provides service definitions for the Amazon Web Services, Inc. (AWS) service offerings included in the G-Cloud 10 framework catalogue (Offered Services). We have broken out service definitions in accordance with Invitation to Tender (ITT) requirements. Please note that we have consolidated common elements of each offered service (e.g. onboarding and off-boarding) and have provided service definitions for these common elements that apply equally to each Offered Service.

3.0 Service Definition – Amazon Elastic Compute Cloud (Amazon EC2)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

3.1 Service Overview

Amazon EC2 is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers. Amazon EC2's simple interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use. Amazon EC2 provides developers the tools to build failure resilient applications and isolate them from common failure scenarios.

Top benefits include:

- **Elastic, Web-Scale Computing** – Amazon EC2 enables you to increase or decrease capacity within minutes, not hours or days. You can commission one, hundreds, or even thousands of server instances simultaneously.
- **Completely Controlled** – You have complete control of your instances. You have root access to each one, and you can interact with them as you would any machine.
- **Flexible Cloud Hosting Services** – You have the choice of multiple instance types, operating systems, and software packages. Amazon EC2 allows you to select a configuration of memory, CPU, instance storage, and the boot partition size that is optimal for your choice of operating system and application.
- **Designed for Use with Other AWS Cloud Services** – Amazon EC2 works in conjunction with Amazon S3, Amazon RDS, and other AWS Cloud services to provide a complete solution for computing, query processing, and storage across a wide range of applications.
- **Reliable** – Amazon EC2 offers a highly reliable environment where replacement instances can be rapidly and predictably commissioned. The service runs within Amazon’s proven network infrastructure and data centres.
- **Secure** – Amazon EC2 works in conjunction with Amazon VPC to provide security and robust networking functionality for your compute resources.
- **Inexpensive** – Amazon EC2 passes on to you the financial benefits of Amazon’s scale. You pay a very low rate for the compute capacity you actually consume. We offer three different ways to buy instances, each with their own cost benefits:
 - **On-Demand Instances** – On-Demand Instances let you pay for compute capacity by the hour with no long-term commitments. This frees you from the costs and complexities of planning, purchasing, and maintaining hardware and transforms what are commonly large, fixed costs into much smaller variable costs. On-Demand Instances also remove the need to buy “safety net” capacity to handle periodic traffic spikes.
 - **Reserved Instances** – A Reserved Instance provides you with a significant discount (up to 75%) compared to On-Demand Instance pricing. There are three Reserved Instance payment options—No Upfront, Partial Upfront, and All Upfront—that enable you to balance the amount you pay up front with your effective hourly price. The Reserved Instance Marketplace is also available, which provides you with the opportunity to sell Reserved Instances if your needs change (e.g., want to move instances to a new AWS Region, change to a new instance type, or sell capacity for projects that end before your Reserved Instance term expires).
 - **Spot Instances** – Spot Instances allow customers to bid on unused Amazon EC2 capacity and run those instances for as long as their bid exceeds the current Spot Price. The Spot Price changes periodically based on supply and demand, and customers whose bids meet or exceed it gain access to the available Spot Instances. If you have flexibility in when your applications can run, Spot Instances can significantly lower your Amazon EC2 costs.

3.1.1 Instance types

Amazon EC2 provides a wide selection of instance types optimised to fit different use cases. Instance types comprise varying combinations of CPU, memory, storage, and networking capacity and give you the flexibility to choose the appropriate mix of resources for your applications. Each instance type includes one or more instance sizes, allowing you to scale your resources to the requirements of your target workload.

3.1.1.1 General Purpose

T2 – [T2 instances](#) are [Burstable Performance Instances](#) that provide a baseline level of CPU performance with the ability to burst above the baseline. T2 Unlimited instances can sustain high CPU performance for as long as a workload needs it. For most general-purpose workloads, T2 Unlimited instances will provide ample performance without any additional charges. If the instance needs to run at higher CPU utilization for a prolonged period, it can also do so at a flat additional charge of 5 cents per vCPU-hour.

The baseline performance and ability to burst are governed by CPU Credits. T2 instances receive CPU Credits continuously at a set rate depending on the instance size, accumulating CPU Credits when they are idle, and consuming CPU credits when they are active. T2 instances are a good choice for a variety of general-purpose workloads including micro-services, low-latency interactive applications, small and medium databases, virtual desktops, development, build and stage environments, code repositories, and product prototypes. For more information see [Burstable Performance Instances](#). Features include:

- High frequency Intel Xeon processors
- Burstable CPU, governed by CPU Credits, and consistent baseline performance
- Lowest-cost general purpose instance type, and Free Tier eligible*
- Balance of compute, memory, and network resources

M5 - [M5 instances](#) are the latest generation of General Purpose Instances. This family provides a balance of compute, memory, and network resources, and it is a good choice for many applications. Features include:

- 2.5 GHz Intel Xeon® Platinum 8175 processors with new Intel Advanced Vector Extension (AVX-512) instruction set
- New larger instance size, m5.24xlarge, offering 96 vCPUs and 384 GiB of memory
- EBS-optimized by default and higher EBS performance on smaller instance sizes
- Up to 25 Gbps network bandwidth using Enhanced Networking
- Requires HVM AMIs that include drivers for ENA and NVMe
- Powered by the new light-weight Nitro system, a combination of dedicated hardware and lightweight hypervisor

M4 – M4 instances provide a balance of compute, memory, and network resources, and it is a good choice for many applications. Features include:

- 2.3 GHz Intel Xeon® E5-2686 v4 (Broadwell) processors or 2.4 GHz Intel Xeon® E5-2676 v3 (Haswell) processors
- EBS-optimized by default at no additional cost
- Support for Enhanced Networking
- Balance of compute, memory, and network resources

3.1.1.2 Compute-Optimised

C5 - [C5 instances](#) are optimized for compute-intensive workloads and deliver very cost-effective high performance at a low price per compute ratio. Features include:

- 3.0 GHz Intel Xeon Platinum processors with new Intel Advanced Vector Extension 512 (AVX-512) instruction set
- Run each core at up to 3.5 GHz using Intel Turbo Boost Technology
- New larger instance size, c5.18xlarge, offering 72 vCPUs and 144 GiB of memory
- Up to 25 Gbps of network bandwidth using Elastic Network Adapter (ENA)-based Enhanced Networking
- EBS optimized by default
- Requires HVM AMIs that include drivers for ENA and NVMe

C4 – C4 instances are optimized for compute-intensive workloads and deliver very cost-effective high performance at a low price per compute ratio. Features include:

- High frequency Intel Xeon E5-2666 v3 (Haswell) processors optimized specifically for EC2
- Default EBS-optimized for increased storage performance at no additional cost
- Higher networking performance with Enhanced Networking supporting Intel 82599 VF
- Requires Amazon VPC, Amazon EBS and 64-bit HVM AMIs

3.1.1.3 Memory Optimised

X1e - [X1e instances](#) are optimized for high-performance databases, in-memory databases and other memory intensive enterprise applications. X1e instances offer one of the lowest price per GiB of RAM among Amazon EC2 instance types. Features include:

- High frequency Intel Xeon E7-8880 v3 (Haswell) processors
- One of the lowest price per GiB of RAM
- Up to 3,904 GiB of DRAM-based instance memory
- SSD storage and EBS-optimized by default and at no additional cost
- Ability to control processor C-state and P-state configurations on x1e.32xlarge, x1e.16xlarge and x1e.8xlarge instances

X1 - [X1 instances](#) are optimized for large-scale, enterprise-class and in-memory applications, and offer one of the lowest price per GiB of RAM among Amazon EC2 instance types. Features include:

- High frequency Intel Xeon E7-8880 v3 (Haswell) processors
- One of the lowest price per GiB of RAM
- Up to 1,952 GiB of DRAM-based instance memory
- SSD storage and EBS-optimized by default and at no additional cost
- Ability to control processor C-state and P-state configuration

R4 - R4 instances are optimized for memory-intensive applications and offer better price per GiB of RAM than R3. Features include:

- High Frequency Intel Xeon E5-2686 v4 (Broadwell) processors
- DDR4 Memory
- Support for [Enhanced Networking](#)

3.1.1.4 Accelerated Computing Instances

P3 - [P3 instances](#) are the latest generation of general purpose GPU instances. Features include:

- Up to 8 NVIDIA Tesla V100 GPUs, each pairing 5,120 CUDA Cores and 640 Tensor Cores
- High frequency Intel Xeon E5-2686 v4 (Broadwell) processors
- Supports NVLink for peer-to-peer GPU communication
- Provide Enhanced Networking using Elastic Network Adapter with up to 25 Gbps of aggregate network bandwidth within a Placement Group

P2 – P2 instances are intended for general-purpose GPU compute applications. Features include:

- High frequency Intel Xeon E5-2686 v4 (Broadwell) processors
- High-performance NVIDIA K80 GPUs, each with 2,496 parallel processing cores and 12GiB of GPU memory
- Supports GPUDirect™ for peer-to-peer GPU communications
- Provides [Enhanced Networking](#) using the Amazon EC2 Elastic Network Adaptor with up to 20Gbps of aggregate network bandwidth within a Placement Group
- Amazon EBS-optimised by default at no additional cost

G3 – [G3 instances](#) are optimized for graphics-intensive applications. Features include:

- High frequency Intel Xeon E5-2686 v4 (Broadwell) processors
- NVIDIA Tesla M60 GPUs, each with 2048 parallel processing cores and 8 GiB of video memory
- Enables NVIDIA GRID Virtual Workstation features, including support for 4 monitors with resolutions up to 4096x2160. Each GPU included in your instance is licensed for one “Concurrent Connected User”
- Enables NVIDIA GRID Virtual Application capabilities for application virtualization software like Citrix XenApp Essentials and VMware Horizon, supporting up to 25 concurrent users per GPU
- Each GPU features an on-board hardware video encoder designed to support up to 10 H.265 (HEVC) 1080p30 streams and up to 18 H.264 1080p30 streams, enabling low-latency frame capture and encoding, and high-quality interactive streaming experiences
- Enhanced Networking using the Elastic Network Adapter (ENA) with 25 Gbps of aggregate network bandwidth within a Placement Group

F1 – [F1 instances](#) offer customizable hardware acceleration with field programmable gate arrays (FPGAs). Features include:

- Instances Features:
 - High frequency Intel Xeon E5-2686 v4 (Broadwell) processors
 - NVMe SSD Storage
 - Support for Enhanced Networking
- FPGA Features:
 - Xilinx Virtex UltraScale+ VU9P FPGAs
 - 64 GiB of ECC-protected memory on 4x DDR4
 - Dedicated PCI-Express x16 interface
 - Approximately 2.5 million logic elements
 - Approximately 6,800 Digital Signal Processing (DSP) engines
 - [FPGA Developer AMI](#)

3.1.1.5 Storage-Optimised

H1 - [H1 instances](#) feature up to 16 TB of HDD-based local storage, deliver high disk throughput, and a balance of compute and memory. Features include:

- Powered by 2.3 GHz Intel® Xeon® E5 2686 v4 processors (codenamed Broadwell)
- Up to 16TB of HDD storage
- High disk throughput
- ENA enabled Enhanced Networking up to 25 Gbps

I3 – [I3](#) is instance family provides Non-Volatile Memory Express (NVMe) SSD-backed Instance storage optimized for low latency, very high random I/O performance, high sequential read throughput and provide high IOPS at a low cost. Features include:

- High Frequency Intel Xeon E5-2686 v4 (Broadwell) Processors with base frequency of 2.3 GHz
- Up to 25 Gbps of network bandwidth using Elastic Network Adapter (ENA)-based Enhanced Networking
- High Random I/O performance and High Sequential Read throughput

D2 – D2 instances feature up to 48 TB of HDD-based local storage, deliver high disk throughput, and offer the lowest price per disk throughput performance on Amazon EC2. Features include:

- High-frequency Intel Xeon E5-2676 v3 (Haswell) processors

- HDD storage
- Consistent high performance at launch time
- High disk throughput
- Support for Enhanced Networking

3.1.1.6 Previous Generation Instances

AWS offers Previous Generation Instances for users who have optimized their applications around these instances and have yet to upgrade. Previous Generation Instances are still fully supported and retain the same features and functionality. Previous Generation Instances are available through the AWS Management Console, AWS CLI, and EC2 API tools. For more information, see [Previous Generation Instances](#).

3.1.1.7 Instance Features

Amazon EC2 instances provide a number of additional features to help you deploy, manage, and scale your applications.

Burstable Performance Instances – Amazon EC2 allows you to choose between Fixed Performance Instances (e.g. M3, C3, and R3) and Burstable Performance Instances (e.g. T2). Burstable Performance Instances provide a baseline level of CPU performance with the ability to burst above the baseline.

T2 Unlimited instances can sustain high CPU performance for as long as a workload needs it. For most general-purpose workloads, T2 Unlimited instances will provide ample performance without any additional charges. The hourly T2 instance price automatically covers all interim spikes in usage when the average CPU utilization of a T2 instance is at or less than the baseline over a 24-hour window. If the instance needs to run at higher CPU utilization for a prolonged period, it can do so at a flat additional charge of 5 cents per vCPU-hour.

T2 instances' baseline performance and ability to burst are governed by CPU Credits. Each T2 instance receives CPU Credits continuously, the rate of which depends on the instance size. T2 instances accrue CPU Credits when they are idle, and use CPU

credits when they are active. A CPU Credit provides the performance of a full CPU core for one minute.

For example, a t2.small instance receives credits continuously at a rate of 12 CPU Credits per hour. This capability provides baseline performance equivalent to 20% of a CPU core (20% x 60 mins = 12 mins). If the instance does not use the credits it receives, they are stored in its CPU Credit balance up to a maximum of 288 CPU Credits. When the t2.small instance needs to burst to more than 20% of a core, it draws from its CPU Credit balance to handle this surge automatically.

With T2 Unlimited enabled, the t2.small instance can burst above the baseline even after its CPU Credit balance is drawn down to zero. For a vast majority of general purpose workloads where the average CPU utilization is at or below the baseline performance, the basic hourly price for t2.small covers all CPU bursts. If the instance happens to run at an average 25% CPU utilization (5% above baseline) over a period of 24 hours after its CPU Credit balance is drawn to zero, it will be charged an additional 6 cents (5 cents/vCPU-hour x 1 vCPU x 5% x 24 hours).

Many applications such as web servers, developer environments and small databases don't need consistently high levels of CPU, but benefit significantly from having full access to very fast CPUs when they need them. T2 instances are engineered specifically for these use cases. If you need consistently high CPU performance for applications such as video encoding, high volume websites or HPC applications, we recommend you use Fixed Performance Instances. T2 instances are designed to perform as if they have dedicated high speed Intel cores available when your application really needs CPU performance, while protecting you from the variable performance or other common side-effects you might typically see from over-subscription in other environments.

Multiple Storage Options – Amazon EC2 allows you to choose between multiple storage options based on your requirements. [Amazon EBS](#) is a durable, block-level storage volume that you can attach to a single, running Amazon EC2 instance. You can use Amazon EBS as a primary storage device for data that requires frequent and

granular updates. For example, Amazon EBS is the recommended storage option when you run a database on Amazon EC2. Amazon EBS volumes persist independently from the running life of an Amazon EC2 instance. Once a volume is attached to an instance you can use it like any other physical hard drive. Amazon EBS provides three volume types to best meet the needs of your workloads: General Purpose (SSD), Provisioned IOPS (SSD), and Magnetic. General Purpose (SSD) is the new, SSD-backed, general purpose EBS volume type that we recommend as the default choice for customers. General Purpose (SSD) volumes are suitable for a broad range of workloads, including small to medium sized databases, development and test environments, and boot volumes. Provisioned IOPS (SSD) volumes offer storage with consistent and low-latency performance, and are designed for I/O intensive applications such as large relational or NoSQL databases. Magnetic volumes provide the lowest cost per gigabyte of all EBS volume types. Magnetic volumes are ideal for workloads where data is accessed infrequently, and applications where the lowest storage cost is important.

Many Amazon EC2 instances can also include storage from disks that are physically attached to the host computer. This disk storage is referred to as instance store. Instance store provides temporary block-level storage for Amazon EC2 instances. The data on an instance store volume persists only during the life of the associated Amazon EC2 instance.

In addition to block level storage via Amazon EBS or instance store, you can also use Amazon S3 for highly durable, highly available object storage. Learn more about Amazon EC2 storage options from the [Amazon EC2 documentation](#).

EBS-Optimised Instances – For an additional, low, hourly fee, customers can launch selected Amazon EC2 instances types as EBS-optimized instances. For C5, C4, M5, M4, P3, P2, G3, and D2 instances, this feature is enabled by default at no additional cost. EBS-optimized instances enable EC2 instances to fully use the IOPS provisioned on an EBS volume. EBS-optimized instances deliver dedicated throughput between Amazon EC2 and Amazon EBS, with options between 500 and 4,000 Megabits per second (Mbps) depending on the instance type used. The dedicated throughput

minimizes contention between Amazon EBS I/O and other traffic from your EC2 instance, providing the best performance for your EBS volumes. EBS-optimized instances are designed for use with both Standard and Provisioned IOPS Amazon EBS volumes. When attached to EBS-optimized instances, Provisioned IOPS volumes can achieve single digit millisecond latencies and are designed to deliver within 10% of the provisioned IOPS performance 99.9% of the time. We recommend using Provisioned IOPS volumes with EBS-optimized instances or instances that support cluster networking for applications with high storage I/O requirements.

EBS-Optimised Instances – For an additional, low, hourly fee, customers can launch selected Amazon EC2 instances types as EBS-optimised instances. For C4, M4, P2, and D2 instances, this feature is enabled by default at no additional cost. Amazon EBS-optimised instances enable Amazon EC2 instances to fully use the IOPS provisioned on an Amazon EBS volume. Amazon EBS-optimised instances deliver dedicated throughput between Amazon EC2 and Amazon EBS, with options between 500 and 4,000 Megabits per second (Mbps) depending on the instance type used. The dedicated throughput minimises contention between Amazon EBS I/O and other traffic from your Amazon EC2 instance, providing the best performance for your EBS volumes. EBS-optimised instances are designed for use with both Standard and Provisioned IOPS Amazon EBS volumes. When attached to EBS-optimised instances, Provisioned IOPS volumes can achieve single digit millisecond latencies and are designed to deliver within 10% of the provisioned IOPS performance 99.9% of the time. We recommend using Provisioned IOPS volumes with Amazon EBS-optimised instances or instances that support cluster networking for applications with high storage I/O requirements.

Cluster Networking – Select EC2 instances support cluster networking when launched into a common cluster placement group. A cluster placement group provides low-latency networking between all instances in the cluster. The bandwidth an EC2 instance can utilize depends on the instance type and its networking performance specification. Inter instance traffic within the same region can utilize up to 5 Gbps for single-flow and up to 25 Gbps for multi-flow traffic in each direction (full duplex). Traffic to and from S3

buckets in the same region can also utilize all available instance aggregate bandwidth. When launched in a placement group, instances can utilize up to 10 Gbps for single-flow traffic and up to 25 Gbps for multi-flow traffic. Network traffic to the Internet is limited to 5 Gbps (full duplex). Cluster networking is ideal for high performance analytics systems and many science and engineering applications, especially those using the MPI library standard for parallel programming.

Dedicated Instances – [Dedicated Instances](#) are Amazon EC2 instances that run in a VPC on hardware that's dedicated to a single customer. They are ideal for workloads where corporate policies or industry regulations require that your Amazon EC2 instances be physically isolated at host hardware level from instances that belong to other AWS accounts. Dedicated Instances let you take full advantage of the benefits of the AWS Cloud: on-demand elastic provisioning, pay only for what you use, all while ensuring that your Amazon EC2 compute instances are isolated at the hardware level. You can also use [Dedicated Hosts](#) to launch Amazon EC2 instances on physical servers that are dedicated for your use. Dedicated Hosts give you additional visibility and control over how instances are placed on a physical server, and you can reliably use the same physical server over time. As a result, Dedicated Hosts enable you to use your existing server-bound software licenses like Windows Server and address corporate compliance and regulatory requirements. Visit this page to [compare Dedicated Instances and Dedicated Hosts](#).

3.2 Backup/Restore and Disaster Recovery

Traditional enterprise backup and recovery strategies typically take an agent-based approach whereby the entire contents of a server are backed up over either the Local Area Network (LAN) or the Storage Area Network (SAN). Traditional architectures have required this approach because replacing failed components is complex, time-consuming, and operationally intensive. This has, in turn, created a backup environment that is complex to manage and resource-intensive to operate—requiring technologies

such as data de-duplication and virtual tape libraries to cope with ever-increasing workloads.

Amazon EC2 enables the full backup and recovery of a standard server, such as a web server or application server, so you can focus on protecting configuration and stateful data rather than on the server itself. This set of data is much smaller than the aggregate set of server data, which typically includes various application files, operating system files, temporary files, and so on. This change of approach means that regular nightly incremental or weekly full backups can take far less time and consume less storage space.

When a compute instance is started in Amazon EC2, it is based upon an Amazon Machine Image (AMI) and can connect to existing storage volumes (e.g., Amazon EBS). In addition, when launching a new instance, it is possible to pass user data to the instance so that it can be accessed internally as dynamic configuration parameters.

3.3 Service Constraints

Please see <http://aws.amazon.com/documentation/ec2/> for more information on service constraints for Amazon EC2.

3.4 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/ec2/> and the following links for comprehensive technical documentation regarding Amazon EC2.

- **Linux Guide** – Describes key concepts of Amazon EC2 and provides instructions for using the features of Amazon EC2. Available in [HTML](#), [PDF](#), and [Kindle formats](#).
- **CLI Reference** – Documents the Amazon EC2 CLI. Available in [HTML](#) and [PDF formats](#).
- **Amazon EC2 Section of the AWS CLI Reference** – Describes the AWS CLI commands that you can use to administer Amazon EC2. Provides syntax, options, and usage examples for each command. Available in [HTML](#) format.
- **Windows Guide** – Describes key concepts for Amazon EC2 and provides instructions for launching and using your Windows instance. Available in [HTML](#), [PDF](#), and [Kindle format](#).
- **API Reference** – Documents the Amazon EC2 Query API. Available in [HTML](#) and [PDF format](#).
- **Amazon EC2 Simple Systems Manager (SSM) API Reference** – Documents the SSM API. Available in [HTML](#) and [PDF](#) format.

4.0 Service Definition – Amazon Virtual Private Cloud (Amazon VPC)

The following subsections provide service definition information.

4.1 Service Overview

Amazon VPC lets you provision a logically isolated section of the AWS Cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways.

You can easily customise the network configuration for your Amazon VPC. For example, you can create a public-facing subnet for your web servers that has access to the Internet and place your backend systems such as databases or application servers in a private-facing subnet with no Internet access. You can leverage multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet.

Additionally, you can create a hardware Virtual Private Network (VPN) connection between your corporate data centre and your VPC and leverage the AWS Cloud as an extension of your corporate data centre.

Top benefits include:

- **Multiple Connectivity Options** – A variety of connectivity options exist for your Amazon VPC. You can connect your VPC to the Internet, to your data centre, or to other VPCs, based on the AWS resources that you want to expose publicly and those that you want to keep private.
- **Secure** – Amazon VPC provides advanced security features such as security groups and network access control lists to enable inbound and outbound filtering at the instance level and subnet level. In addition, you can store data in Amazon S3 and restrict access so that it is only accessible from instances in your VPC.
- **Simple** – You can create a VPC quickly and easily using the AWS Management Console.
- **All the Scalability and Reliability of AWS** – Amazon VPC provides all the same benefits as the rest of the AWS platform. You can instantly scale your resources up or down, select Amazon EC2 instance types and sizes that are right for your applications, and pay only for the resources you use—all within Amazon's proven infrastructure.

4.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon VPC. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation>.

4.3 Service Constraints

Please see <http://aws.amazon.com/documentation/vpc/> for more information.

4.4 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/vpc/> and the following links for comprehensive technical documentation regarding Amazon VPC.

- **Getting Started Guide** – Provides instructions to create a VPC and launch an instance into the Amazon VPC. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **CLI Reference** – Documents the Amazon VPC CLI. Available in [HTML](#) and [PDF](#) formats.
- **API Reference** – Documents the Amazon VPC Query API. Available in [HTML](#) and [PDF](#) formats.
- **User Guide** – Describes key concepts for Amazon VPC and provides instructions for using the features of Amazon VPC. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **Network Administrator Guide** – Describes customer gateways and helps network administrators configure them. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **Quick Reference Card** – Briefly covers the essential commands for using Amazon VPC from the CLI. Available in [PDF](#) format.

5.0 Service Definition – AWS Direct Connect

The following subsections provide service definition information.

5.1 Service Overview

AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your data centre, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections.

AWS Direct Connect lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using industry-standard 802.1q Virtual LANS (VLANs), this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space and private resources such as Amazon EC2 instances running within an Amazon VPC using private IP space, all while maintaining network separation between the public and private environments. Virtual interfaces can be reconfigured at any time to meet your changing needs.

Top benefits include:

- **Reduces Your Bandwidth Costs** – If you have bandwidth-heavy workloads that you wish to run in AWS, AWS Direct Connect reduces your network costs into and out of AWS in two ways. First, by transferring data to and from AWS directly, you can reduce your bandwidth commitment to your Internet service provider. Second, all data transferred over your dedicated connection is charged at the reduced AWS Direct Connect data transfer rate rather than Internet data transfer rates.
- **Consistent Network Performance** – Network latency over the Internet can vary because the Internet is constantly changing how data gets from point A to B. With AWS Direct Connect, you choose the data that uses the dedicated connection and how that data is routed, which can provide a more consistent network experience over Internet-based connections.
- **Compatible with all AWS Cloud Services** – AWS Direct Connect is a network service and works with all AWS Cloud services that are accessible over the Internet, such as Amazon S3, Amazon EC2, and Amazon VPC.
- **Private Connectivity to Your Amazon VPC** – You can use AWS Direct Connect to establish a private virtual interface from your on-premises network directly to your Amazon VPC, providing you with a private, high-bandwidth network connection between your network and your VPC.
- **Elastic** – AWS Direct Connect makes it easy to scale your connection to meet your needs.
- **Simple** – You can sign up for AWS Direct Connect quickly and easily using the [AWS Management Console](#).

5.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Direct Connect. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation>.

5.3 Service Constraints

Please see <http://aws.amazon.com/documentation/direct-connect/> for more information.

5.4 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/directconnect/> and the following links for comprehensive technical documentation regarding AWS Direct Connect.

- **User Guide** – Provides a conceptual overview of AWS Direct Connect and includes instructions on using the various features with the CLI. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **API Reference** – Describes all the API operations for AWS Direct Connect in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#) formats.

6.0 Service Definition – Amazon Simple Storage Service (Amazon S3)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

6.1 Service Overview

Amazon S3 provides developers and IT teams with secure, durable, and highly scalable object storage. Amazon S3 is easy to use, with a simple interface to store and retrieve any amount of data from anywhere on the web. With Amazon S3, you pay only for the storage you actually use. There is no minimum fee and no setup cost.

Amazon S3 can be used alone or together with other AWS Cloud services such as Amazon EC2, Amazon EBS, and Amazon Glacier, as well as third-party storage repositories and gateways. Amazon S3 provides cost-effective object storage for a wide variety of use cases, including cloud applications, content distribution, backup and archiving, disaster recovery, and big data analytics.

Top benefits include:

- **Durable** – Amazon S3 provides durable infrastructure to store important data and is designed for durability of 99.999999999% of objects.

- **Low Cost** – Amazon S3 allows you to store large amounts of data at a very low cost. You pay for what you need, with no minimum commitments or up-front fees.
- **Available** – Amazon S3 is designed for 99.99% availability of objects over a given year.
- **Secure** – Amazon S3 supports data transfer over SSL and automatic encryption of your data once it is uploaded.
- **Scalable** – With Amazon S3, you can store as much data as you want and access it when you need it.
- **Send Event Notifications** – Amazon S3 can send event notifications when objects are uploaded to Amazon S3.
- **High Performance** – Amazon S3 supports multi-part uploads to help maximise network throughput and resiliency and lets you choose the AWS Region in which to store your data, minimising network latency.
- **Integrated** – Amazon S3 is integrated with other AWS Cloud services to simplify uploading and downloading data from Amazon S3 and to make it easier to build solutions that use a range of AWS Cloud services.
- **Easy to Use** – Amazon S3 is easy to use with a web-based management console and mobile app and full Representational State Transfer (REST) APIs and SDKs for easy integration with third-party technologies.

6.2 Backup/Restore and Disaster Recovery

Amazon S3 offers a highly durable, scalable, and secure solution for backing up and archiving your critical data. You can use Amazon S3's versioning capability to provide even further protection for your stored data. You can also define rules to archive sets of Amazon S3 objects to Amazon Glacier's extremely low-cost storage service based on object lifetimes. As your data ages, these rules enable you to ensure that it is automatically stored on the storage option that is most cost effective for your needs.

6.3 Service Constraints

Please see <http://aws.amazon.com/documentation/s3/> for more information.

6.4 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/s3/> and the following links for comprehensive technical documentation regarding Amazon S3.

- **Getting Started Guide** – Introduces you to Amazon S3, helps you set up an account, and walks you through a simple example to help you use Amazon

- S3 for the first time. Also provides tips and links to advanced product features and resources. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **API Reference** – Describes all the Amazon S3 API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#) formats.
 - **Quick Reference Card** – Briefly covers the essential commands for using Amazon S3 from the CLI. Available in [PDF](#) format.
 - **Developer Guide** – Provides a conceptual overview of Amazon S3 and includes detailed instructions for using the various features. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
 - **Console User Guide** – Provides information to help you use Amazon S3 with the AWS Management Console. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.

7.0 Service Definition – Amazon Glacier

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

7.1 Service Overview

Amazon Glacier is a secure, durable, and extremely low-cost storage service for data archiving and online backup. Customers can reliably store large or small amounts of data. To keep costs low yet suitable for varying retrieval needs, Amazon Glacier provides three options for access to archives, from a few minutes to several hours.

Top benefits include:

- **Low Cost** – Amazon Glacier allows you to archive large amounts of data at a very low cost.
- **Secure** – Amazon Glacier supports data transfer over SSL and automatically encrypts your data at rest.
- **Durable** – Amazon Glacier provides a highly durable storage infrastructure designed for online backup and archival. Your data is redundantly stored across multiple facilities and multiple devices in each facility.
- **Simple** – Amazon Glacier allows you to offload the administrative burden of operating storage infrastructure to AWS.
- **Flexible** – Amazon Glacier scales to meet your storage needs. There is no limit to how much data you can store, and you can choose to store your data in the AWS Region that supports your regulatory and business criteria.
- **Integrated** – Through Amazon S3 life cycle policies, you can optimise your storage costs by moving infrequently accessed objects from Amazon S3 to Amazon Glacier (or vice versa).

7.2 Backup/Restore and Disaster Recovery

Amazon S3 enables you to use Amazon Glacier's extremely low-cost storage service as a storage option for data archival. Amazon Glacier is optimised for data that is infrequently accessed and for which retrieval times of several hours are suitable. Examples include digital media archives, financial and healthcare records, raw genomic sequence data, long-term database backups, and data that must be retained for regulatory compliance.

For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation>.

7.3 Service Constraints

Please see <http://aws.amazon.com/documentation/glacier/> for more information.

7.4 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/glacier> and the following links for comprehensive technical documentation regarding Amazon Glacier.

- **Developer Guide** – Provides detailed information about setting up and working with Amazon Glacier using the REST API and the AWS SDKs for Java and Microsoft .NET. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.

8.0 Service Definition – Amazon Elastic Block Store (Amazon EBS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

8.1 Service Overview

Amazon EBS provides persistent, available, and durable block-level storage volumes for use with Amazon EC2 instances in the AWS Cloud. Amazon EBS volumes offer the consistent and low-latency performance needed to run your workloads. With Amazon EBS, you can scale your usage up or down within minutes, all while paying a low price for only what you provision.

Top benefits include:

- **Reliable, Secure Storage** – Each Amazon EBS volume is automatically replicated within its Availability Zone to protect you from component failure.
- **Consistent and Low-Latency Performance** – Amazon EBS General Purpose volumes and Amazon EBS Provisioned IOPS volumes deliver low-latency through SSD technology and consistent I/O performance scaled to the needs of your application.
- **Backup, Restore, and Innovate** – Back up your data by taking point-in-time snapshots of your Amazon EBS volumes. Boost the agility of your business by using Amazon EBS snapshots to create new Amazon EC2 instances.
- **Geographic Flexibility** – Amazon EBS provides the ability to copy snapshots across AWS Regions, enabling geographical expansion, data centre migration, and disaster recovery.
- **Quickly Scale Up and Easily Scale Down** – Increase or decrease block storage and performance within minutes, enjoying the freedom to adjust as your needs evolve.

8.2 Backup/Restore and Disaster Recovery

An Amazon EBS snapshot is a point-in-time backup copy of an Amazon EBS volume that is stored in Amazon S3. Snapshots are incremental backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved. When you delete a snapshot, only the data exclusive to that snapshot is removed. Active snapshots contain all of the information needed to restore your data (from the time the snapshot was taken) to a new Amazon EBS volume.

When you create a new Amazon EBS volume, you can create it based on an existing snapshot; the new volume begins as an exact replica of the original volume that was used to create the snapshot. New volumes created from existing Amazon S3 snapshots load lazily in the background, so you can begin using them right away. If your instance accesses a piece of data that hasn't yet been loaded, the volume immediately downloads the requested data from Amazon S3 and then continues loading the rest of the volume's data in the background. For more information about creating snapshots, see the [Creating an Amazon EBS Snapshot](#) user guide.

You can share your snapshots with specific individuals or make them public to share them with the entire AWS community. Users with access to your snapshots can create their own Amazon EBS volumes from your snapshot, but your snapshots remain

completely intact. For more information about how to share snapshots, see the [Sharing Snapshots](#) user guide.

Amazon EBS snapshots are constrained to the region in which they are created. Once you have created a snapshot of an Amazon EBS volume, you can use it to create new volumes in the same region. For more information, see the [Restoring an Amazon EBS Volume from a Snapshot](#) user guide. You can also copy snapshots across AWS Regions, making it easier to leverage multiple AWS Regions for geographical expansion, data centre migration, and disaster recovery. You can copy any accessible snapshots that are in the “available” status. For more information, see the [Copying an Amazon EBS Snapshot](#) user guide.

8.3 Service Constraints

Please see <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html> for more information.

8.4 Technical Requirements

Please see <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AmazonEBS.html> and the following links for comprehensive technical documentation regarding Amazon EBS.

- [Creating an Amazon EBS Snapshot](#)
- [Deleting an Amazon EBS Snapshot](#)
- [Copying an Amazon EBS Snapshot](#)
- [Describing Snapshots](#)
- [Sharing Snapshots](#)

9.0 Service Definition – Amazon Relational Database Service (Amazon RDS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

9.1 Service Overview

Amazon RDS is a web service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while managing time-consuming database management tasks, freeing you up to focus on your applications and business.

Top benefits include:

- **Easy to Administer** – Amazon RDS makes it easy to go from project conception to deployment. Use the [AWS Management Console](#), the [Amazon RDS CLI](#), or simple [API calls](#) to access the capabilities of a production-ready relational database in minutes. There is no need for infrastructure provisioning and no need for installing and maintaining database software.
- **Scalable** – You can scale your database's compute and storage resources with only a few mouse clicks or an API call, often with no downtime.
- **Available and Durable** – Amazon RDS runs on the same highly reliable infrastructure used by other AWS Cloud services. When you provision a Multi-AZ Database Instance (DB Instance), Amazon RDS synchronously replicates the data to a standby instance in a different Availability Zone. Amazon RDS has many other features that enhance reliability for critical production databases, including automated backups, Database Snapshots (DB Snapshots), and automatic host replacement.
- **Fast** – Amazon RDS offers database server sizing choices up to 32 vCPUs and 244 GiB, as well as storage choices for a wide range of application performance requirements.
- **Secure** – Amazon RDS makes it easy to control network access to your database. Amazon RDS also lets you run your database instances in Amazon VPC, which enables you to isolate your database instances and to connect to your existing IT infrastructure through an industry-standard encrypted IPsec VPN. Many Amazon RDS engine types offer encryption at rest and encryption in transit.
- **Inexpensive** – You pay very low rates and only for the resources you actually consume. In addition, you benefit from the option of on-demand pricing with no up-front or long-term commitments or even lower hourly rates via our reserved pricing option.

9.2 Backup/Restore and Disaster Recovery

The following sections provide information regarding backup/restore and disaster recovery for Amazon RDS.

9.2.1 Database Instance Backups

Amazon RDS provides two different methods for backing up and restoring your Amazon database instances: automated backups and DB Snapshots. Automated backups automatically back up your DB instance during a specific, user-definable backup window and keeps the backups for a limited, user-specified period of time (called the backup retention period); you can later recover your database to any point in time during that retention period. DB Snapshots are user-initiated backups that enable you to back up your DB instance to a known state and to restore to that specific state at any time. Amazon RDS keeps all DB Snapshots until you delete them. A brief I/O freeze, typically lasting a few seconds, occurs during both automated backups and DB Snapshot operations on Single-AZ DB instances.

9.2.2 Automated Backup

Automated backup is an Amazon RDS feature that automatically creates a backup of your database. Automated backups are enabled by default for a new DB instance. An automated backup occurs during a daily user-configurable period of time known as the preferred backup window. Backups created during the preferred backup window are retained for a user-configurable number of days (the backup retention period).

The preferred backup window is the user-defined period of time during which your DB instance is backed up. Amazon RDS uses these periodic data backups in conjunction with your transaction logs to enable you to restore your DB instance to any second during your retention period, up to the LatestRestorableTime (typically up to the last five minutes). During the backup window, storage I/O may be suspended while your data is being backed up and you may experience elevated latency. This I/O suspension typically lasts for the duration of the snapshot. This period of I/O suspension is shorter for Multi-AZ database deployments because the backup is taken from the standby, but latency can still occur during the backup process.

When the backup retention changes to a non-zero value, the first backup occurs immediately. Changing the backup retention period to 0 turns off automatic backups for the DB instance and deletes all existing automated backups for the instance. If you

don't specify a preferred backup window when you create the DB instance, Amazon RDS assigns a default 30-minute backup window that is selected at random from an 8-hour block of time per region. Changes to the backup window take effect immediately. The backup window cannot overlap with the weekly maintenance window for the DB instance.

When you delete a DB instance, you can create a final DB Snapshot upon deletion; if you take this action, you can use that DB Snapshot to restore the deleted DB instance at a later date. Amazon RDS retains this final, user-created DB Snapshot along with all other manually created DB Snapshots after the DB instance is deleted. All automated backups are deleted and cannot be recovered when you delete a DB instance.

For more information on working with automated backups, see the [Working With Automated Backups](#) user guide.

9.2.3 Point-In-Time Recovery

In addition to the daily automated backup, Amazon RDS archives database change logs. This enables you to recover your database to any point in time during the backup retention period, up to the last five minutes of database usage.

Amazon RDS stores multiple copies of your data, but for Single-AZ DB instances, these copies are stored in a single Availability Zone. If for any reason a Single-AZ DB instance becomes unusable, you can use point-in-time recovery to launch a new DB instance with the latest restorable data. For more information on working with point-in-time recovery, see the [Restoring a DB Instance to a Specified Time](#) user guide. Multi-AZ deployments store copies of your data in different Availability Zones for greater levels of data durability. For more information on Multi-AZ deployments, see the [High Availability \(Multi-AZ\)](#) user guide.

9.2.4 Automated Backups with Unsupported MySQL Storage Engines

Amazon RDS automated backups and DB Snapshots are currently supported for all database engines. For the MySQL database engine, only the InnoDB storage engine is supported; use of these features with other MySQL storage engines, including MyISAM,

may lead to unreliable behaviour while restoring from backups. Specifically, since storage engines like MyISAM do not support reliable crash recovery, your tables can be corrupted in the event of a crash. For this reason, we encourage you to use the InnoDB storage engine.

If you choose to use MyISAM, you can attempt to manually repair tables that become damaged after a crash by using the REPAIR command (refer to <http://dev.mysql.com/doc/refman/5.5/en/repair-table.html> for detailed instructions). However, as noted in the MySQL documentation, there is a good chance that you will not be able to recover all of your data.

9.2.5 DB Snapshots

DB Snapshots are user-initiated and enable you to back up your DB instance in a known state as frequently as you wish and then restore to that specific state at any time. DB Snapshots can be created with the Amazon RDS console or the `CreateDBSnapshot` action in the Amazon RDS API. DB Snapshots are kept until you explicitly delete them with the Amazon RDS console or the `DeleteDBSnapshot` action in the Amazon RDS API. For more information on working with DB Snapshots, see the [Creating a DB Snapshot](#) and [Restoring From a DB Snapshot](#) user guides.

Please refer to the [Amazon RDS Backing Up and Restoring](#) user guide for additional information.

9.3 Service Constraints

Please see <http://aws.amazon.com/documentation/rds/> for more information.

9.4 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/rds/> and the following links for comprehensive technical documentation regarding Amazon RDS.

- **User Guide** – Describes all Amazon RDS concepts and provides instructions on using the various features with both the console and the CLI. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.

- **API Reference** – Describes all the API operations for Amazon RDS in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#) formats.
- **CLI Reference** – Describes all the API operations for Amazon RDS in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#) formats.
- **Quick Reference Card** – Briefly covers the essential commands for using [Amazon RDS from the command line. Available in PDF format.](#)
- [Service Definition – AWS Identity and Access Management \(IAM\)](#)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

9.5 Service Overview

AWS IAM enables you to securely control access to AWS Cloud services and resources for your users. Using AWS IAM, you can create and manage AWS users and groups and use permissions to allow and deny their access to AWS resources.

With AWS IAM you can:

- **Manage AWS IAM Users and Their Access** – You can create users in AWS IAM, assign them individual security credentials (e.g., access keys, passwords, multi-factor authentication devices) or request temporary security credentials to provide users access to AWS Cloud services and resources. You can manage permissions in order to control which operations a user can perform.
- **Manage AWS IAM Roles and Their Permissions** – You can create roles in AWS IAM and manage permissions to control which operations can be performed by the entity or AWS Cloud service that assumes the role. You can also define which entity is allowed to assume the role.
- **Manage Federated Users and Their Permissions** – You can enable identity federation to allow existing identities in your enterprise to access the AWS Management Console, to call AWS APIs, and to access resources without the need to create an AWS IAM user for each identity.

9.6 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS IAM. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

9.7 Service Constraints

Please see <http://aws.amazon.com/documentation/iam/> for more information.

9.8 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/iam> and the following links for comprehensive technical documentation regarding AWS IAM.

- **Using AWS IAM** – Introduces you to AWS IAM, helps you set up an account, and walks you through a simple example to help you use AWS IAM for the first time. Also provides tips and links to advanced product features and resources. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **AWS IAM Section of AWS CLI Reference** – Describes the AWS CLI commands that you can use to administer AWS IAM. Provides syntax, options, and usage examples for each command. Available in [HTML](#) format.
- **AWS IAM API Reference** – Describes all the API operations for AWS IAM in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#) formats.
- **Security Token Service (STS) Section of AWS CLI Reference** – Describes the AWS CLI commands that you can use to generate temporary security credentials. Provides syntax, options, and usage examples for each command. Available in [HTML](#) format.
- **STS API Reference** – Describes all the API operations for AWS STS in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#) formats.

10.0 Service Definition – Amazon Cloud Directory

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

10.1 Service Overview

Amazon Cloud Directory enables you to build flexible cloud-native directories for organizing hierarchies of data along multiple dimensions. With Cloud Directory, you can create directories for a variety of use cases, such as organizational charts, course catalogs, and device registries. While traditional directory solutions, such as Active Directory Lightweight Directory Services (AD LDS) and other LDAP-based directories, limit you to a single hierarchy, Cloud Directory offers you the flexibility to create directories with hierarchies that span multiple dimensions. For example, you can create an organizational chart that can be navigated through separate hierarchies for reporting structure, location, and cost center.

Amazon Cloud Directory automatically scales to hundreds of millions of objects and provides an extensible schema that can be shared with multiple applications. As a fully-managed service, Cloud Directory eliminates time-consuming and expensive administrative tasks, such as scaling infrastructure and managing servers. You simply define the schema, create a directory, and then populate your directory by making calls to the Cloud Directory API.

Top benefits include::

- **Efficiently organize hierarchies of data across multiple dimensions** - Many applications require building directories to store hierarchies of data with multiple dimensions. For example, an organizational chart could have one hierarchy based on reporting structure, a second based on location, and a third based on cost center. Amazon Cloud Directory enables you to combine different schemas within a single directory to create multiple hierarchies without having to duplicate data.
- **Scale automatically on managed infrastructure** - Amazon Cloud Directory is designed to automatically scale to hundreds of millions of objects. Amazon built Cloud Directory on highly scalable, managed infrastructure that provides high availability and built-in server management. Unlike traditional databases and directories, Cloud Directory does not require upfront capacity planning or continuous, manual optimization.
- **Search your directory for objects and relationships** - Searching large sets of highly connected data has traditionally required complex and computationally expensive queries. With Amazon Cloud Directory's built-in search capabilities, you can search for all the parent objects along a dimension without creating multiple, nested queries. For example, looking up a management chain for an employee can be performed with a single query in Cloud Directory where traditional solutions would require multiple queries.
- **Easily adapt to changing data requirements** - Application data requirements evolve over time. Amazon Cloud Directory is built around a model that explicitly makes room for change. Cloud Directory has a flexible schema so you can extend your schema with new attributes at any time. Multiple applications can independently extend the schema. This enables applications to share a single directory and avoid duplication of data.
- **Simplify policy management** - With Amazon Cloud Directory, you can attach application-specific policies to objects and define policy inheritance rules along different hierarchies. For example, you can create a custom policy for an organizational chart to allocate budget limits through your cost center hierarchy, and enforce naming conventions for your device hierarchy. Cloud Directory enables your application to evaluate and enforce these policies on demand to meet the unique requirements of your application. You are not limited by the native policies of a directory.
- **Integrate with AWS CloudTrail and tagging** - Amazon Cloud Directory is integrated with AWS CloudTrail and resource tagging. Using AWS CloudTrail you can log the date, time, and identity of users who accesses your directory data. With resource tagging, you can tag your directories and schemas to better track and manage resources.

10.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Cloud Directory. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

10.3 Service Constraints

Please see <http://aws.amazon.com/documentation/iam/> for more information.

10.4 Technical Requirements

AWS Directory Service provides multiple ways to set up and run Amazon Cloud Directory, Amazon Cognito, and Microsoft AD with other AWS services. Amazon Cloud Directory provides a highly scalable directory store for your application's multihierarchical data. Amazon Cognito helps you create a directory store that authenticates your users either through your own user pools or through federated identity providers. AWS Directory Service for Microsoft Active Directory (Enterprise Edition), also known as Microsoft AD, enables your directory-aware workloads and AWS resources to use a managed Active Directory [in the AWS Cloud](#). Please refer to <https://aws.amazon.com/documentation/directory-service/> for more information on which directories you can choose from and additional documentation available on this service.

11.0 Service Definition – Amazon CloudWatch

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

11.1 Service Overview

Amazon CloudWatch is a monitoring service for AWS Cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, and set alarms. Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances and Amazon RDS DB instances, as well as custom metrics generated by your applications and services and any log files your applications generate. You can use Amazon CloudWatch to gain system-wide

visibility into resource utilisation, application performance, and operational health. You can use these insights to react and keep your application running smoothly.

Top benefits include:

- **Monitor Amazon EC2** – View metrics for CPU utilisation, data transfer, and disk usage activity from Amazon EC2 instances (Basic Monitoring) for no additional charge.
- **Monitor Other AWS Resources** – Monitor metrics on Amazon EBS volumes, Amazon RDS DB instances, Elastic Load Balancers, and more for no additional charge.
- **Monitor Custom Metrics** – Submit custom metrics generated by your own applications via a simple API request and have them monitored by Amazon CloudWatch.
- **Monitor and Store Logs** – You can use CloudWatch Logs to monitor and troubleshoot your systems and applications using your existing system, application, and custom log files.
- **Set Alarms** – Set alarms on any of your metrics to send you notifications or take other automated actions.
- **View Graphs and Statistics** – View graphs and statistics for any of your metrics on the Amazon CloudWatch dashboard and get a quick overview of all your alarms and monitored AWS resources in one location.

11.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon CloudWatch. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

11.3 Service Constraints

Please see <http://aws.amazon.com/documentation/cloudwatch/> for more information.

11.4 Technical Requirements

Please refer to <http://aws.amazon.com/documentation/cloudwatch/> and the following links for comprehensive technical documentation regarding Amazon CloudWatch.

- **Developer Guide** – Provides a conceptual overview of Amazon CloudWatch and includes detailed development instructions for using the various features. Available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **CLI Reference** – Describes the Amazon CloudWatch CLI commands that you can use to monitor AWS. Provides syntax, options, and usage examples for each command. Available in [HTML](#) and [PDF](#) formats.
- **Amazon CloudWatch Section of AWS CLI Reference** – Describes the Amazon CloudWatch commands in the AWS CLI that you can use to monitor AWS. Provides syntax, options, and usage examples for each command. Available in [HTML](#) format.

- **API Reference** – Describes all the API operations for Amazon CloudWatch in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#) format.
- **CloudWatch Logs API Reference** – Describes all the API operations for Amazon CloudWatch Logs in detail. Also provides sample requests, responses, and errors for the supported web services protocols. Available in [HTML](#) and [PDF](#).
- **CloudWatch Logs Section of AWS CLI Reference** – Describes the AWS CLI commands that you can use to administer CloudWatch Logs. Provides syntax, options, and usage examples for each command. Available in [HTML](#) format.

12.0 Service Definition – Amazon EMR

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

12.1 Service Overview

Amazon EMR provides a managed Hadoop framework that makes it easy, fast, and cost-effective to process vast amounts of data across dynamically scalable Amazon EC2 instances. You can also run other popular distributed frameworks such as [Apache Spark](#), [HBase](#), [Presto](#), and [Flink](#) in Amazon EMR, and interact with data in other AWS data stores such as Amazon S3 and Amazon DynamoDB.

Amazon EMR securely and reliably handles a broad set of big data use cases, including log analysis, web indexing, data transformations (Extract, Transform, Load [ETL]), ML, financial analysis, scientific simulation, and bioinformatics.

Top features include:

- **Easy to Use** – You can launch an Amazon EMR cluster in minutes. You don't need to worry about node provisioning, cluster setup, Hadoop configuration, or cluster tuning. Amazon EMR takes care of these tasks so you can focus on analysis.
- **Low Cost** – Amazon EMR pricing is simple and predictable: You pay an hourly rate for every instance hour you use. Because Amazon EMR has native support for Amazon EC2 Spot and Reserved Instances, you can also save 50-80% on the cost of the underlying instances.
- **Elastic** – With Amazon EMR, you can provision one, hundreds, or thousands of compute instances to process data at any scale. You can easily increase or decrease the number of instances and you only pay for what you use.
- **Reliable** – You can spend less time tuning and monitoring your cluster. Amazon EMR has tuned Hadoop for the cloud; it also monitors your cluster, retrying failed tasks and automatically replacing poorly performing instances.

- **Secure** – Amazon EMR automatically configures Amazon EC2 firewall settings that control network access to instances, and you can launch clusters in an Amazon VPC, a logically isolated network you define. For objects stored in Amazon S3, you can use Amazon S3 [server-side encryption](#) or Amazon S3 [client-side encryption](#) with Amazon EMR File System (EMRFS), using AWS KMS or customer-managed keys.
- **Flexible** – You have complete control over your cluster. You have root access to every instance, you can easily install additional applications, and you can customise every cluster. Amazon EMR also supports multiple Hadoop distributions and applications.

12.2 Backup/Restore and Disaster Recovery

With Amazon EMR, you can leverage multiple data stores, including Amazon S3, Amazon RDS, Amazon Glacier, and Amazon Redshift. Refer to those services for details on their backup/restore and disaster recovery.

12.3 Service Constraints

Please see <https://aws.amazon.com/documentation/elastic-mapreduce/> for more information.

12.4 Technical Requirements

Please refer to the links below and <https://aws.amazon.com/documentation/elastic-mapreduce/> and the following links for comprehensive technical documentation regarding Amazon EMR.

- **Management Guide** – Provides information about the Amazon EMR service and platform. [HTML](#) | [PDF](#)
- **Release Guide** – Provides information about Amazon EMR releases, including software installed on your Amazon EMR cluster like Hadoop and Spark. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for Amazon EMR in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

13.0 Service Definition – AWS CloudHSM

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

13.1 Service Overview

The AWS CloudHSM service helps you meet corporate, contractual and regulatory compliance requirements for data security by using dedicated HSM appliances within

the AWS Cloud. With CloudHSM, you control the encryption keys and cryptographic operations performed by the HSM.

The AWS CloudHSM service allows you to protect your encryption keys within HSMs designed and validated to government standards for secure key management. You can securely generate, store, and manage the cryptographic keys used for data encryption such that they are accessible only by you. AWS CloudHSM helps you comply with strict key management requirements without sacrificing application performance.

The AWS CloudHSM service works with Amazon VPC. CloudHSM instances are provisioned inside your Amazon VPC with an IP address that you specify, providing simple and private network connectivity to your Amazon EC2 instances. Placing AWS CloudHSM instances near your Amazon EC2 instances decreases network latency, which can improve application performance. AWS provides dedicated and exclusive (single tenant) access to AWS CloudHSM instances, isolated from other AWS customers.

Top benefits include:

- **Secure Key Storage** – As part of the service, you have dedicated access to HSM capabilities in the cloud. AWS CloudHSM protects your cryptographic keys with tamper-resistant HSM appliances that are designed to comply with international (Common Criteria EAL4+) and U.S. Government (National Institute of Standards and Technology [NIST] Federal Information Processing Standards [FIPS] 140-2) regulatory standards for cryptographic modules. You retain full control of your keys and cryptographic operations on the HSM, while Amazon manages and maintains the hardware without having access to your keys.
- **Contractual and Regulatory Compliance** – By protecting your keys in hardware and preventing them from being accessed by third parties, AWS CloudHSM can help you comply with the most stringent regulatory and contractual requirements for key protection.
- **On-Demand** – The AWS CloudHSM API, CLI Tools, and SDK let you start and stop dedicated AWS CloudHSM instances whenever you want.
- **Secure Connectivity** – AWS CloudHSM instances are in your VPC, so it is easy to use them with your Amazon EC2 applications. You use standard Amazon VPC security mechanisms to control access to AWS CloudHSM instances.
- **Improved Application Performance** – By placing AWS CloudHSM instances in your Amazon VPC near your Amazon EC2 instances, you can reduce network latency and increase the performance of your AWS applications that use HSMs.
- **Integrated** – You can use AWS CloudHSM with Amazon Redshift, Amazon RDS Oracle, or third-party applications such as SafeNet ProtectV volume encryption for EBS, Apache (SSL termination), or Microsoft SQL Server (transparent data encryption). You can also use AWS CloudHSM when writing your own applications and continue to use the standard cryptographic libraries you're familiar with, including PKCS#11, Java JCA/JCE, and Microsoft CAPI and CNG.

- **Auditable** – If you need to track resource changes, or audit activities for security and compliance purposes, you can review all of the AWS CloudHSM API calls made from your account through AWS CloudTrail. Additionally, you can audit operations on the HSM appliance using syslog or send syslog log messages to your own collector.

13.2 Backup/Restore and Disaster Recovery

AWS CloudHSM is available in multiple regions and Availability Zones to help you build highly available applications that require strong key protection. The AWS CloudHSM CLI Tools can help you configure high availability groups that span multiple Availability Zones, so you can build resilient applications. In the unlikely event of a hardware failure, you can launch a new AWS CloudHSM instance and replicate the keys to the new HSM with a few commands. You can also use AWS CloudHSM with your compatible on-premises HSMs to securely store keys in your datacentre. This increases key durability and gives you the flexibility to securely migrate keys in and out of AWS.

13.3 Service Constraints

Please see <https://aws.amazon.com/cloudhsm/> for more information.

13.4 Technical Requirements

Please refer to the links below and <https://aws.amazon.com/documentation/cloudhsm/> and the following links for comprehensive technical documentation regarding AWS CloudHSM.

- **User Guide** – Provides key concepts of AWS CloudHSM and documents advanced product features and the AWS CLI. [HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes the API operations for AWS CloudHSM. [HTML](#) | [PDF](#)

14.0 Service Definition – AWS CloudTrail

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

14.1 Service Overview

AWS CloudTrail is a web service that records AWS API calls for your account and delivers log files to you. The recorded information includes the identity of the API caller,

the time of the API call, the source IP address of the API caller, the request parameters, and the response elements returned by the AWS service.

With AWS CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS Cloud services (such as AWS CloudFormation). The AWS API call history produced by AWS CloudTrail enables security analysis, resource change tracking, and compliance auditing.

Top features and benefits include:

- **Increased Visibility** – AWS CloudTrail provides increased visibility into your user activity by recording AWS API calls. You can answer questions such as, what actions did a given user take over a given time period? For a given resource, which user has taken actions on it over a given time period? What is the source IP address of a given activity? Which activities failed due to inadequate permissions?
- **Durable and Inexpensive Log File Storage** – AWS CloudTrail uses Amazon S3 for log file storage and delivery, so log files are stored durably and inexpensively. You can use Amazon S3 lifecycle configuration rules to further reduce storage costs. For example, you can define rules to automatically delete old log files or archive them to Amazon Glacier for additional savings.
- **Easy Administration** – AWS CloudTrail is a fully managed service; you simply turn on AWS CloudTrail for your account using the AWS Management Console, the AWS CLI, or the AWS CloudTrail SDK and start receiving AWS CloudTrail log files in the Amazon S3 bucket that you specify.
- **Notifications for Log File Delivery** – AWS CloudTrail can be configured to publish a notification for each log file delivered, thus enabling you to automatically take action upon log file delivery. CloudTrail uses the Amazon SNS for notifications.
- **Log File Aggregation** – AWS CloudTrail can be configured to aggregate log files across multiple accounts and regions so that log files are delivered to a single bucket. For detailed instructions, refer to the Aggregating CloudTrail Log Files to a Single Amazon S3 Bucket section of the user guide.
- **Reliable and Timely Delivery** – AWS CloudTrail continuously transports events from AWS Cloud services using a highly available and fault tolerant processing pipeline. CloudTrail typically delivers events within 15 minutes of the API call.
- **Troubleshoot operational or security issues** – You can troubleshoot operational issues or perform security analysis by looking up API activity that was captured for your AWS account. Using the AWS CloudTrail console, AWS CLI, or AWS SDKs, you can quickly and easily answer questions related to API activity for the last 7 days and take immediate action.
- **Receive SNS Notifications of API activity** – AWS CloudTrail can be configured to be deliver API activity to an Amazon CloudWatch Logs log group you specify. You can then create CloudWatch Alarms to receive Amazon SNS notifications when specific API activity occurs.

14.2 Backup/Restore and Disaster Recovery

AWS CloudTrail uses Amazon S3 for log file storage and delivery, so log files are stored durably and inexpensively. Refer to Amazon S3 for more detailed information.

14.3 Service Constraints

Please see <https://aws.amazon.com/cloudtrail/> for more information.

14.4 Technical Requirements

Please refer to the links below and <https://aws.amazon.com/documentation/cloudtrail/> and the following links for comprehensive technical documentation regarding AWS CloudTrail.

- **User Guide** – Provides detailed descriptions of product concepts and includes instructions for using the various features with both the console and the CLI. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for AWS CloudTrail in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

15.0 Service Definition – AWS Config

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

15.1 Service Overview

AWS Config is a fully managed service that provides you with an AWS resource inventory, configuration history, and configuration change notifications to enable security and governance. With AWS Config you can discover existing AWS resources, export a complete inventory of your AWS resources with all configuration details, and determine how a resource was configured at any point in time. These capabilities enable compliance auditing, security analysis, resource change tracking, and troubleshooting.

Top benefits include:

- **Configuration Visibility** – You can view continuously updated details of all configuration attributes associated with AWS resources. You are notified via Amazon SNS of every configuration change and you can process these notifications programmatically.
- **Fully Managed** – With AWS Config, there are no software agents to install, and no databases to manage. AWS Config automatically manages the underlying infrastructure required to record, store and report on the configuration details of your AWS resources.
- **Easy to get started** – You can enable AWS Config with a few clicks in the AWS Management Console. AWS Config will discover your AWS resources and start recording configuration changes. You can access information about the configuration of any resource using the AWS Management Console, CLI, or SDKs.

- **Low cost** – With AWS Config, there are no upfront costs. You pay-as-you-go based on the number of resources and configuration changes recorded for your AWS account.

15.2 Backup/Restore and Disaster Recovery

AWS Config uses an Amazon S3 bucket you specify to store the information recorded. Refer to Amazon S3 for more detailed information.

15.3 Service Constraints

Please see <https://aws.amazon.com/config/> for more information.

15.4 Technical Requirements

AWS Config provides a detailed view of the resources associated with your AWS account, including how they are configured, how they are related to one another, and how the configurations, and their relationships have changed over time.

- **Developer Guide** – Provides a conceptual overview of AWS Config and includes detailed development instructions for using the various features. [HTML](#) | [PDF](#)
- **CLI Reference** – Documents the AWS Config CLI. [HTML](#)
- **API Reference** – Describes all the API operations for AWS Config in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

16.0 Service Definition – Amazon MQ

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

16.1 Service Overview

Amazon MQ is a managed message broker service for Apache ActiveMQ that makes it easy to set up and operate message brokers in the cloud. Message brokers allow different software systems—often using different programming languages, and on different platforms—to communicate and exchange information. Messaging is the communications backbone that connects and integrates the components of distributed applications, such as order processing, inventory management, and order fulfillment for e-commerce. Amazon MQ manages the administration and maintenance of ActiveMQ, a popular open-source message broker. The underlying infrastructure is automatically

provisioned for high availability and message durability to support the reliability of your applications. With Amazon MQ, you get direct access to the ActiveMQ console and industry standard APIs and protocols for messaging, including JMS, NMS, AMQP, STOMP, MQTT, and WebSocket. You can easily move from any message broker that uses these standards to Amazon MQ because you don't have to rewrite any messaging code in your applications.

Top features include:

1. **Fully Managed** - Amazon MQ makes it easy to setup and operate message brokers in the cloud. You can use the AWS Management Console, CLI, or API calls to launch a production-ready message broker in minutes. There is no need to provision hardware, and no need to install and maintain ActiveMQ software. Amazon MQ manages the set up and ongoing administrative tasks such as software upgrades, security updates, and failure detection and recovery. It's integrated with Amazon CloudWatch so you can monitor metrics and generate alarms to get alerts about potential issues. For example, you can monitor the depth of a queue, or generate alarms if messages are not getting through.
 - **Easy Migration** - Amazon MQ makes it easy to migrate messaging to the cloud while preserving the existing connections between your applications. It supports industry-standard APIs and protocols for messaging, including JMS, NMS, AMQP, STOMP, MQTT, and WebSocket. This enables you to move from any message broker that uses these standards to Amazon MQ without having to rewrite any messaging code in your applications. In most cases, you can simply update the endpoints of your Amazon MQ broker to connect to your existing applications, and start sending messages.
 - **Available and Durable** - Amazon MQ provides high availability and message durability. It runs on the same highly reliable infrastructure used by other Amazon Web Services. Amazon MQ always stores messages redundantly across multiple Availability Zones (AZs). Active/standby brokers are designed for high availability. In the event of a failure of the broker, or even a full AZ outage, Amazon MQ automatically fails over to the standby instance, so you can continue sending and receiving messages.
 - **Low Cost** - With Amazon MQ, you pay only for what you use. There are no minimum fees or upfront commitments. You pay for the number of hours your broker instance runs, and the storage you use monthly. Amazon MQ is free to try. The AWS Free Tier includes up to 750 hours of a single-instance mq.t2.micro broker and up to 1GB of storage per month for one year. There is no infrastructure to maintain, so you can repurpose or retire on-premises servers, storage, and load balancers associated with your message broker.

16.2 Backup/Restore and Disaster Recovery

Amazon MQ runs on the same highly reliable infrastructure used by other Amazon Web Services. Amazon MQ stores your messages redundantly across multiple Availability Zones (AZs). Active/standby brokers are designed for high availability. In the event of a

failure of the broker, or even a full AZ outage, Amazon MQ automatically fails over to the standby broker so you can continue sending and receiving messages. For more information, please see <https://docs.aws.amazon.com/amazon-mq/latest/developer-guide/welcome.html>.

16.3 Service Constraints

Please see <https://aws.amazon.com/amazon-mq/> for more information.

16.4 Technical Requirements

Amazon MQ is a managed message broker service for Apache ActiveMQ that makes it easy to set up and operate message brokers in the cloud. Amazon MQ provides interoperability with your existing applications and services. Amazon MQ works with your existing applications and services without the need to manage, operate, or maintain your own messaging system.

- **Developer Guide** - Provides a conceptual overview of Amazon MQ and includes detailed instructions for creating and managing brokers and migrating from on-premises brokers.
[HTML](#) | [PDF](#) | [Kindle](#)
- **REST API Reference** - Describes all REST APIs for Amazon MQ and provides HTTP methods, example request and response schemas, and a full list of parameters.
[HTML](#) | [PDF](#)

17.0 Service Definition – Amazon Simple Notification Service (Amazon SNS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

17.1 Service Overview

Amazon SNS is a fast, flexible, fully managed push notification service that lets you send individual messages or to fan-out messages to large numbers of recipients. Amazon SNS makes it simple and cost effective to send push notifications to mobile device users, email recipients or even send messages to other distributed services.

With Amazon SNS, you can send notifications to Apple, Google, Fire OS, and Windows devices, as well as to Android devices in China with Baidu Cloud Push. You can use Amazon SNS to send SMS messages to mobile device users in the US or to email recipients worldwide.

Beyond these endpoints, Amazon SNS can also deliver messages to Amazon SQS, AWS Lambda functions, or to any HTTP endpoint.

Top features include:

- Amazon SNS lets you push messages to mobile devices or distributed services via API or an easy-to-use management console. You can seamlessly scale from a handful of messages per day to millions of messages or higher.
- With Amazon SNS you can publish a message once and deliver it one or more times. So you can choose to direct unique messages to individual Apple, Google, or Amazon devices or broadcast deliveries to many mobile devices with a single publish request.
- Amazon SNS allows you to group multiple recipients using topics. A topic is an “access point” for allowing recipients to dynamically subscribe for identical copies of the same notification. One topic can support deliveries to multiple endpoint types—for example, you can group together iOS, Android and SMS recipients. When you publish once to a topic, Amazon SNS delivers appropriately formatted copies of your message to each subscriber.
- Amazon SNS has no upfront costs and you can pay as you go.

17.2 Backup/Restore and Disaster Recovery

Amazon SNS runs within Amazon’s proven network infrastructure and datacentres, so topics will be available whenever applications need them. To prevent messages from being lost, all messages published to Amazon SNS are stored redundantly across multiple servers and data centres.

17.3 Service Constraints

Please see <https://aws.amazon.com/sns/> for more information.

17.4 Technical Requirements

Amazon SNS is a web service that enables applications, end-users, and devices to instantly send and receive notifications from the cloud.

- **Developer Guide** – Provides a conceptual overview of Amazon SNS and includes detailed development instructions for using the various features.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for Amazon SNS in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

- **Quick Reference** – Briefly covers the essential commands for using Amazon SNS from the command line.
[PDF](#)

18.0 Service Definition – AWS Key Management Service (AWS KMS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

18.1 Service Overview

AWS KMS is a managed service that makes it easy for you to create and control the encryption keys used to encrypt your data, and uses HSMS to protect the security of your keys. AWS Key Management Service is integrated with other AWS Cloud services including Amazon EBS, Amazon S3, Amazon RDS, and Amazon Redshift. AWS KMS is also integrated with AWS CloudTrail to provide you with logs of all key usage to help meet your regulatory and compliance needs.

Top benefits include:

- **Centralised Key Management** – AWS Key Management Service provides you with centralised control of your encryption keys. AWS KMS presents a single view into all of the key usage in your organisation. You can easily create keys, implement key rotation, create usage policies, and enable logging from the AWS Management Console, or by using the API.
- **Integrated with AWS Cloud Services** – AWS KMS is integrated with Amazon S3, Amazon EBS, Amazon Redshift, Amazon RDS, and Amazon EMR to make it easy to encrypt the data you store with these services using keys that you manage.
- **Encryption for all your applications** – AWS KMS makes it easy to manage encryption keys used to encrypt data stored by your applications regardless of where you store it. AWS KMS provides an SDK for programmatic integration of encryption and key management into your applications.
- **Built-in Auditing** – AWS KMS works with AWS CloudTrail to provide you with logs of API calls made to or by AWS KMS. These logs help you meet compliance and regulatory requirements by providing details of when keys were accessed and who accessed them.
- **Fully Managed** – AWS KMS is a fully managed service, so you can focus on the encryption needs of your applications while AWS handles availability, physical security, and hardware maintenance of the underlying infrastructure.
- **Low-cost** – There is no charge for the storage of default keys in your account. You pay only for additional master keys that you create and your key usage.
- **Secure** – AWS KMS provides you a secure location to store and use encryption keys, using hardened systems where your unencrypted keys are only used in memory. AWS KMS keys are never transmitted outside of the AWS Regions in which they were created.

18.2 Backup/Restore and Disaster Recovery

AWS KMS is a managed service. As your usage of AWS KMS encryption keys grows, you do not have to buy additional key management hardware or software or manage any infrastructure. AWS KMS automatically scales to meet your encryption key needs.

Key storage is highly durable. AWS KMS stores multiple copies of encrypted versions of your keys in systems that are designed for 99.999999999% durability to help assure you that your keys will be available when you need to access them.

AWS KMS is deployed in multiple Availability Zones within an AWS Region to provide high availability for your encryption keys.

18.3 Service Constraints

Please see <https://aws.amazon.com/kms/> for more information.

18.4 Technical Requirements

AWS KMS is an encryption and key management service scaled for the cloud. KMS keys and functionality are used by other AWS Cloud services, and you can use them to protect data in your own applications that use AWS.

- **Developer Guide** – Provides conceptual overviews of AWS KMS and explains how to use it to protect data in your own applications that use AWS.
[HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for AWS KMS in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

19.0 Service Definition – Amazon CloudSearch

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

19.1 Service Overview

Amazon CloudSearch is a managed service in the AWS Cloud that makes it simple and cost-effective to set up, manage, and scale a search solution for your website or application. Amazon CloudSearch supports 34 languages and popular search features such as highlighting, autocomplete, and geospatial search.

With Amazon CloudSearch, you can quickly add rich search capabilities to your website or application. You don't need to become a search expert or worry about hardware provisioning, setup, and maintenance. With a few clicks in the AWS Management Console, you can create a search domain and upload the data that you want to make searchable, and Amazon CloudSearch will automatically provision the required resources and deploy a highly tuned search index.

You can easily change your search parameters, fine tune search relevance, and apply new settings at any time. As your volume of data and traffic fluctuates, Amazon CloudSearch seamlessly scales to meet your needs.

Amazon CloudSearch supports powerful search features such as:

- Free text, Boolean, and Faceted search
- Autocomplete suggestions
- Customizable relevance ranking and query-time rank expressions
- Field weighting
- Geospatial search
- Highlighting
- Support for 34 languages

Top benefits include:

- **Simple** – You can configure and manage an Amazon CloudSearch domain through the AWS Management Console, AWS CLI, and AWS SDKs. Simply point to a sample of your data and Amazon CloudSearch will automatically recommend how to configure your domain's indexing options. You can easily add or delete index fields and customise search options such as faceting and highlighting. Configuration changes do not require you to re-upload your data.
- **Scalable** – Amazon CloudSearch offers powerful auto scaling for all search domains. As your data or query volume changes, Amazon CloudSearch can scale your search domain's resources up or down as needed. You can control scaling if you know that you need more capacity for bulk uploads or are expecting a surge in search traffic.
- **Reliable** – Amazon CloudSearch provides automatic monitoring and recovery for your search domains. When Multi-AZ is enabled, Amazon CloudSearch provisions and maintains resources for a search domain in two Availability Zones to ensure high availability. Updates are automatically applied to the search instances in both Availability Zones. Search traffic is distributed across both Availability Zones and the instances in either zone are capable of handling the full load in the event of a failure.
- **High Performance** – Amazon CloudSearch ensures low latency and high throughput performance, even at large scale through automatic sharding and horizontal and vertical auto scaling.
- **Fully Managed** – Amazon CloudSearch is a fully managed custom search service. Hardware and software provisioning, setup and configuration, software patching, data partitioning, node monitoring, scaling, and data durability are handled for you.
- **Cost-Effective** – Amazon CloudSearch is designed to be cost-effective. You pay low hourly rates, and only for the resources you use. Amazon CloudSearch offers low total cost of ownership for your search applications compared to operating a search environment on your own.

- **Secure** – Amazon CloudSearch uses strong cryptographic methods to authenticate users and prevent unauthorised access to your domains. Amazon CloudSearch supports HTTPS and integrates with Identity and Access Management (IAM) to control access to the Amazon CloudSearch configuration service and each domain's document, search, and suggest services.

19.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon CloudSearch. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

19.3 Service Constraints

Please see <https://aws.amazon.com/cloudsearch/> for more information.

19.4 Technical Requirements

Amazon CloudSearch is a fully managed service in the cloud that makes it easy to set up, manage, and scale a search solution for your website. Amazon CloudSearch enables you to search large collections of data such as web pages, document files, forum posts, or product information. With Amazon CloudSearch, you can quickly add search capabilities to your website without having to become a search expert or worry about hardware provisioning, setup, and maintenance. As your volume of data and traffic fluctuates, Amazon CloudSearch automatically scales to meet your needs.

- **Developer Guide** – Describes how to create, manage, and use an Amazon CloudSearch domain to implement a search solution for your website or application. Also provides instructions for using the various features through the AWS Management Console, AWS CLI, and API.
[HTML](#) | [PDF](#) | [Kindle](#)

20.0 Service Definition – Amazon Redshift

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

20.1 Service Overview

Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyse all your data using your existing business intelligence tools. It is optimised for datasets ranging from a few hundred

gigabytes to a petabyte or more and costs a tenth the cost of most traditional data warehousing solutions.

Amazon Redshift delivers fast query performance by using columnar storage technology to improve I/O efficiency and parallelizing queries across multiple nodes. Amazon Redshift has custom JDBC and ODBC drivers that you can download from the Connect Client tab of our Console, allowing you to use a wide range of familiar SQL clients. You can also use standard PostgreSQL JDBC and ODBC drivers. Data load speed scales linearly with cluster size, with integrations to Amazon S3, Amazon DynamoDB, Amazon EMR, Amazon Kinesis, or any SSH-enabled host.

Amazon Redshift's data warehouse architecture allows you to automate most of the common administrative tasks associated with provisioning, configuring and monitoring a cloud data warehouse. Backups to Amazon S3 are continuous, incremental, and automatic. Restores are fast—you can start querying in minutes while your data is spooled down in the background. Enabling disaster recovery across regions takes just a few clicks.

Top features and benefits include:

- **Optimised for Data Warehousing** – Amazon Redshift uses a variety of innovations to obtain very high query performance on datasets ranging in size from a hundred gigabytes to a petabyte or more. It uses columnar storage, data compression, and zone maps to reduce the amount of I/O needed to perform queries. Amazon Redshift has a Massively Parallel Processing (MPP) data warehouse architecture, parallelizing and distributing SQL operations to take advantage of all available resources. The underlying hardware is designed for high performance data processing, using local attached storage to maximise throughput between the CPUs and drives, and a 10GigE mesh network to maximise throughput between nodes.
- **Scalable** – With a few clicks of the AWS Management Console or a simple API call, you can easily change the number or type of nodes in your cloud data warehouse as your performance or capacity needs change. Dense Storage (DS) nodes allow you to create very large data warehouses using HDDs for a very low price point. Dense Compute (DC) nodes allow you to create very high performance data warehouses using fast CPUs, large amounts of RAM and SSDs.
- **No Up-Front Costs** – You pay only for the resources you provision. You can choose On-Demand pricing with no up-front costs or long-term commitments, or obtain significantly discounted rates with Reserved Instance pricing.
- **Get Started in Minutes** – With a few clicks in the AWS Management Console or simple API calls, you can create a cluster, specifying its size, underlying node type, and security profile. Amazon Redshift will provision your nodes, configure the connections between them, and secure the cluster. Your data warehouse should be up and running in minutes.
- **Fully Managed** – Amazon Redshift handles all of the work needed to manage, monitor, and scale your data warehouse, from monitoring cluster health and taking backups to applying

patches and upgrades. You can easily resize your cluster as your performance and capacity needs change. By handling all these time-consuming, labour-intensive tasks, Amazon Redshift frees you up to focus on your data and business.

- **Fault Tolerant** – Amazon Redshift has multiple features that enhance the reliability of your data warehouse cluster. All data written to a node in your cluster is automatically replicated to other nodes within the cluster and all data is continuously backed up to Amazon S3. Amazon Redshift continuously monitors the health of the cluster and automatically re-replicates data from failed drives and replaces nodes as necessary.
- **Encryption** – With just a couple of parameter settings, you can set up Amazon Redshift to use SSL to secure data in transit and hardware-accelerated AES-256 encryption for data at rest. If you choose to enable encryption of data at rest, all data written to disk will be encrypted as well as any backups. By default, Amazon Redshift takes care of key management but you can choose to manage your keys using your own HSMs, AWS CloudHSM, or AWS KMS.
- **Network Isolation** – Amazon Redshift enables you to configure firewall rules to control network access to your data warehouse cluster. You can run Amazon Redshift inside Amazon VPC to isolate your data warehouse cluster in your own virtual network and connect it to your existing IT infrastructure using industry-standard encrypted IPsec VPN.
- **Audit and Compliance** – Amazon Redshift integrates with AWS CloudTrail to enable you to audit all Amazon Redshift API calls. Amazon Redshift also logs all SQL operations, including connection attempts, queries and changes to your database. You can access these logs using SQL queries against system tables or choose to have them downloaded to a secure location on Amazon S3. Amazon Redshift is compliant with Service Organization Controls (SOC) 1, SOC 2, SOC 3, and Payment Card Industry Data Security Standard (PCI DSS) Level 1 requirements. For more details, please visit the [AWS Compliance Center](#).
- **SQL** – Amazon Redshift is a SQL data warehouse solution and uses industry standard ODBC and JDBC connections. You can download our custom JDBC and ODBC drivers from the Connect Client tab of the service console. Many popular software vendors have certified Amazon Redshift with their offerings to enable you to continue to use the tools you do today.
- **Integrated** – Amazon Redshift is integrated with other AWS Cloud services and has built in commands to load data in parallel to each node from Amazon S3 or your Amazon EC2 and on premise servers using SSH.

20.2 Backup/Restore and Disaster Recovery

Amazon Redshift's automated snapshot feature continuously backs up new data on the cluster to Amazon S3. Snapshots are continuous, incremental and automatic. Amazon Redshift stores your snapshots for a user-defined period, which can be from one to thirty-five days. You can take your own snapshots at any time, which leverage all existing system snapshots and are retained until you explicitly delete them. Redshift can also asynchronously replicate your snapshots to Amazon S3 in another region for disaster recovery. Once you delete a cluster, your system snapshots are removed but your user snapshots are available until you explicitly delete them.

You can use any system or user snapshot to restore your cluster using the AWS Management Console or the Amazon Redshift APIs. Your cluster is available as soon

as the system metadata has been restored and you can start running queries while user data is spooled down in the background.

20.3 Service Constraints

Please see <https://aws.amazon.com/redshift/> for more information.

20.4 Technical Requirements

Amazon Redshift is a fast, fully managed, petabyte-scale data warehouse service that makes it simple and cost-effective to efficiently analyse all your data using your existing business intelligence tools. It is optimised for datasets ranging from a few hundred gigabytes to a petabyte or more.

- **Getting Started** – Introduces you to Amazon Redshift, helps you set up an account, and walks you through a simple example to use Amazon Redshift for the first time. Also provides tips and links to advanced product features and resources.
[HTML](#) | [PDF](#) | [Kindle](#)
- **Cluster Management Guide** – Shows you how to create and manage Amazon Redshift clusters.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for Amazon Redshift in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)
- **Database Developer Guide** – Explains how to design, build, query, and maintain the databases that make up your data warehouse, and includes syntax for Amazon Redshift SQL commands and functions.
[HTML](#) | [PDF](#)
- **CLI Guide** – Describes the CLI for Amazon Redshift in detail. Provides basic syntax, options, and usage examples for each command.
[HTML](#)

21.0 Service Definition – Amazon EC2 Container Registry (Amazon ECR)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

21.1 Service Overview

Amazon EC2 Container Registry (Amazon ECR) is a fully managed [Docker](#) container registry that makes it easy for developers to store, manage, and deploy Docker container images. Amazon ECR is integrated with [Amazon ECS](#), simplifying your development-to-production workflow. Amazon ECR eliminates the need to operate your

own container repositories or worry about scaling the underlying infrastructure. Amazon ECR hosts your images in a highly available and scalable architecture, allowing you to reliably deploy containers for your applications. Integration with AWS IAM provides resource-level control of each repository. With Amazon ECR, there are no up-front fees or commitments. You pay only for the amount of data you store in your repositories and data transferred to the Internet.

Top benefits include:

- **Fully Managed** – Amazon ECR eliminates the need to operate and scale the infrastructure required to power your container registry. There is no software to install and manage or infrastructure to scale. Just push your container images to Amazon ECR and pull the images when you need to deploy.
- **Secure** – Amazon ECR transfers your container images over HTTPS and automatically encrypts your images at rest. You can configure policies to manage permissions and control access to your images using IAM users and roles without having to manage credentials directly on your Amazon EC2 instances.
- **Highly Available** – Amazon ECR has a highly scalable, redundant, and durable architecture. Your container images are highly available and accessible, allowing you to reliably deploy new containers for your applications.
- **Simplified Workflow** – Amazon ECR integrates with Amazon ECS and the Docker CLI, allowing you to simplify your development and production workflows. You can easily push your container images to Amazon ECR using the Docker CLI from your development machine, and Amazon ECS can pull them directly for production deployments.

21.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon ECR. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

21.3 Service Constraints

Please see <https://aws.amazon.com/ecr/> for more information.

21.4 Technical Requirements

Amazon ECR is a fully managed Docker container registry that makes it easy for developers to store, manage, and deploy Docker container images.

- **User Guide** – Describes key concepts of Amazon ECR and provides instructions for using the features of Amazon ECR.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for Amazon ECR in detail.
[HTML](#) | [PDF](#)

- **ECR section of AWS CLI Reference** – Documents the Amazon ECR commands available in the AWS CLI.
[HTML](#)

22.0 Service Definition – Amazon EC2 Container Service (Amazon ECS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

22.1 Service Overview

Amazon ECS is a highly scalable, high-performance container management service that supports [Docker](#) containers and allows you to easily run applications on a managed cluster of [Amazon EC2](#) instances. Amazon ECS eliminates the need for you to install, operate, and scale your own cluster management infrastructure. With simple API calls, you can launch and stop Docker-enabled applications, query the complete state of your cluster, and access many familiar features like security groups, [Elastic Load Balancing](#), [Amazon EBS](#) volumes, and [AWS IAM](#) roles. You can use Amazon ECS to schedule the placement of containers across your cluster based on your resource needs and availability requirements. You can also integrate your own scheduler or third-party schedulers to meet business or application specific requirements.

Top benefits include:

- **Easily Manage Clusters for Any Scale** – Amazon ECS eliminates the need to operate cluster management software or design fault-tolerant cluster architectures. There is no software to install, manage, or scale, allowing you to focus on building Dockerised applications.
- **Flexible Container Placement** – You can use Amazon ECS to schedule long-running applications and batch jobs. You can also use the APIs to get up-to-date cluster state information to integrate your own custom, as well as third-party, schedulers. Amazon ECS is a shared-state, optimistic concurrency system that supports multiple schedulers on the same cluster for each business or application-specific requirement.
- **Designed for Use with Other AWS Cloud Services** – Amazon ECS is integrated with Elastic Load Balancing, Amazon EBS, Amazon VPC, IAM, and AWS CloudTrail, providing you a complete solution for running a wide range of containerised applications or services.
- **Extensible** – Amazon ECS can easily be integrated or extended through simple APIs. Amazon ECS provides complete visibility and control into your AWS resources, so you can easily integrate and use your own schedulers or connect Amazon ECS into your existing software delivery process (e.g., continuous integration and delivery systems).
- **Performance at Scale** – Amazon ECS is built on technology developed from many years of experience running highly scalable services. You can launch tens or tens of thousands of Docker containers in seconds using Amazon ECS with no additional complexity.
- **Secure** – Amazon ECS launches your containers on your own Amazon EC2 instances. No compute resources are shared with other customers. Your clusters run in an Amazon VPC, allowing you to use your own Amazon VPC security groups and network Access Control Lists

(ACLs). These features provide you a high level of isolation and help you use Amazon ECS to build highly secure and reliable applications.

22.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon ECS. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

22.3 Service Constraints

Please see <https://aws.amazon.com/ecs/> for more information.

22.4 Technical Requirements

Amazon ECS is a highly scalable, fast, container management service that makes it easy to run, stop, and manage Docker containers on a cluster of Amazon EC2 instances.

- **Developer Guide** – Describes key concepts of Amazon ECS and provides instructions for using the features of Amazon ECS.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Documents the Amazon ECS Query API.
[HTML](#) | [PDF](#)
- **ECS section of AWS CLI Reference** – Documents the Amazon ECS commands available in the AWS CLI.
[HTML](#)

23.0 Service Definition – Amazon EKS (in preview)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

23.1 Service Overview

Amazon Elastic Container Service for Kubernetes (Amazon EKS) is a managed service that makes it easy for customers to run Kubernetes on AWS without needing to install and operate their own Kubernetes clusters. With Amazon EKS, upgrades and high availability are managed by AWS. Amazon EKS runs three Kubernetes masters across three Availability Zones in order to ensure high availability. Amazon EKS automatically detects and replaces unhealthy masters, and it provides automated version upgrades and patching for the masters. Top benefits include:

- **No Masters to Manage** - Amazon EKS runs master nodes across multiple AWS Availability Zones, automatically detects and replaces unhealthy masters, and provides on-demand upgrades and patching. You simply provision worker nodes and connect them to the AWS cloud.
- **Secure by Default** - Secure and encrypted communication channels are automatically setup between your worker nodes and the EKS-managed masters, making your infrastructure running on Amazon EKS secure by default.
- **Conformant and Compatible** - We run the latest upstream version of Kubernetes and are Certified Kubernetes Conformant, so applications running on Amazon EKS are fully compatible with applications running on any standard Kubernetes environment.
- **Built with the Community** - We actively work with the community, making contributions to the Kubernetes code base that help Amazon EKS users take advantage of AWS services and features.

23.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon EKS. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

23.3 Service Constraints

Amazon EKS features and default service limits are found at <https://aws.amazon.com/eks/faq/>. Note that Amazon EKS is currently available in preview.

23.4 Technical Requirements

Not available this time.

24.0 Service Definition – AWS Fargate

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

24.1 Service Overview

AWS Fargate is a technology for Amazon ECS and EKS that allows you to run containers without having to manage servers or clusters. With AWS Fargate, you no longer have to provision, configure, and scale clusters of virtual machines to run containers. This removes the need to choose server types, decide when to scale your clusters, or optimize cluster packing. AWS Fargate removes the need for you to interact with or think about servers or clusters. Fargate lets you focus on designing and building

your applications instead of managing the infrastructure that runs them. AWS Fargate will also support Amazon EKS in 2018.

Amazon ECS and EKS have two modes: Fargate launch type and EC2 launch type. With Fargate launch type, all you have to do is package your application in containers, specify the CPU and memory requirements, define networking and IAM policies, and launch the application. EC2 launch type allows you to have server-level, more granular control over the infrastructure that runs your container applications. With EC2 launch type, you can use Amazon ECS and EKS to manage a cluster of servers and schedule placement of containers on the servers. Amazon ECS and EKS keeps track of all the CPU, memory and other resources in your cluster, and also finds the best server for a container to run on based on your specified resource requirements. You are responsible for provisioning, patching, and scaling clusters of servers. You can decide which type of server to use, which applications and how many containers to run in a cluster to optimize utilization, and when you should add or remove servers from a cluster. EC2 launch type gives you more control of your server clusters and provides a broader range of customization options, which might be required to support some specific applications or possible compliance and government requirements.

Top benefits include:

- **No Clusters to Manage** - With AWS Fargate, you only have to think about the containers so you can just focus on building and operating your application. AWS Fargate eliminates the need to manage a cluster of Amazon EC2 instances. You no longer have to pick the instance types, manage cluster scheduling, or optimize cluster utilization. All of this goes away with Fargate.
- **Seamless Scaling** - AWS Fargate makes it easy to scale your applications. You no longer have to worry about provisioning enough compute resources for your container applications. After you define your application requirements (e.g., CPU, memory, etc.), AWS Fargate manages all the scaling and infrastructure needed to run your containers in a highly-available manner. You no longer have to decide when to scale your clusters or pack them for optimal utilization. With Fargate, you can launch tens or tens of thousands of containers in seconds and easily scale to run your most mission-critical applications.
- **Integrated with Amazon ECS and EKS** - AWS Fargate seamlessly integrates with Amazon ECS. You just define your application as you do for Amazon ECS. You package your application into task definitions, specify the CPU and memory needed, define the networking and IAM policies that each container needs, and upload everything to Amazon ECS. After everything is setup, AWS Fargate launches and manages your containers for you.

24.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Fargate. For additional information beyond what is described herein, please refer to https://docs.aws.amazon.com/AmazonECS/latest/developerguide/AWS_Fargate.html

24.3 Service Constraints

AWS Fargate features and default service limits are found at <https://aws.amazon.com/fargate/faqs/>.

24.4 Technical Requirements

AWS Fargate is a technology that you can use with Amazon ECS to run [containers](#) without having to manage servers or clusters of EC2 instances. With AWS Fargate, you no longer have to provision, configure, and scale clusters of virtual machines to run containers. This removes the need to choose server types, decide when to scale your clusters, or optimize cluster packing.

When you run your tasks and services with the Fargate launch type, you package your application in containers, specify the CPU and memory requirements, define networking and IAM policies, and launch the application. For more information, please refer to https://docs.aws.amazon.com/AmazonECS/latest/developerguide/AWS_Fargate.html.

25.0 Service Definition – AWS Elastic Beanstalk

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

25.1 Service Overview

AWS Elastic Beanstalk is an easy-to-use service for deploying and scaling web applications and services developed with Java, [.NET](#), PHP, Node.js, Python, Ruby, Go, and [Docker](#) on familiar servers such as Apache, Nginx, Passenger, and [IIS](#).

You can simply upload your code and AWS Elastic Beanstalk automatically handles the deployment, from capacity provisioning, load balancing, and auto scaling to application

health monitoring. At the same time, you retain full control over the AWS resources powering your application and can access the underlying resources at any time.

Top benefits include:

- **Fast and Simple to Begin** – AWS Elastic Beanstalk is the fastest and simplest way to deploy your application on AWS. You simply use the AWS Management Console, a Git repository, or an Integrated Development Environment (IDE) such as Eclipse or Visual Studio to upload your application, and AWS Elastic Beanstalk automatically handles the deployment details of capacity provisioning, load balancing, auto scaling, and application health monitoring. Within minutes, your application will be ready to use without any infrastructure or resource configuration work on your part.
- **Developer Productivity** – AWS Elastic Beanstalk provisions and operates the infrastructure and manages the application stack (platform) for you, so you don't have to spend the time or develop the expertise. It will also keep the underlying platform running your application up to date with the latest patches and updates. Instead, you can focus on writing code rather than spending time managing and configuring servers, databases, load balancers, firewalls, and networks.
- **Impossible to Outgrow** – AWS Elastic Beanstalk automatically scales your application up and down based on your application's specific need using easily adjustable Auto Scaling settings. For example, you can use CPU utilization metrics to trigger Auto Scaling actions. With AWS Elastic Beanstalk, your application can handle peaks in workload or traffic while minimizing your costs.
- **Complete Resource Control** – You have the freedom to select the AWS resources, such as Amazon EC2 instance type, that are optimal for your application. Additionally, AWS Elastic Beanstalk lets you “open the hood” and retain full control over the AWS resources powering your application. If you decide you want to take over some (or all) of the elements of your infrastructure, you can do so seamlessly by using AWS Elastic Beanstalk's management capabilities.

25.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elastic Beanstalk. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

25.3 Service Constraints

Please see <https://aws.amazon.com/elasticbeanstalk/> for more information.

25.4 Technical Requirements

With AWS Elastic Beanstalk, you can quickly deploy and manage applications in the AWS Cloud without worrying about the infrastructure that runs those applications. AWS Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and AWS Elastic Beanstalk automatically handles

the details of capacity provisioning, load balancing, scaling, and application health monitoring.

- **Developer Guide** – Provides conceptual overviews of AWS Elastic Beanstalk as well as detailed instructions on the variety of tools for deploying applications and managing applications and environments.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS Elastic Beanstalk in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)
- **Getting Started Walkthrough** – Walks developers through the use of the console to create, view, deploy, and update an application for the first time, as well as the steps for editing and terminating an environment.
[HTML](#)

26.0 Service Definition – AWS Lambda

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

26.1 Service Overview

AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume—there is no charge when your code is not running. With AWS Lambda, you can run code for virtually any type of application or backend service, all with zero administration. Just upload your code and AWS Lambda takes care of everything required to run and scale your code with high availability. You can set up your code to automatically trigger from other AWS Cloud services or call it directly from any web or mobile application.

Top benefits include:

- **No Servers to Manage** – AWS Lambda automatically runs your code without requiring you to provision or manage servers. Just write the code and upload it to AWS Lambda.
- **Continuous Scaling** – AWS Lambda automatically scales your application by running code in response to each trigger. Your code runs in parallel and processes each trigger individually, scaling precisely with the size of the workload.
- **Subsecond Metering** – With AWS Lambda, you are charged for every 100ms your code executes and the number of times your code is triggered. You don't pay anything when your code isn't running.

26.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Lambda. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

26.3 Service Constraints

Please see <https://aws.amazon.com/lambda/> for more information.

26.4 Technical Requirements

AWS Lambda is a zero-administration compute platform for back-end web developers that runs your code for you in the AWS Cloud and provides you with a fine-grained pricing structure. AWS Lambda runs your back-end code on its own AWS compute fleet of Amazon EC2 instances across multiple Availability Zones in a region, which provides the high availability, security, performance, and scalability of the AWS infrastructure.

- **Developer Guide** – Provides a conceptual overview of AWS Lambda, includes detailed instructions for using the various features, and provides a complete API reference for developers. [HTML](#) | [PDF](#)

27.0 Service Definition – AWS Auto Scaling

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

27.1 Service Overview

AWS Auto Scaling is a web service designed to launch or terminate Amazon EC2 instances automatically based on user-defined policies, schedules, and health checks

Top benefits include:

- **Setup Scaling Quickly** - AWS Auto Scaling lets you set target utilization levels for multiple resources in a single, intuitive interface. You can quickly see the average utilization of all of your scalable resources without having to navigate to other consoles. For example, if your application uses Amazon EC2 and Amazon DynamoDB, you can use AWS Auto Scaling to manage resource provisioning for all of the EC2 Auto Scaling groups and database tables in your application.
- **Make Smart Scaling Decisions** - AWS Auto Scaling lets you build scaling plans that automate how groups of different resources respond to changes in demand. You can optimize availability, costs, or a balance of both. AWS Auto Scaling automatically creates all of the scaling policies and sets targets for you based on your preference. AWS Auto Scaling monitors your application and automatically adds or removes capacity from your resource groups in real-time as demands change.
- **Automatically Maintain Performance** - Using AWS Auto Scaling, you maintain optimal application performance and availability, even when workloads are periodic, unpredictable, or continuously changing. AWS Auto Scaling continually monitors your applications to make sure

that they are operating at your desired performance levels. When demand spikes, AWS Auto Scaling automatically increases the capacity of constrained resources so you maintain a high quality of service.

- **Pay Only For What You Need** - AWS Auto Scaling can help you optimize your utilization and cost efficiencies when consuming AWS services so you only pay for the resources you actually need. When demand drops, AWS Auto Scaling will automatically remove any excess resource capacity so you avoid overspending. AWS Auto Scaling is free to use, and allows you to optimize the costs of your AWS environment.

27.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Auto Scaling. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>

27.3 Service Constraints

AWS Auto Scaling features and default service limits are found at <https://aws.amazon.com/autoscaling/faqs/>

27.4 Technical Requirements

With AWS Auto Scaling, you can quickly discover the scalable AWS resources for your application and set up dynamic scaling. It uses [Amazon EC2 Auto Scaling](#) to scale your EC2 instances and Application Auto Scaling to scale resources from other services.

- **Administrator Guide**
Provides a conceptual overview of AWS Auto Scaling and includes detailed instructions for using the various features.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference**
Describes all the API operations for AWS Auto Scaling in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)
- **AWS Auto Scaling section of the AWS CLI Reference**
Describes the AWS CLI commands that you can use to administer AWS Auto Scaling. Provides syntax, options, and usage examples for each command.
[HTML](#)

28.0 Service Definition – Amazon API Gateway

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

28.1 Service Overview

Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. With a few clicks in the AWS Management Console, you can create an API that acts as a “front door” for applications to access data, business logic, or functionality from your back-end services, such as workloads running on Amazon EC2, code running on AWS Lambda, or any Web application. Amazon API Gateway handles all the tasks involved in accepting and processing up to hundreds of thousands of concurrent API calls, including traffic management, authorization and access control, monitoring, and API version management. Amazon API Gateway has no minimum fees or start-up costs. You pay only for the API calls you receive and the amount of data transferred out.

Top benefits include:

- **Low-Cost and Efficient** – With Amazon API Gateway, you pay only for calls made to your APIs and data transfer out. There are no minimum fees or up-front commitments.
- **Performance at Any Scale** – With Amazon CloudFront integration, Amazon API Gateway allows you to take advantage of the worldwide network of edge locations to provide your end users with the lowest possible latency for API requests and responses. Amazon API Gateway also helps you manage traffic through throttling, so that back-end operations can withstand traffic spikes. Additionally, Amazon API Gateway helps you improve the performance of your APIs by caching the output of API calls to avoid calling your backend systems unnecessarily.
- **Easily Monitor API Activity** – After your API is deployed, Amazon API Gateway provides you with a dashboard to visually monitor calls to your services using [Amazon CloudWatch](#), so you see performance metrics and information on API calls, data latency, and error rates.
- **Streamline API Development** – Amazon API Gateway lets you simultaneously run multiple versions of the same API, allowing you to quickly iterate, test, and release new versions.
- **Flexible Security Controls** – Amazon API Gateway provides you with tools to authorise access to your APIs and control service operation access. You can use AWS administration and security tools, such as AWS IAM and Amazon Cognito, to authorise access to your APIs. Amazon API Gateway can verify signed API calls on your behalf using the same technology AWS uses for its own APIs. If you already use [OAuth](#) tokens or other authorization mechanisms, Amazon API Gateway can use AWS Lambda to execute a custom authoriser to help you verify incoming requests.
- **Create RESTful Endpoints for Existing Services** – With Amazon API Gateway, you can create modern, resource-based APIs and then use the dynamic and flexible data transformation capabilities to generate the requests in the language your target services expect. Amazon API Gateway also helps you protect your existing services by setting throttling rules to avoid overwhelming your back-end infrastructure during unpredictable traffic spikes.
- **Run Your APIs Without Servers** – Amazon API Gateway tightly integrates with AWS Lambda to allow you to create completely server-less APIs. With Amazon API Gateway, you create REST APIs that your mobile and web applications can use to call publicly available AWS Cloud services through your code running in AWS Lambda. AWS Lambda runs your code on a high-availability

compute infrastructure and performs all the heavy lifting and administration of your compute resources.

28.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Lambda. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

28.3 Service Constraints

Please see <https://aws.amazon.com/api-gateway/> for more information.

28.4 Technical Requirements

Amazon API Gateway is a fully managed service that makes it easy for developers to publish, maintain, monitor, and secure APIs at any scale. Create an API to access data, business logic, or functionality from your back-end services, such as applications running on Amazon EC2, code running on AWS Lambda, or any web application.

- **Developer Guide** – Provides a conceptual overview of Amazon API Gateway and includes detailed instructions for using the service.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the Amazon API Gateway operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#)

29.0 Service Definition – Amazon Route 53

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

29.1 Service Overview

Amazon Route 53 is a highly available and scalable cloud DNS web service. It is designed to give developers and businesses an extremely reliable and cost-effective way to route end users to Internet applications by translating names like `www.example.com` into the numeric IP addresses like `192.0.2.1` that computers use to connect to each other.

Amazon Route 53 effectively connects user requests to infrastructure running in AWS—such as Amazon EC2 instances, Elastic Load Balancing load balancers, or Amazon S3 buckets—and can also be used to route users to infrastructure outside of AWS. You can

use Amazon Route 53 to configure DNS health checks to route traffic to healthy endpoints or to independently monitor the health of your application and its endpoints. Amazon Route 53 Traffic Flow makes it easy for you to manage traffic globally through a variety of routing types, including Latency Based Routing, Geo DNS, and Weighted Round Robin—all of which can be combined with DNS Failover in order to enable a variety of low-latency, fault-tolerant architectures. Using Amazon Route 53 Traffic Flow's simple visual editor, you can easily manage how your end users are routed to your application's endpoints—whether in a single AWS Region or distributed around the globe. Amazon Route 53 also offers Domain Name Registration—you can purchase and manage domain names such as example.com and Amazon Route 53 will automatically configure DNS settings for your domains.

Top benefits include:

- **Highly Available and Reliable** – Amazon Route 53 is built using AWS's highly available and reliable infrastructure. The distributed nature of our DNS servers helps ensure a consistent ability to route your end users to your application. Features such as Amazon Route 53 Traffic Flow help you improve reliability with easy configuration of failover to re-route your users to an alternate location if your primary application endpoint becomes unavailable. Amazon Route 53 is designed to provide the level of dependability required by important applications. Amazon Route 53 is backed by the [Amazon Route 53 Service Level Agreement](#).
- **Scalable** – Amazon Route 53 is designed to automatically scale to handle very large query volumes without any intervention from you.
- **Designed for Use with Other AWS Cloud Services** – Amazon Route 53 is designed to work well with other AWS features and offerings. You can use Amazon Route 53 to map domain names to your Amazon EC2 instances, Amazon S3 buckets, Amazon CloudFront distributions, and other AWS resources. By using the AWS Identity and Access Management (IAM) service with Amazon Route 53, you get fine-grained control over who can update your DNS data. You can use Amazon Route 53 to map your zone apex (example.com versus www.example.com) to your Elastic Load Balancing instance, Amazon CloudFront distribution, AWS Elastic Beanstalk environment, or Amazon S3 website bucket using a feature called Alias record.
- **Simple** – With self-service sign-up, Amazon Route 53 can start to answer your DNS queries within minutes. You can configure your DNS settings with the AWS Management Console or our easy-to-use API. You can also programmatically integrate the Amazon Route 53 API into your overall web application. For instance, you can use Amazon Route 53's API to create a new DNS record whenever you create a new Amazon EC2 instance. Amazon Route 53 Traffic Flow makes it easy to set up sophisticated routing logic for your applications by using the simple visual policy editor.
- **Fast** – Using a global anycast network of DNS servers around the world, Amazon Route 53 is designed to automatically route your users to the optimal location depending on network conditions. As a result, the service offers low query latency for your end users, as well as low update latency for your DNS record management needs. Amazon Route 53 Traffic Flow lets you further improve your customers' experience by running your application in multiple locations

around the world and using traffic policies to ensure your end users are routed to the closest healthy endpoint for your application.

- **Cost Effective** – Amazon Route 53 passes on the benefits of AWS’s scale to you. You pay only for the resources you use, such as the number of queries that the service answers for each of your domains, hosted zones for managing domains through the service, and optional features such as traffic policies and health checks, all at a low cost and without minimum usage commitments or any up-front fees.
- **Secure** – By integrating Amazon Route 53 with AWS IAM, you can grant unique credentials and manage permissions for every user within your AWS account and specify who has access to which parts of the Amazon Route 53 service.
- **Flexible** – Amazon Route 53 Traffic Flow routes traffic based on multiple criteria, such as endpoint health, geographic location, and latency. You can configure multiple traffic policies and decide which policies are active at any given time. You can create and edit traffic policies using the simple visual editor in the Amazon Route 53 console, AWS Software Development Kits (SDKs), or the Amazon Route 53 API. Traffic Flow’s versioning feature maintains a history of changes to your traffic policies, so you can easily roll back to a previous version using the console or API.

29.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Route 53. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

29.3 Service Constraints

Please see <https://aws.amazon.com/route53/> for more information.

29.4 Technical Requirements

Amazon Route 53 is a highly available and scalable DNS web service.

- **Developer Guide** – Provides an overview of Amazon Route 53, detailed feature descriptions, procedures for using the console, and an explanation of how to use the API.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for Amazon Route 53 in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

30.0 Service Definition – AWS Storage Gateway

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

30.1 Service Overview

The AWS Storage Gateway is a service connecting an on-premises software appliance with [cloud-based storage](#) to provide seamless and secure integration between an organisation's on-premises IT environment and AWS's storage infrastructure. The service allows you to securely [store data in the AWS Cloud](#) for scalable and cost-effective storage. The AWS Storage Gateway supports industry-standard storage protocols that work with your existing applications. It provides low-latency performance by maintaining frequently accessed data on-premises while securely storing all of your data encrypted in [Amazon S3](#) or [Amazon Glacier](#).

Top benefits include:

- **Secure** – The AWS Storage Gateway securely transfers your data to AWS over SSL and securely stores your data in Amazon S3 and Amazon Glacier using one of the strongest block ciphers available, 256-bit Advanced Encryption Standard (AES-256) encryption. A multi-factor scheme is used where each block of data is encrypted with a unique key, and the key itself is encrypted with a regularly rotated master key.
- **Durably Backed by Amazon S3 and Amazon Glacier** – The AWS Storage Gateway durably stores your on-premises application data by uploading it to Amazon S3 and Amazon Glacier. Amazon S3 and Amazon Glacier redundantly store data in multiple facilities and on multiple devices within each facility. Amazon S3 and Amazon Glacier also perform regular, systematic data integrity checks and are built to be automatically self-healing.
- **Compatible** – There is no need to re-architect your on-premises applications. Gateway-cached volumes and gateway-stored volumes expose a standard iSCSI block disk device interface, and gateway-VTL presents a standard iSCSI virtual tape library interface.
- **Cost Effective** – By making it easy for your on-premises applications to store data on Amazon S3 or Amazon Glacier, AWS Storage Gateway reduces the cost, maintenance, and scaling challenges associated with managing primary, backup, and archive storage environments. You pay only for what you use with no long-term commitments.
- **Optimised for Network Efficiency** – The AWS Storage Gateway efficiently uses your Internet bandwidth to speed up the upload of your on-premises application data to AWS. The AWS Storage Gateway compresses and only uploads data that has changed, minimizing the amount of data sent over the Internet. You can also use [AWS Direct Connect](#) to further increase throughput and reduce your network costs by establishing a dedicated network connection between your on-premises gateway and AWS.

30.2 Backup/Restore and Disaster Recovery

Gateway-stored volumes and gateway-cached volumes are designed to seamlessly integrate with [Amazon S3](#), [Amazon EBS](#), and [Amazon EC2](#) by enabling you to store point-in-time snapshots of your on-premises application data in Amazon S3 as [Amazon EBS snapshots](#) for future recovery on-premises or in Amazon EC2. This integration allows you to easily mirror data from your on-premises applications to applications running on Amazon EC2 for disaster recovery and on-demand compute capacity cases. Gateway-VTL integrates with [Amazon Glacier](#) and allows you to cost-effectively and durably store your archive and long-term backup data. All three gateway configurations integrate with [Amazon CloudWatch](#) to provide metrics, monitoring, and alarms; with [AWS CloudTrail](#) to provide a record of API calls for resource change tracking and compliance auditing; and implement tags to simplify management of your AWS resources for [cost allocation](#) or [resource grouping](#). For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

30.3 Service Constraints

Please see <https://aws.amazon.com/storagegateway/> for more information.

30.4 Technical Requirements

AWS Storage Gateway is a service that connects an on-premises software appliance with cloud-based storage to provide seamless and secure integration between your on-premises IT environment and the AWS storage infrastructure in the cloud.

- **User Guide** – Describes AWS Storage Gateway concepts and provides instructions on using the various features with both the console and the API.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes the AWS Storage Gateway API and CLI operations. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

31.0 Service Definition – Amazon CloudFront

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

31.1 Service Overview

Amazon CloudFront is a global Content Delivery Network (CDN) service that accelerates delivery of your websites, APIs, video content, or other web assets. It integrates with other AWS products to give developers and businesses an easy way to accelerate content to end users with no minimum usage commitments.

Top benefits include:

- **Fast** – Using a network of edge locations around the world, Amazon CloudFront caches copies of your static content close to viewers, lowering latency when they download your objects and giving you the high, sustained data transfer rates needed to deliver large popular objects to end users at scale. Requests for your dynamic content are carried back to your origin servers running in AWS (e.g., [Amazon EC2](#), [Elastic Load Balancing](#)) over optimised network paths for a more reliable and consistent experience. These network paths are constantly monitored by Amazon and connections from Amazon CloudFront edge locations to the origin are reused to serve your dynamic content from our CDN with the best possible performance.
- **Simple** – A single API call lets you get started distributing content from your Amazon S3 bucket or Amazon EC2 instance or other origin server through the Amazon CloudFront network. Alternatively, interact with Amazon CloudFront through the AWS Management Console's simple graphical user interface. There is no need to create separate domains for your static and dynamic content. With Amazon CloudFront, you can just use the same domain name to point to all of your website content. Any changes you make to your existing configuration take effect across the entire global network within minutes. Plus, since there's no need to negotiate with a sales person, you can get started quickly and begin delivering your entire website using Amazon CloudFront.
- **Designed for Use with Other AWS Cloud Services** – Amazon CloudFront is designed for use with other AWS Cloud services, including Amazon S3, where you can durably store the definitive versions of your static files, and Amazon EC2, where you can run your application server for dynamically generated content. If you are using Amazon S3 or Amazon EC2 as an origin server, data transferred from the origin server to edge locations (Amazon CloudFront "origin fetches") will be billed at a lower price than Internet data transfer out of Amazon S3 or Amazon EC2. Amazon CloudFront also integrates with Elastic Load Balancing. For instance, you can deploy your web application on Amazon EC2 servers behind Elastic Load Balancing and use Amazon CloudFront to deliver your entire website.
- **Cost Effective** – Amazon CloudFront passes on the benefits of Amazon's scale to you. You pay only for the content that you deliver through the network, without minimum commitments or up-front fees. This applies for any type of content that you deliver—static, dynamic, streaming media, or a web application with any combination of these.
- **Elastic** – With Amazon CloudFront, you don't need to worry about maintaining expensive web server capacity to meet the demand from potential traffic spikes for your content. The service automatically responds as demand increases or decreases without any intervention from you. Amazon CloudFront also uses multiple layers of caching at each edge location and collapses simultaneous requests for the same object before contacting your origin server. These

optimizations further help reduce the need to scale your origin infrastructure as your website becomes more popular.

- **Reliable** – Amazon CloudFront is built using Amazon’s highly reliable infrastructure. The distributed nature of edge locations used by Amazon CloudFront automatically routes end users to the closest available location as required by network conditions. Origin requests from the edge locations to AWS origin servers (e.g., Amazon EC2, Amazon S3) are carried over network paths that Amazon constantly monitors and optimises for both availability and performance.
- **Global** – Amazon CloudFront uses a global network of edge locations, located near your end users in the United States, Europe, Asia, South America, and Australia.

31.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon CloudFront. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

31.3 Service Constraints

Please see <https://aws.amazon.com/cloudfront/> for more information.

31.4 Technical Requirements

Amazon CloudFront is a web service that speeds up distribution of your static and dynamic web content, for example, .html, .css, .php, image, and media files, to end users. CloudFront delivers your content through a worldwide network of edge locations. When an end user requests content that you’re serving with CloudFront, the user is routed to the edge location that provides the lowest latency, so content is delivered with the best possible performance. If the content is already in that edge location, CloudFront delivers it immediately. If the content is not currently in that edge location, CloudFront retrieves it from an Amazon S3 bucket or an HTTP server (for example, a web server) that you have identified as the source for the definitive version of your content.

- **Developer Guide** – Provides an overview of Amazon CloudFront, detailed feature descriptions, procedures for using the console, and an explanation of the API.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for Amazon CloudFront in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

32.0 Service Definition – AWS Web Application Firewall (AWS WAF)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

32.1 Service Overview

AWS WAF is a web application firewall that helps protect your web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. AWS WAF gives you control over which traffic to allow or block to your web applications by defining customizable web security rules. You can use AWS WAF to create custom rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that are designed for your specific application. New rules can be deployed within minutes, letting you respond quickly to changing traffic patterns. In addition, AWS WAF includes a full-featured API that you can use to automate the creation, deployment, and maintenance of web security rules.

Top benefits include:

- **Increased Protection Against Web Attacks** – AWS WAF protects web applications from attacks by filtering traffic based on rules that you create. For example, you can filter web requests based on IP addresses, HTTP headers, HTTP body, or Uniform Resource Identifier (URI) strings, which allows you to block common attack patterns, such as SQL injection or cross-site scripting.
- **Security Integrated with How You Develop Applications** – Every feature in AWS WAF can be configured using either the AWS WAF API or the AWS Management Console. This allows you to define application-specific rules that increase web security as you develop your application. This lets you put web security at multiple points in the development chain, from the hands of the developer initially writing code, to the DevOps engineer deploying software, to the security experts conducting an audit.
- **Ease of Deployment and Maintenance** – AWS WAF is easy to deploy, protects any application deployed on Amazon CloudFront content delivery service, and there is no additional software to deploy. You can centrally define your rules and reuse them across all the web applications that you need to protect.
- **Improved Web Traffic Visibility** – You can set up AWS WAF to just monitor requests that match your filter criteria. AWS WAF gives near real-time visibility into your web traffic, which you can use to create new rules or alerts in Amazon CloudWatch.
- **Cost-Effective Web Application Protection** – With AWS WAF you pay only for what you use. AWS WAF provides a customizable, self-service offering, and pricing is based on how many rules you deploy and how many web requests your web application receives. There are no minimum fees and no up-front commitments.

32.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS WAF. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

32.3 Service Constraints

Please see <http://aws.amazon.com/waf/> for more information.

32.4 Technical Requirements

AWS WAF is a web application firewall service that lets you monitor web requests for Amazon CloudFront distributions and restrict access to your content. Use AWS WAF to block or allow requests based on conditions that you specify, such as the IP addresses that requests originate from or values in the requests.

- **Developer Guide** – Describes how to get started with AWS WAF, explains key concepts, and provides step-by-step instructions that show you how to use the features.

[HTML](#) | [PDF](#)

- **API Reference** – Describes all the API operations for AWS WAF in detail.

[HTML](#) | [PDF](#)

33.0 Service Definition – AWS Artifact

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

33.1 Service Overview

An audit and compliance portal for on-demand access to download AWS' compliance reports and manage select agreements.

AWS Artifact provides on-demand access to AWS' security and compliance reports and select online agreements. Reports available in AWS Artifact include our Service Organization Control (SOC) reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals that validate the implementation and operating effectiveness of AWS security controls.

Agreements available in AWS Artifact include the Business Associate Addendum (BAA) and the Nondisclosure Agreement (NDA).

Top benefits include:

- **Generate Compliance Reports with Ease and On-Demand** - Save time with on-demand access to AWS' auditor issued reports, certifications, accreditations and other third-party attestations
- **Globally Available** Save time with on-demand access to AWS' auditor issued reports, certifications, accreditations and other third-party attestations
- **Easy Identification** - Validate the implementation and operating effectiveness of the AWS security control environment from any geography or vertical.
- **Quick Assessments** - Easily identify the scope of each of the audit artifacts, including services, regions and applicable audit dates
- **Continuous monitoring** - Continuously monitor the security and compliance of AWS with immediate access when new reports are released
- **Enhanced Transparency** - Perform due-diligence anytime with enhanced transparency into the AWS control environment.

33.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Artifact. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

33.3 Service Constraints

Please see <https://aws.amazon.com/artifact/> for more information.

33.4 Technical Requirements

AWS Artifact is a web service that enables you to download AWS security and compliance documents such as ISO certifications and SOC reports.

- **User Guide** - Describes key concepts of AWS Artifact and provides instructions for using the features of AWS Artifact.
[HTML](#) | [PDF](#)

34.0 Service Definition – AWS Mobile Hub

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

34.1 Service Overview

AWS Mobile Hub lets customers easily add and configure features for mobile apps, including user authentication, data storage, back-end logic, push notifications, content

delivery, and analytics. AWS Mobile Hub gives customers easy access to testing on real devices, as well as analytics dashboards to track app usage, all from a single, integrated console. AWS Mobile Hub includes integration with Amazon Lex, so customers can build mobile apps that use speech and text to a mobile app. Amazon Lex lets customers use the technology that powers Amazon Alexa to create engaging speech- and text-based conversational bots in their own app. AWS Mobile Hub's SaaS Connectors make it easy to securely access data in third-party enterprise SaaS applications from a mobile apps.

Top benefits include:

- **Start for free** - Get started building your app for free. Every AWS service has a free tier. Pay only for the services you use beyond the free tier.
- **Add cloud services fast** - Create amazing cloud-enabled apps in minutes. Use serverless technology to run code without provisioning or managing servers.
- **Deliver quality apps** - Automate your DevOps pipeline with build, test, and deploy services for your iOS, Android, and web applications.
- **Engage your audience** - Analyze app usage, and engage in meaningful conversations with your users using email, two-way SMS, and mobile push.

34.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Mobile. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

34.3 Service Constraints

Please see <http://aws.amazon.com/waf/> for more information.

34.4 Technical Requirements

AWS Mobile Hub provides an integrated console that helps you build, test, and monitor your mobile apps. Use the console to choose the features you want to include in your app. Mobile Hub then provisions and configures the necessary AWS services on your behalf and creates a working sample app for you.

- **Developer Guide** – Walks you through how to get started and use the AWS Mobile Hub console. [HTML](#) | [PDF](#)

35.0 Service Definition – AWS Import/ Export Disk

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

35.1 Service Overview

AWS Import/Export Disk accelerates moving large amounts of data into and out of the AWS Cloud using portable storage devices for transport. AWS Import/Export Disk transfers your data directly onto and off of storage devices using Amazon's high-speed internal network and bypassing the Internet. For significant data sets, AWS Import/Export Disk is often faster than Internet transfer and more cost effective than upgrading your connectivity.

AWS Import/Export Disk supports data transfer into and out of [Amazon S3](#) buckets in the US East (Northern Virginia), US West (Oregon), US West (Northern California), EU (Ireland), and Asia Pacific (Singapore) regions. The service also supports importing data into [Amazon EBS](#) snapshots in the same regions. In addition, AWS Import/Export supports importing data into [Amazon Glacier](#) in the US East (Northern Virginia), US West (Oregon), US West (Northern California), and EU (Ireland) regions.

35.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Import/Export Disk. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

35.3 Service Constraints

Please see <http://aws.amazon.com/importexport/disk/> for more information.

35.4 Technical Requirements

AWS Import/Export is a service that accelerates transferring large amounts of data into and out of AWS using physical storage appliances, bypassing the Internet. AWS Import/Export consists of AWS Snowball, which uses on-demand, Amazon-provided secure storage appliances to physically transport terabytes to many petabytes of data,

and AWS Import/Export Disk, which uses customer-provided portable devices to transfer smaller datasets.

- **Developer Guide** – Provides an overview of AWS Import/Export, including detailed instructions for using Snowball to transfer large amounts of data into AWS.

[HTML](#) | [PDF](#) | [Kindle](#)

36.0 Service Definition – Amazon Elastic File System (Amazon EFS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

36.1 Service Overview

Amazon EFS is a file storage service for Amazon EC2 instances. Amazon EFS is easy to use and provides a simple interface that allows you to create and configure file systems quickly and easily. With Amazon EFS, storage capacity is elastic, growing and shrinking automatically as you add and remove files, so your applications have the storage they need, when they need it.

Amazon EFS supports the Network File System version 4 (NFSv4) protocol, so the applications and tools that you use today work seamlessly with Amazon EFS. Multiple Amazon EC2 instances can access an Amazon EFS file system at the same time, providing a common data source for workloads and applications running on more than one instance.

With Amazon EFS, you pay only for the storage used by your file system. You don't need to provision storage in advance and there is no minimum fee or setup cost. Amazon EFS is designed for a wide variety of use cases like content repositories, development environments, and home directories. With on-demand scaling and performance, Amazon EFS is an ideal solution for big data applications.

Top benefits include:

- **Seamless Integration** – Amazon EFS supports the NFSv4 protocol, so the applications and tools that you use today work seamlessly with Amazon EFS. Amazon EFS file systems support

standard file system semantics, and [Amazon EC2](#) instances mount Amazon EFS file systems using standard file system mount tools.

- **Scale Up and Down Seamlessly** – Amazon EFS automatically scales your file system storage capacity up or down as you add or remove files without disrupting your applications, ensuring you always have the storage you need while reducing time-consuming administration tasks.
- **Fully Managed Service** – Amazon EFS has a simple interface that allows you to create and configure file systems quickly and easily. The service manages all the file storage infrastructure for you, avoiding the complexity of deploying, patching, and maintaining complex file system deployments.
- **Share File Storage Across Instances** – Multiple Amazon EC2 instances can access an Amazon EFS file system, so applications that scale beyond a single instance can access a file system. Amazon EC2 instances running in multiple Availability Zones within the same region can access the file system, so that many users can access and share a common data source.
- **Consistent, Scalable Performance** – Amazon EFS is SSD-based and is designed to provide the throughput, IOPS, and low latency needed for a broad range of workloads. With Amazon EFS, throughput and IOPS scale as a file system grows, and file operations are delivered with consistent, low latencies.
- **Low Cost** – Amazon EFS provides the capacity you need, when you need it, without having to provision storage in advance. You pay for what you use, with no minimum commitments or up-front fees.
- **Highly Available and Durable** – Amazon EFS is designed to be durable and highly available. Each Amazon EFS file system object (i.e., directory, file, and link) is redundantly stored across multiple Availability Zones.
- **Secure** – Amazon EFS allows you to tightly control access to your file systems. Only Amazon EC2 instances within the [Amazon VPC](#) you specify can directly access your Amazon EFS file systems. Amazon VPC security groups and network access control lists allow you to manage network access to your Amazon EFS file systems. At the file and directory level, Amazon EFS supports user and group read/write/execute permissions. The service is also integrated with [AWS IAM](#), which can be used to control access to Amazon EFS APIs as well as manage resource-level permissions.

36.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon EFS. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

36.3 Service Constraints

Please see <https://aws.amazon.com/efs/> for more information.

36.4 Technical Requirements

Amazon EFS provides file storage for your Amazon EC2 instances. With Amazon EFS, you can create a file system, mount the file system on your Amazon EC2 instances, and then read and write data from your Amazon EC2 instances to and from your file system.

- **User Guide** – Walks through how to set up Amazon EFS and integrate it with other services. Includes the API reference.
[HTML](#) | [PDF](#)

37.0 Service Definition – Amazon Aurora

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

37.1 Service Overview

Amazon Aurora is a MySQL-compatible relational database engine that combines the speed and availability of high-end commercial databases with the simplicity and cost effectiveness of open-source databases. Amazon Aurora provides up to five times better performance than MySQL with the security, availability, and reliability of a commercial database at one tenth the cost.

Top benefits include:

- **High Performance** – Amazon Aurora provides five times the throughput of standard MySQL running on the same hardware. This consistent performance is on par with commercial databases, at 1/10th the cost and no license constraints. On the largest Amazon Aurora instance you can achieve up to 500,000 reads and 100,000 writes per second. You can further scale read operations using read replicas that have a very low 10 ms latency.
- **Highly Secure** – Amazon Aurora provides multiple levels of security for your database. These include network isolation using [Amazon VPC](#), encryption at rest using keys you create and control through [AWS KMS](#), and encryption of data in transit using SSL. On an encrypted Amazon Aurora instance, data in the underlying storage is encrypted, as are the automated backups, snapshots, and replicas in the same cluster.
- **MySQL-Compatible** – The Amazon Aurora database engine is fully compatible with MySQL 5.6 using the InnoDB storage engine. This means the code, applications, drivers, and tools you already use today with your MySQL databases can be used with Amazon Aurora with little or no change. This also allows for easy migration of existing MySQL databases using standard MySQL import and export tools or MySQL binlog replication.
- **Highly Scalable** – You can scale your Amazon Aurora database from an instance with 2 vCPU and 15 GiB of memory up to an instance with 32 vCPUs and 244 GiB memory. You can also add up to 15 low-latency read replicas across three Availability Zones to further scale read capacity. Amazon Aurora automatically grows storage as needed, from 10 GB up to 64 TB.
- **High Availability and Durability** – Amazon Aurora is designed to offer greater than 99.99% availability. Recovery from physical storage failures is transparent, and instance failover typically requires less than 30 seconds. Amazon Aurora's storage is fault tolerant and self-healing. Six copies of your data is replicated across three Availability Zones and continuously backed up to Amazon S3.
- **Fully Managed** – Amazon Aurora is a fully managed database service. You no longer need to worry about database management tasks such as hardware provisioning, software patching,

setup, configuration, monitoring, or backups. Amazon Aurora automatically and continuously monitors and backs up your database to Amazon S3, enabling granular, point-in-time recovery.

37.2 Backup/Restore and Disaster Recovery

Amazon Aurora is designed to offer greater than 99.99% availability, increasing MySQL and PostgreSQL performance and availability by tightly integrating the database engine with an SSD-backed virtualized storage layer purpose-built for database workloads. Amazon Aurora's storage is fault-tolerant and self-healing, and disk failures are repaired in the background without loss of database availability. Amazon Aurora is designed to automatically detect database crashes and restart without the need for crash recovery or to rebuild the database cache. If the entire instance fails, Amazon Aurora will automatically fail over to one of up to 15 read replicas. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

37.3 Service Constraints

Please see <http://aws.amazon.com/rds/aurora/> for more information.

37.4 Technical Requirements

If you are new to Amazon RDS, you must first complete the tasks in [Setting Up for Amazon RDS](#) and then review [Getting Started with Amazon RDS](#).

Once you are ready to create an Amazon Aurora DB Instance, visit [Creating a DB Cluster and Connecting to a Database on an Amazon Aurora DB Instance](#). You can use the [AWS Management Console](#) or Amazon RDS APIs to launch a DB Instance in minutes. If you wish to run more than 40 DB Instances, please complete this [form](#) and we will promptly respond to your request.

38.0 Service Definition – AWS Database Migration Service

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

38.1 Service Overview

AWS Database Migration Service helps you migrate databases to AWS easily and securely. The source database remains fully operational during the migration, minimizing downtime to applications that rely on the database. AWS Database Migration Service can migrate your data to and from most widely used commercial and open-source databases. The service supports homogenous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle to Amazon Aurora or Microsoft SQL Server to MySQL. It also allows you to stream data to Amazon Redshift from any of the supported sources including Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle, and SQL Server, enabling consolidation and easy analysis of data in the petabyte-scale data warehouse.

Top benefits include:

- **Simple to Use** – AWS Database Migration Service is simple to use. There is no need to install any drivers or applications, and it does not require changes to the source database in most cases. You can begin a database migration with just a few clicks in the AWS Management Console. Once the migration has started, AWS manages all the complexities of the migration process including automatically replicating data changes that occur in the source database during the migration process.
- **Zero Downtime** – AWS Database Migration Service helps you migrate your databases to AWS with virtually no downtime. All data changes to the source database that occur during the migration are continuously replicated to the target, allowing the source database to be fully operational during the migration process. After the database migration is complete, the target database will remain synchronised with the source for as long as you choose, allowing you to switch over the database at a convenient time.
- **Supports Most Widely Used Databases** – AWS Database Migration Service can migrate your data to and from most of the widely used commercial and open-source databases. It supports homogeneous migrations such as Oracle to Oracle, as well as heterogeneous migrations between different database platforms, such as Oracle to Amazon Aurora. Migrations can be from on-premises databases to Amazon RDS or Amazon EC2, databases running on Amazon EC2 to Amazon RDS, or vice versa, as well as from one Amazon RDS database to another Amazon RDS database.
- **Low Cost** – AWS Database Migration Service is a low-cost service. You only pay for the compute resources used during the migration process and any additional log storage. Migrating a terabyte-sized database can be done for as little as a few dollars per month (see Supplier Pricing Document for pricing details). This applies to both homogeneous and heterogeneous migrations of any supported databases. This is in stark contrast to conventional database migration methods, which can be very expensive.
- **Fast and Easy to Set Up** – You can set up a migration task within minutes in the AWS Management Console. A migration task is where you define the parameters AWS Database Migration Service uses to execute the migration. This includes setting up connections to the source and target databases, as well as choosing the replication instance used to run the

migration process. Once set up, the same task can be used for test runs before performing the actual migration.

- **Reliable** – AWS Database Migration Service is highly resilient and self-healing. It continually monitors source and target databases, network connectivity, and the replication instance. In case of interruption, it automatically restarts the process and continues the migration from where it was halted. Detailed diagnostic information is available for you to take the necessary corrective action for errors that cannot be automatically resolved.

38.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Database Migration Service. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

38.3 Service Constraints

Please see <https://aws.amazon.com/dms/> for more information.

38.4 Technical Requirements

AWS Database Migration Service is a web service you can use to migrate data from your database that is on-premises, on an Amazon RDS DB instance, or in a database on an Amazon EC2 instance to a database on an AWS Cloud service. These services can include a database on Amazon RDS or a database on an Amazon EC2 instance. You can also migrate a database from an AWS Cloud service to an on-premises database. You can migrate data between heterogeneous or homogenous database engines.

- **User Guide** – Describes all AWS DS concepts and provides instructions on using the various features with both the console and the AWS CLI. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for AWS Database Migration Service in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

39.0 Service Definition – AWS Migration Hub

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

39.1 Service Overview

AWS Migration Hub provides a single location to track the progress of application migrations across multiple AWS and partner solutions. Using Migration Hub allows you to choose the AWS and partner migration tools that best fit your needs, while providing visibility into the status of migrations across your portfolio of applications. Migration Hub also provides key metrics and progress for individual applications, regardless of which tools are being used to migrate them. For example, you might use [AWS Database Migration Service](#), [AWS Server Migration Service](#), and partner migration tools such as ATADATA ATAmotion or CloudEndure Live Migration to migrate an application comprised of a database, virtualized web servers, and a bare metal server. Using Migration Hub, you can view the migration progress of all the resources in the application. This allows you to quickly get progress updates across all of your migrations, easily identify and troubleshoot any issues, and reduce the overall time and effort spent on your migration projects.

AWS Migration Hub provides a single place to monitor migrations in any AWS region where your migration tools are available. There is no additional cost for using Migration Hub. You only pay for the cost of the individual migration tools you use, and any resources being consumed on AWS.

Top benefits include:

- **Centralized Tracking** - Migrations involve many components that need to be tracked such as the status of servers or databases being migrated, and these are typically tracked across different tools. AWS Migration Hub helps address this by providing a central location to track the status of all these components, making it easier to view overall migration progress and reducing the time spent determining current status and next steps.
- **Migration Flexibility** - AWS Migration Hub provides the flexibility to use the migration tools that work best for your organization. Whether you use AWS migration tools like AWS Server Migration Service, AWS Database Migration Service, or partner tools like ATADATA and CloudEndure, Migration Hub makes it easy for you to track migrations from all of those tools.
- **Improved Visibility** - AWS Migration Hub helps plan your migrations by allowing you to group related servers and resources that should be migrated together. Once grouped into a single

application, you can easily track the progress of these resources through their entire migration lifecycle to help ensure they all migrate successfully.

39.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Serverless Application Repository. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/documentation/>

39.3 Service Constraints

Please see <https://docs.aws.amazon.com/migrationhub/latest/ug/policy-templates.html> and <https://aws.amazon.com/documentation/> for more information on service constraints for AWS Migration Hub.

39.4 Technical Requirements

AWS Migration Hub provides a single location to track migration tasks across multiple AWS tools and partner solutions. With Migration Hub, you can choose the AWS and partner migration tools that best fit your needs while providing visibility into the status of your migration projects. Migration Hub also provides key metrics and progress information for individual applications, regardless of which tools are used to migrate them.

- **User Guide**
Walks through how to set up AWS Migration Hub and integrate it with other services and includes the API reference.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference**
Describes the API operations available for Migration Hub along with sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

40.0 Service Definition – AWS Server Migration Service (SMS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

40.1 Service Overview

AWS Server Migration Service (SMS) is an agentless service which makes it easier and faster for you to migrate thousands of on-premises workloads to AWS. AWS SMS allows you to automate, schedule, and track incremental replications of live server volumes, making it easier for you to coordinate large-scale server migrations.

Top benefits include:

- **Easy to Get Started** - Start and manage server migration with a few clicks via the AWS Management Console. AWS Server Migration Service will automatically replicate live server volumes to AWS and create Amazon Machine Images (AMI) as needed.
- **Control** - Create and manage a customized replication schedule designed for large-scale migrations, and track the progress of each migration.
- **Agility** - Perform migrations faster while minimizing network bandwidth, by migrating only incremental changes made to on-premises servers.
- **Cost-Effective** - AWS Server Migration Service is free to use; pay only for the storage resources used during the migration process.
- **Minimize Downtime** - Incremental server replication allows you to reduce server downtime significantly.

40.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Server Migration Service (SMS). For additional information beyond what is described herein, please refer to <https://aws.amazon.com/documentation/>

40.3 Service Constraints

Please see <https://aws.amazon.com/server-migration-service/faqs/> for more information on service constraints for AWS Server Migration Service (SMS).

40.4 Technical Requirements

AWS Server Migration Service (SMS) combines data collection tools with automated server replication to speed the migration of on-premises servers to AWS.

- **User Guide**
Describes key concepts of AWS SMS and provides instructions for using the features of AWS SMS.
[HTML](#) | [PDF](#)
- **SMS section of AWS CLI Reference**
Documents the AWS SMS commands available in the AWS CLI.

[HTML](#)

41.0 Service Definition – AWS Snowball

The following subsections provide service definition information.

41.1 Service Overview

AWS Snowball is a petabyte-scale data transport solution that uses secure appliances to [transfer large amounts of data](#) into and out of the [AWS cloud](#). Using Snowball addresses common challenges with large-scale data transfers including high network costs, long transfer times, and security concerns. Transferring data with Snowball is simple, fast, secure, and can be as little as one-fifth the cost of high-speed Internet.

With Snowball, you don't need to write any code or purchase any hardware to transfer your data. Simply create a job in the AWS Management Console and a Snowball appliance will be automatically shipped to you. Once it arrives, attach the appliance to your local network, download and run the Snowball client to establish a connection, and then use the client to select the file directories that you want to transfer to the appliance. The client will then encrypt and transfer the files to the appliance at high speed. Once the transfer is complete and the appliance is ready to be returned, the E Ink shipping label will automatically update and you can track the job status via [Amazon Simple Notification Service \(SNS\)](#), text messages, or directly in the Console.

Top benefits include:

- **High Speed** - Even with high-speed Internet connections, it can take months to transfer large amounts of data. For example, 100 terabytes of data will take more than 100 days to transfer over a dedicated 100 Mbps connection. That same transfer can be accomplished in less than one day, plus shipping time, using two Snowball appliances.
- **Extremely Scalable** - A single Snowball appliance can transport multiple terabytes of data and multiple appliances can be used in parallel to transfer petabytes of data into or out of AWS. [Learn more »](#)

- **Tamper Resistant and Secure** - Snowball appliances use tamper-resistant enclosures, 256-bit encryption, and an industry-standard Trusted Platform Module (TPM) designed to ensure both security and full chain-of-custody for your data. Encryption keys are managed with the [AWS Key Management Service \(KMS\)](#), and they are never sent to, or stored on the appliance.
- **Simple and Compatible** - Transfer jobs are created right from the AWS Management Console. Once a job is created, AWS automatically ships a Snowball appliance to you. When you receive the appliance, simply attach it to your local network, download and run the Snowball client to establish a connection, and then use the client to select the data that you want to transfer to the appliance. The client will then encrypt and transfer the files to the appliance at high speed. Once the transfer is complete and the appliance is ready to be returned, the E Ink shipping label will automatically update to the correct AWS facility, and the job status can be tracked via [Amazon SNS](#), text messages, or directly in the Console.
- **Low Cost** - It can cost thousands of dollars to transfer 100 terabytes of data using high-speed Internet. The same 100 terabytes of data can be transferred using two Snowball appliances for as little as one-fifth the cost of using the Internet.
- **Easy Data Retrieval** - Many organizations are concerned that once they have moved all their data to the cloud that it will be both expensive and time-consuming to retrieve the data if needed. Snowball offers all customers a fast and inexpensive way to ensure data can be quickly transferred both into and out of AWS.

41.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Snowball. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation> or

<https://aws.amazon.com/snowball/faqs/>.

41.3 Service Constraints

Please see <https://aws.amazon.com/documentation/snowball/> for more information.

41.4 Technical Requirements

AWS Snowball is a service for customers who want to transport terabytes or petabytes of data to and from AWS, or who want to access the storage and compute power of the AWS Cloud locally and cost effectively in places where connecting to the internet may not be an option.

- **AWS Snowball User Guide** - Provides a conceptual overview of how to use AWS Snowball with a Snowball appliance, including guidance for importing and exporting data into the AWS Cloud, shipping considerations, and other features of a Snowball appliance.
[HTML](#) | [PDF](#)
- **API Reference** - Describes all the API operations for AWS Snowball in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)
- **AWS Snowball Edge Developer Guide** - Provides a conceptual overview of how to use AWS Snowball with an Edge appliance, and includes guidance for local storage and compute, clustering, importing and exporting data into Amazon S3, and other features of an Edge appliance.
[HTML](#) | [PDF](#)

42.0 Service Definition – AWS Snowball Edge

The following subsections provide service definition information.

42.1 Service Overview

AWS Snowball Edge is a 100TB data transfer device with on-board storage and compute capabilities. You can use Snowball Edge to move large amounts of data into and out of AWS, as a temporary storage tier for large local datasets, or to support local workloads in remote or offline locations.

Snowball Edge connects to your existing applications and infrastructure using standard

storage interfaces, streamlining the data transfer process and minimizing setup and integration. Snowball Edge can cluster together to form a local storage tier and process your data on-premises, helping ensure your applications continue to run even when they are not able to access the cloud.

Top benefits include:

- **Process Data Locally** - You can deploy AWS Lambda code on Snowball Edge to perform tasks such as analyzing data streams or processing data locally. Data is collected, stored, and processed using the Snowball Edge independent from any other storage or compute resource.
- **Stand-Alone Storage** - You can use Snowball Edge as stand-alone storage. You can point existing applications to Snowball Edge devices as a file or object storage destination without changing application logic. You can also cluster Snowball Edge devices together into a single, larger, storage tier with increased durability. When a Snowball Edge is full, it can be removed from the cluster, returned to AWS, and replaced with a new Snowball Edge.
- **Secure** - Snowball Edge devices use tamper-resistant enclosures, 256-bit encryption, and industry-standard Trusted Platform Modules (TPM) designed to ensure both security and full chain-of-custody for your data. Encryption keys are managed with the [AWS Key Management Service \(KMS\)](#) and they are never stored on the device.
- **Scalable** - Snowball Edge devices can transport multiple terabytes of data and multiple devices can be used in parallel to transfer petabytes of data into or out of AWS. Snowball Edge is currently available in select regions and your location will be verified once you create a job in the AWS Management Console.

42.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Snowball Edge. For additional information beyond what is described herein, please refer to <https://docs.aws.amazon.com/snowball/latest/developer-guide/whatisedge.html>.

42.3 Technical Requirements

AWS Snowball Edge is a service for customers who want to transport terabytes or petabytes of data to and from AWS, or who want to access the storage and compute

power of the AWS Cloud locally and cost effectively in places where connecting to the internet may not be an option.

- **AWS Snowball Edge Developer Guide** - Provides a conceptual overview of how to use AWS Snowball with an Edge appliance, and includes guidance for local storage and compute, clustering, importing and exporting data into Amazon S3, and other features of an Edge appliance.

[HTML](#) | [PDF](#)

43.0 Service Definition – AWS Snowmobile

The following subsections provide service definition information.

43.1 Service Overview

AWS Snowmobile is an Exabyte-scale data transfer service used to move extremely large amounts of data to AWS. You can transfer up to 100PB per Snowmobile, a 45-foot long ruggedized shipping container, pulled by a semi-trailer truck. Snowmobile makes it easy to move massive volumes of data to the cloud, including video libraries, image repositories, or even a complete data center migration. Transferring data with Snowmobile is secure, fast and cost effective.

After an initial assessment, a Snowmobile will be transported to your data center and AWS personnel will configure it for you so it can be accessed as a network storage target. When your Snowmobile is on site, AWS personnel will work with your team to connect a removable, high-speed network switch from Snowmobile to your local network and you can begin your high-speed data transfer from any number of sources within your data center to the Snowmobile. After your data is loaded, Snowmobile is driven back to AWS where your data is imported into [Amazon S3](#) or [Amazon Glacier](#).

Snowmobile uses multiple layers of security designed to protect your data including dedicated security personnel, GPS tracking, alarm monitoring, 24/7 video surveillance, and an optional escort security vehicle while in transit. All data is encrypted with 256-bit encryption keys managed through the [AWS Key Management Service](#) (KMS) and designed to ensure both security and full chain-of-custody of your data.

Top benefits include:

- **Fast Transfer Even at Massive Scale** - Even with high-speed internet connections, it can take decades to transfer extremely large amounts of data. With Snowmobile, you can move 100 petabytes of data in as little as a few weeks, plus transport time. That same transfer could take more than 20 years to accomplish over a direct connect line with a 1Gbps connection.
- **Strong Encryption** - Your data is encrypted with keys that you provide before it is written to Snowmobile. All data is encrypted with 256-bit encryption and you can manage your encryption keys with AWS Key Management Service (AWS KMS). Encryption keys used by encrypting servers on the Snowmobile are never written to disk. Should power be removed for any reason, the keys are securely erased.
- **Rugged, Durable, and Secure** - Snowmobile is tamper-resistant, water-resistant and temperature controlled. The data container is only operated by AWS personnel and physical access is limited via security access hardware controls. Snowmobile is protected by 24/7 video surveillance and alarm monitoring, GPS tracking and may optionally be escorted by a security vehicle while in transit.
- **Customized for Your Needs** - As physical sites can have different migration requirements, AWS will work with you to help ensure all needs are factored in before Snowmobile arrives on site.
- **Massively Scalable** - One Snowmobile can transport up to one hundred petabytes of data in a single trip, the equivalent of using about 1,250 [AWS Snowball devices](#).

43.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Snowmobile. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/snowmobile/faqs/>.

43.3 Service Constraints

Please see <https://aws.amazon.com/snowmobile/faqs/> for more information.

43.4 Technical Requirements

A Snowmobile comes with a removable connector cabinet that needs to be mounted on one of your data center racks where it can be connected directly to your high-speed network backbone. The connector rack provides multiple 40Gb/s interfaces that can transfer up to 1 Tb/s in aggregate. Please refer to <https://aws.amazon.com/snowmobile/> for more information on how to get started.

44.0 Service Definition – Amazon DynamoDB

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

44.1 Service Overview

Amazon DynamoDB is a fast and flexible [NoSQL database](#) service for all applications that need consistent, single-digit millisecond latency at any scale. It is a fully managed cloud database and supports both document and key-value store models. Its flexible data model and reliable performance make it a great fit for mobile, web, gaming, ad tech, Internet of Things (IoT), and many other applications.

Top benefits include:

- **Fast, Consistent Performance** – Amazon DynamoDB is designed to deliver consistent, fast performance at any scale for all applications. Average service-side latencies are typically single-digit milliseconds. As your data volumes grow and application performance demands increase, Amazon DynamoDB uses automatic partitioning and SSD technologies to meet your throughput requirements and deliver low latencies at any scale.
- **Highly Scalable** – When creating a table, simply specify how much request capacity you require. If your throughput requirements change, simply update your table's request capacity using the AWS Management Console or the Amazon DynamoDB APIs. Amazon DynamoDB manages all the scaling behind the scenes, and you are still able to achieve your prior throughput levels while scaling is underway.
- **Fully Managed** – Amazon DynamoDB is a fully managed cloud NoSQL database service—you simply create a database table, set your throughput, and let the service handle the rest. You no longer need to worry about database management tasks such as hardware or software provisioning, setup and configuration, software patching, operating a reliable, distributed database cluster, or partitioning data over multiple instances as you scale.
- **Event-Driven Programming** – Amazon DynamoDB integrates with [AWS Lambda](#) to provide Triggers, which enables you to architect applications that automatically react to data changes.
- **Fine-Grained Access Control** – Amazon DynamoDB integrates with AWS IAM for fine-grained access control for users within your organisation. You can assign unique security credentials to each user and control each user's access to services and resources.
- **Flexible** – Amazon DynamoDB supports both document and key-value data structures, giving you the flexibility to design the best architecture that is optimal for your application.

44.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon DynamoDB. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

44.3 Service Constraints

Please see <https://aws.amazon.com/dynamodb/> for more information.

44.4 Technical Requirements

Amazon DynamoDB is a fully managed [NoSQL database](#) service that provides fast and predictable performance with seamless scalability. You can use Amazon DynamoDB to create a database table that can store and retrieve any amount of data and serve any level of request traffic. Amazon DynamoDB automatically spreads the data and traffic for the table over a sufficient number of servers to handle the request capacity specified by the customer and the amount of data stored, while maintaining consistent and fast performance.

- **Getting Started Guide** – Introduces you to Amazon DynamoDB, helps you set up DynamoDB Local, and walks you through how to use DynamoDB for the first time.

[HTML](#) | [PDF](#) | [Kindle](#)

- **Developer Guide** – Provides a conceptual overview of Amazon DynamoDB and includes detailed development instructions for using the various features.

[HTML](#) | [PDF](#) | [Kindle](#)

- **API Reference** – Describes all the API operations for Amazon DynamoDB in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

[HTML](#) | [PDF](#)

- **Streams API Reference** – Describes the API operations for DynamoDB Streams.

[HTML](#) | [PDF](#)

45.0 Service Definition – Amazon ElastiCache

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

45.1 Service Overview

Amazon ElastiCache is a web service that makes it easy to deploy, operate, and scale an in-memory cache in the cloud. The service improves the performance of web applications by allowing you to retrieve information from fast, managed, in-memory caches, instead of relying entirely on slower, disk-based databases. ElastiCache supports two open-source, in-memory caching engines:

- **Memcached** – A widely adopted memory object caching system. ElastiCache is protocol compliant with Memcached, so popular tools that you use today with existing Memcached environments will work seamlessly with the service.
- **Redis** – A popular open-source, in-memory key-value store that supports data structures such as sorted sets and lists. ElastiCache supports Master/Slave replication and Multi-AZ, which can be used to achieve cross-AZ redundancy.

Amazon ElastiCache automatically detects and replaces failed nodes, reducing the overhead associated with self-managed infrastructures and provides a resilient system that mitigates the risk of overloaded databases, which slow website and application load times. Through integration with [Amazon CloudWatch](#), Amazon ElastiCache provides enhanced visibility into key performance metrics associated with your Memcached or Redis nodes.

Using Amazon ElastiCache, you can add an in-memory caching layer to your infrastructure in a matter of minutes by using the [AWS Management Console](#).

Top benefits include:

- **Simple to Deploy** – Amazon ElastiCache makes it very easy to deploy a Memcached or Redis-compliant cache environment. Use the [AWS Management Console](#) or simple API calls to access the capabilities of a production-ready cloud cache cluster in minutes without worrying about infrastructure provisioning, or installing and maintaining cache software.
- **Managed** – Amazon ElastiCache automates time-consuming management tasks—such as patch management, failure detection, and recovery—allowing you to pursue higher-value application development.
- **Compatible** – With Amazon ElastiCache, you get native access to the Memcached or Redis in-memory caching environments. This facilitates compatibility with your existing tools and applications.
- **Elastic** – With a simple API call or a few clicks of the AWS Management Console, you can add or delete Cache Nodes to your cloud cache cluster to meet your application load. [Auto Discovery](#) for Memcached enables automatic discovery of Cache Nodes by ElastiCache Clients when the nodes are added to or removed from an Amazon ElastiCache Cluster.

- **Reliable** – Amazon ElastiCache has multiple features that enhance reliability for critical production deployments, including automatic failure detection and recovery. Amazon ElastiCache runs on the same highly reliable infrastructure used by other AWS Cloud services.
- **Integrated** – Amazon ElastiCache is designed for seamless use with other AWS Cloud services, including Amazon RDS, Amazon DynamoDB, Amazon EC2, Amazon CloudWatch, and Amazon SNS.
- **Secure** – Amazon ElastiCache provides a number of mechanisms to secure your Cache Cluster.
 - Amazon ElastiCache provides interfaces that allow you to configure firewall settings that control network access to your Cache Cluster.
 - Amazon ElastiCache allows you to run your Cache Cluster in [Amazon VPC](#). Amazon VPC enables you to isolate your Cache Cluster by specifying the IP ranges you wish to use for your Cache Nodes, and connect to your existing applications inside Amazon VPC. To learn more about Amazon ElastiCache in Amazon VPC, refer to the [Amazon ElastiCache User Guide](#).
- **Cost Effective** – Amazon ElastiCache saves you the administrative cost of setting up and managing a multi-node Cache Cluster. You can scale up and scale down the number of Cache Nodes in your Cache Cluster to deliver optimum performance as your application usage pattern changes, paying only for the resources you actually consume. The on-demand pricing allows you to pay for memory/compute capacity by the hour with no long-term commitments. This makes the use of Amazon ElastiCache very cost effective and frees you from the costs and complexities of planning, purchasing, and maintaining hardware.
- **Multi-AZ** – Amazon ElastiCache provides replication features for the Redis engine and Multi-AZ functionality. You can take advantage of multiple AWS Availability Zones to gain availability and scale beyond the capacity constraints of a single cache node. In case of primary node loss, Amazon ElastiCache will automatically detect the failure and failover to a read replica to provide higher availability without the need for manual intervention.

45.2 Backup/Restore and Disaster Recovery

Amazon ElastiCache for Redis helps you protect your data by creating snapshots of your clusters. Via a few clicks on the console or simple API calls, you can set up automatic snapshots, as well as initiate a backup whenever you choose. The snapshots can then be used for seeding new ElastiCache for Redis clusters. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

45.3 Service Constraints

Please see <https://aws.amazon.com/elasticache/> for more information.

45.4 Technical Requirements

Amazon ElastiCache is a web service that makes it easy to set up, manage, and scale distributed in-memory cache environments in the cloud. It provides a high-performance,

resisable, and cost-effective in-memory cache, while removing the complexity associated with deploying and managing a distributed cache environment.

- **User Guide** – Helps you understand the components and features that Amazon ElastiCache provides and how to use them. You'll learn how to access ElastiCache through a web-based GUI, Command Line Tools, and the Amazon ElastiCache API.
[HTML](#) | [PDF](#) | [Kindle](#)
- **ElastiCache section of the AWS CLI Reference** – Describes the CLI for Amazon ElastiCache in detail. Provides basic syntax, options, and usage examples for each command.
[HTML](#)
- **API Reference** – Describes all the API operations for Amazon ElastiCache in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

46.0 Service Definition – AWS Directory Service

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

46.1 Service Overview

AWS Directory Service makes it easy to setup and run Microsoft Active Directory (AD) in the AWS Cloud or connect your AWS resources with an existing on-premises Microsoft Active Directory. Once your directory is created, you can use it to manage users and groups, provide single sign-on to applications and services, create and apply group policy, domain join Amazon EC2 instances, as well as simplify the deployment and management of cloud-based Linux and Microsoft Windows workloads.

AWS Directory Service provides you with three directory types to choose from, including AWS Directory Service for Microsoft Active Directory (Enterprise Edition), also referred to as Microsoft AD, as well as Simple AD and AD Connector.

Top benefits include:

- **Simplify Deployment of Linux and Windows Workloads on AWS** – You can use AWS Directory Service to provision a managed Microsoft Active Directory, enabling you to manage users and group memberships, domain-joining Amazon EC2 Linux and Windows computers, perform Kerberos Single Sign-On (SSO), apply Group Policies and create trust relationships between domains. This makes it easier to deploy and manage Amazon EC2 instances and deploy directory-aware Windows-based workloads, including SharePoint, custom .NET, and SQL Server-based applications.
- **Easy to Get Started; Pay as You Go** – Getting started is easy. You can use the AWS Management Console, or the API to provision the directory type that meets your needs. Once your directory is running, you pay only for the directory hours you use, whether you require a directory for a handful of users or tens of thousands of users.
- **Seamless End User Access to Applications** – AWS Directory Service enables your end users to use their existing corporate credentials when accessing AWS applications, such as Amazon WorkSpaces, Amazon WorkDocs, and Amazon WorkMail, as well as directory-aware Microsoft applications, including SharePoint, custom .NET, and SQL Server-based applications.
- **Managed Service** – AWS Directory Service helps to reduce management tasks. There is no need to build out your own complex, highly available directory topology because each directory is deployed across multiple Availability Zones, and monitoring automatically detects and replaces domain controllers that fail. In addition, data replication and automated daily snapshots are configured for you. There is no software to install and AWS handles all of the patching and software updates.
- **Seamless Administration of AWS Resources** – AWS Directory Service enables your IT administrators to use their existing corporate credentials to access the AWS Management Console via AWS IAM roles to manage all your AWS resources (e.g., Amazon EC2 instances or Amazon S3 buckets). There is no need to stand up and manage federation infrastructure of your own.

46.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Directory Service. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

46.3 Service Constraints

Please see <http://aws.amazon.com/directoryservice/> for more information.

46.4 Technical Requirements

AWS Directory Service is a managed service that makes it easy to connect AWS Cloud services to your existing on-premises Microsoft Active Directory (AD Connector), or to set up and operate a new directory in the AWS Cloud (Simple AD and AWS Directory Service for Microsoft Active Directory). Your directory users and groups can access the

AWS Management Console and AWS applications, such as Amazon WorkSpaces and Amazon WorkDocs, using their existing credentials.

- **Administration Guide** – Describes how to create and manage an AWS Directory Service directory.
[HTML](#) | [PDF](#)
- **API Reference** – Describes the API operations for AWS Directory Service.
[HTML](#) | [PDF](#)

47.0 Service Definition – Amazon Inspector

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

47.1 Service Overview

Amazon Inspector is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assesses applications for vulnerabilities or deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritised by level of severity.

To help you get started quickly, Amazon Inspector includes a knowledge base of hundreds of rules mapped to common security best practices and vulnerability definitions. Examples of built-in rules include checking for remote root login being enabled or vulnerable software versions installed. These rules are regularly updated by AWS security researchers.

Top benefits include:

- **Identify Application Security Issues** – Amazon Inspector helps you to identify security vulnerabilities as well as deviations from security best practices in applications, both before they are deployed and while they are running in a production environment. This helps improve the overall security posture of your applications deployed on AWS.
- **Integrate Security into DevOps** – Amazon Inspector is agent-based, API-driven, and delivered as a service. This makes it easy for you to build right into your existing DevOps process, decentralizing and automating vulnerability assessments and empowering your development and operations teams to make security assessment an integral part of the deployment process.
- **Increase Development Agility** – Amazon Inspector helps you reduce the risk of introducing security issues during development and deployment by automating the security assessment of

your applications and proactively identifying vulnerabilities. This allows you to develop and iterate on new applications quickly and assess compliance with best practices and policies.

- **Leverage AWS Security Expertise** – The AWS security organisation is continuously assessing the AWS environment and updating a knowledge base of security best practices and rules. Amazon Inspector makes this expertise available to you in the form of a service that simplifies the process of establishing and enforcing best practices within your AWS environment.
- **Streamline Security Compliance** – Amazon Inspector gives security teams and auditors visibility into the security testing that is being performed during development of applications on AWS. This streamlines the process of validating and demonstrating that security and compliance standards and best practices are being followed throughout the development process.
- **Enforce Security Standards** – Amazon Inspector allows you to define standards and best practices for your applications and validate adherence to these standards. This simplifies enforcement of your organisation's security standards and best practices and helps to proactively manage security issues before they impact your production application.

47.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Inspector. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

47.3 Service Constraints

Please see <http://aws.amazon.com/inspector/> for more information.

47.4 Technical Requirements

Amazon Inspector is a security vulnerability assessment service that helps improve the security and compliance of your AWS resources. Amazon Inspector automatically assesses resources for vulnerabilities or deviations from best practices and then produces a detailed list of security findings prioritised by level of severity. Amazon Inspector includes a knowledge base of hundreds of rules mapped to common security standards and vulnerability definitions that are regularly updated by AWS security researchers.

- **User Guide** – Walks through how to set up Amazon Inspector and evaluate your security configuration.
[HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for Amazon Inspector in detail. Also provides sample requests, responses, and errors for the supported web

services

protocols.

[HTML](#) | [PDF](#)

- **Inspector section of the AWS CLI Reference** – Describes the AWS CLI commands that you can use to administer Amazon Inspector. Provides syntax, options, and usage examples for each command.

[HTML](#)

48.0 Service Definition – Amazon GuardDuty

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

48.1 Service Overview

Amazon GuardDuty is a managed threat detection service that continuously monitors for malicious or unauthorized behavior to help you protect your AWS accounts and workloads. It monitors for activity such as unusual API calls or potentially unauthorized deployments that indicate a possible account compromise. GuardDuty also detects potentially compromised instances or reconnaissance by attackers.

Enabled with a few clicks in the AWS Management Console, Amazon GuardDuty can immediately begin analyzing billions of events across your AWS accounts for signs of risk. GuardDuty identifies suspected attackers through integrated threat intelligence feeds and uses machine learning to detect anomalies in account and workload activity. When a potential threat is detected, the service delivers a detailed security alert to the GuardDuty console and AWS CloudWatch Events. This makes alerts actionable and easy to integrate into existing event management and workflow systems.

Amazon GuardDuty is cost effective and easy. It does not require you to deploy and maintain software or security infrastructure, meaning it can be enabled quickly with no risk of negatively impacting existing application workloads. There are no upfront costs with GuardDuty, no software to deploy, and no threat intelligence feeds required. Customers pay for the events analyzed by GuardDuty and there is a 30-day free trial available for every new account to the service..

Top benefits include:

- **Intelligent threat detection** - Amazon GuardDuty gives you intelligent threat detection by collecting, analyzing, and correlating billions of events from AWS CloudTrail, Amazon VPC Flow Logs, and DNS Logs across all of your associated AWS accounts. GuardDuty detections are made more accurate by incorporating threat intelligence (such as lists of known malicious IP addresses provided by AWS Security and 3rd party threat intelligence partners). GuardDuty also uses machine learning to detect anomalous account and network activities. For example, GuardDuty will alert you if it detects remote API calls from a known malicious IP address indicating potentially compromised AWS credentials. GuardDuty also detects direct threats to your AWS environment indicating a compromised instance, such as an Amazon EC2 instance sending encoded data within DNS queries.
- **Centralize threat detection across accounts** - Many organizations use multiple AWS accounts to help provide proper cost allocation, agility, and security. With a few clicks in the AWS Management Console, you can centralize your threat detection by enabling Amazon GuardDuty across any of your AWS accounts. With GuardDuty, there is no need to install additional security software or infrastructure to analyze your account and workload activity data. Your security operations center team can easily manage and triage threats from a single console view and automate security responses using a single security account.
- **Strengthens security through automation** - in addition to detecting threats, Amazon GuardDuty also makes it easy to automate how you respond to these threats, reducing your remediation and recovery time. You can set up your remediation scripts or AWS Lambda functions to trigger based on GuardDuty findings. GuardDuty security findings include the affected resource's details, such as tags, security groups, or credentials. GuardDuty findings also include attacker information, such as IP address and geo-location. This makes GuardDuty security findings informative and actionable. For example, account compromise can be difficult to detect quickly if you are not continuously monitoring account activities in near real-time. With GuardDuty, when an instance is suspected of having data stolen the service will alert you to be able to automatically create an access control entry restricting outbound access for that instance..

48.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon GuardDuty. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

48.3 Service Constraints

Please see <https://aws.amazon.com/guardduty/resources/> for more information.

48.4 Technical Requirements

- **User Guide** – Walks through how to set up Amazon GuardDuty and evaluate the security of your AWS environment.

[HTML](#) | [PDF](#)

- **API Reference** – Describes all of the API operations for Amazon GuardDuty. Also provides sample requests, responses, and errors for the supported web services protocols.

[HTML](#) | [PDF](#)

49.0 Service Definition – Amazon Macie

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

49.1 Service Overview

Amazon Macie is a security service that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. Amazon Macie recognizes sensitive data such as personally identifiable information (PII) or intellectual property, and provides you with dashboards and alerts that give visibility into how this data is being accessed or moved. The fully managed service continuously monitors data access activity for anomalies, and generates detailed alerts when it detects risk of unauthorized access or inadvertent data leaks. Today, Amazon Macie is available to protect data stored in Amazon S3, with support for additional AWS data stores coming later this year.

Top benefits include:

- **Superior Visibility Of Your Data** - Amazon Macie makes it easy for security administrators to have management visibility into data storage environments, beginning with S3, with additional AWS data stores coming soon.
- **Simple To Set Up, Easy To Manage** - Getting started with Amazon Macie is fast and easy. Simply log into the AWS console, select the Amazon Macie service, and provide the AWS accounts you would like to protect.
- **Data Security Automation Through Machine Learning** - Amazon Macie uses machine learning to automate the process of discovering, classifying, and protecting data stored in AWS. This helps you better understand where sensitive information is stored and how it's being accessed, including user authentications and access patterns.
- **Custom Alert Monitoring With Cloudwatch** - Amazon Macie can send all findings to Amazon CloudWatch Events. This allows you to build custom remediation and alert management for your existing security ticketing systems

49.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Macie. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

49.3 Service Constraints

Please see <http://aws.amazon.com/inspector/> for more information.

49.4 Technical Requirements

Amazon Macie is a security service that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. Amazon Macie recognizes sensitive data such as personally identifiable information (PII) or intellectual property, and provides you with dashboards and alerts that give visibility into how this data is being accessed or moved.

- **User Guide** – Describes key concepts for Amazon Macie and provides instructions for using the Amazon Macie console.

[HTML](#) | [PDF](#) | [Kindle](#).

50.0 Service Definition – AWS Certificate Manager

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

50.1 Service Overview

AWS Certificate Manager is a service that lets you easily provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and your internal connected resources. SSL/TLS certificates are used to secure network communications and establish the identity of websites over the Internet as well as resources on private networks. AWS Certificate Manager removes the time-consuming manual process of purchasing, uploading, and renewing SSL/TLS certificates. With AWS Certificate Manager, you can quickly request a certificate, deploy it on ACM-integrated AWS resources, such as Elastic Load Balancers, Amazon CloudFront distributions, and APIs on API Gateway, and let AWS

Certificate Manager handle certificate renewals. It also enables you to create private certificates for your internal resources and manage the certificate lifecycle centrally. Public and private certificates provisioned through AWS Certificate Manager for use with ACM-integrated services are free. You pay only for the AWS resources you create to run your application. For private certificates, you pay monthly for the operation of the private CA and for the private certificates you issue..

Top benefits include:

- **Protect and Secure Your Website** - SSL, and its successor TLS, are industry standard protocols for encrypting network communications and establishing the identity of websites over the Internet. SSL/TLS provides encryption for sensitive data in transit and authentication using SSL/TLS certificates to establish the identity of your site and secure connections between browsers and applications and your site. AWS Certificate Manager provides an easy way to provision and manage these certificates so you can configure a website or application to use the SSL/TLS protocol.
- **Protect and Secure Your Internal Resources** - Private certificates are used for identifying and securing communication between connected resources on private networks, such as servers, mobile and IoT devices, and applications. AWS Certificate Manager (ACM) Private Certificate Authority (CA) is a managed private CA service that helps you easily and securely manage the lifecycle of your private certificates. ACM Private CA provides you a highly-available private CA service without the upfront investment and ongoing maintenance costs of operating your own private CA. ACM Private CA extends ACM's certificate management capabilities to private certificates, enabling you to create and manage public and private certificates centrally. ACM Private CA allows developers to be more agile by providing them APIs to create and deploy private certificates programmatically. You also have the flexibility to create private certificates for applications that require custom certificate lifetimes or resource names. Learn more about ACM Private Certificate Authority.
- **Get Certificates Easily** - AWS Certificate Manager removes many of the time-consuming and error-prone steps to acquire an SSL/TLS certificate for your website or application. There is no need to generate a key pair or certificate signing request (CSR), submit a CSR to a Certificate Authority, or upload and install the certificate once received. With a few clicks in the AWS Management Console, you can request a trusted SSL/TLS certificate from AWS. Once the certificate is created, AWS Certificate Manager takes care of deploying certificates to help you enable SSL/TLS for your website or application.
- **Free Public Certificates for ACM-integrated Services** - With AWS Certificate Manager, there is no additional charge for provisioning public or private SSL/TLS certificates you use with ACM-integrated services, such as Elastic Load Balancing and API Gateway. You pay for the AWS resources you create to run your application. For private certificates, ACM Private CA provides you the ability to pay monthly for the service and certificates you create. You pay less per certificate as you create more private certificates.
- **Managed Certificate Renewal** - AWS Certificate Manager manages the renewal process for the certificates managed in ACM and used with ACM-integrated services, such as Elastic Load Balancing and API Gateway. ACM can automate renewal and deployment of these certificates. With ACM Private CA APIs, ACM enables you to automate creation and renewal of private certificates for on-premises resources, EC2 instances, and IoT devices.

- **Secure Key Management** - AWS Certificate Manager is designed to protect and manage the private keys used with SSL/TLS certificates. Strong encryption and key management best practices are used when protecting and storing private keys.
- **Centrally Manage Certificates on the AWS Cloud** - You will find it easy to centrally manage AWS Certificate Manager SSL/TLS certificates provided by AWS Certificate Manager in an AWS Region from the AWS Management Console, AWS CLI, or AWS Certificate Manager APIs. You can also audit the use of each certificate by reviewing your Amazon CloudTrail logs.
- **Integrated with Other AWS Cloud Services** - AWS Certificate Manager is integrated with other AWS services, so you can provision an SSL/TLS certificate and deploy it with your Elastic Load Balancer, Amazon CloudFront distribution or API in Amazon API Gateway. AWS Certificate Manager also works with AWS Elastic Beanstalk and AWS CloudFormation for public email-validated certificates to help you manage public certificates and use them with your applications in the AWS Cloud. To deploy a certificate with an AWS resource, you simply select the certificate you want from a drop-down list in the AWS Management Console. Alternatively, you can call an AWS API or CLI to associate the certificate with your resource. AWS Certificate Manager then deploys the certificate to the selected resource for you.
- **Import Third-Party Certificates** - AWS Certificate Manager makes it easy to import SSL/TLS certificates issued by third-party Certificate Authorities (CAs) and deploy them with your Elastic Load Balancers, Amazon CloudFront distributions and APIs on Amazon API Gateway. You can monitor the expiration date of an imported certificate, and import a replacement when the existing certificate is nearing expiration. Alternatively, you can request a free certificate from AWS Certificate Manager and let AWS manage future renewals for you. Importing certificates doesn't cost anything.

50.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Certificate Manager. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

50.3 Service Constraints

Please see <https://aws.amazon.com/certificate-manager/> for more information.

50.4 Technical Requirements

AWS Certificate Manager (ACM) makes it easy to provision, manage, and deploy SSL/TLS certificates on AWS-managed resources.

- **User Guide** – Provides conceptual overviews and explains how to provision ACM Certificates on AWS-based websites.

[HTML](#) | [PDF](#) | [Kindle](#).

- **API Reference** - Describes the API operations available for ACM along with sample requests, responses, and errors for the supported web services protocols.

[HTML](#) | [PDF](#)

51.0 Service Definition – AWS OpsWorks

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

51.1 Service Overview

AWS OpsWorks is a configuration management service that helps you configure and operate applications of all shapes and sizes using Chef. You can define the application's architecture and the specification of each component, including package installation, software configuration, and resources such as storage. Start from templates for common technologies like application servers and databases or build your own to perform any task that can be scripted. AWS OpsWorks includes automation to scale your application based on time or load and dynamic configuration to orchestrate changes as your environment scales.

Top benefits include:

- **Supports Any Application** – AWS OpsWorks supports a wide variety of architectures, from simple web applications to highly complex custom applications and any software that has a scripted installation. Since AWS OpsWorks supports Chef recipes and Bash/PowerShell scripts, you can leverage community-built configurations such as MongoDB and Elasticsearch. You start by modelling and visualizing your application with layers that define resource and software configuration. You control every aspect of your application's configuration to match your needs, processes, and tools.
- **Configuration as Code** – AWS OpsWorks lets you define configurations for your entire environment in a format that you can maintain and version just like your application source code. You can reproduce the software configuration on new instances and apply changes to all running instances, ensuring consistent configuration at any time. You can deploy your application from source repositories such as Git and Subversion to one or hundreds of instances with the click of a button.
- **Automation to Run at Scale** – AWS OpsWorks provides dynamic configuration and orchestration that allows you to efficiently manage your applications over their lifetime, including support for automatic instance scaling and auto healing. Each new instance that comes online is built to specification automatically and can change its configuration in response to system events.

This lets you, for example, change the configuration of a web application firewall when a new web server comes online without any manual steps.

- **Resource Organisation** – AWS OpsWorks lets you model and visualise your application using concepts such as stacks, layers, and apps. The AWS OpsWorks dashboard shows the status of your stacks across all AWS Regions. AWS OpsWorks tags your resources with the stack and layer names to make discovery easier and to support logical groupings for monitoring, cost allocation, and permissions. You can grant AWS IAM users access to specific stacks, making management of multi-user environments easier.
- **Supports Any Server** – AWS OpsWorks helps you automate operational tasks like software configurations, package installations, database setups, and code deployment on any Linux or Windows server, including existing Amazon EC2 instances or servers running in your own data centre. You can use a single configuration management service to deploy and operate applications across your hybrid architecture.

51.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS OpsWorks. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

51.3 Service Constraints

Please see <https://aws.amazon.com/opsworks/> for more information.

51.4 Technical Requirements

AWS OpsWorks provides a simple and flexible way to create and manage stacks and applications. With AWS OpsWorks, you can provision AWS resources, manage their configuration, deploy applications to those resources, and monitor their health.

- **User Guide** – Provides detailed descriptions of all the AWS OpsWorks concepts and provides instructions on using the various features with both the console and the CLI.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS OpsWorks in detail. In addition, it provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

52.0 Service Definition – AWS Systems Manager

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

52.1 Service Overview

AWS Systems Manager gives you visibility and control of your infrastructure on AWS. Systems Manager provides a unified user interface so you can view operational data from multiple AWS services and allows you to automate operational tasks across your AWS resources. With Systems Manager, you can group resources, like [Amazon EC2](#) instances, [Amazon S3](#) buckets, or [Amazon RDS](#) instances, by application, view operational data for monitoring and troubleshooting, and take action on your groups of resources. Systems Manager simplifies resource and application management, shortens the time to detect and resolve operational problems, and makes it easy to operate and manage your infrastructure securely at scale.

Top benefits include:

- **Shorten the Time to Detect Problems** - AWS Systems Manager helps you quickly view operational data for groups of resources, so you can quickly identify any issues that might impact applications that use those resources. Resources can be grouped by applications, application layers, production versus development environments, or anything else you choose. Systems Manager presents the operational data for your resource groups in a single, easy to read dashboard so you don't have to navigate to other AWS consoles. For example, if you have an application that uses [Amazon EC2](#), [Amazon S3](#), and [Amazon RDS](#), you can use Systems Manager to create a resource group for the application and easily see the software installed on your [Amazon EC2](#) instances, any changes in your [Amazon S3](#) objects, or database instances that have stopped.
- **Easy to Use Automation** - AWS Systems Manager allows you to automate operational actions to help make your teams more efficient. You can automate maintenance and deployment tasks on Amazon EC2 and on-premises instances, or automatically apply patches, updates, and configuration changes across any resource group. Systems Manager provides predefined automation documents for common operational tasks, such as stopping and restarting an EC2 instance,

that you can customize to your own specific use cases. Systems Manager also has built-in safety controls, allowing you to incrementally roll out new changes and automatically halt the roll-out when errors occur.

- **Improve Visibility and Control** - AWS Systems Manager helps you easily understand and control the current state of your resource groups. With Systems Manager, you can view detailed system configurations, operating system patch levels, software installations, application configurations, and other details about your environment through the Systems Manager dashboard. Systems Manager is integrated with [AWS Config](#) so you can easily view changes across your resources as they occur over time.
- **Manage Hybrid Environments** - With AWS Systems Manager, you can manage servers running on AWS and in your on-premises data center through a single interface. Systems Manager securely communicates with a lightweight agent installed on your servers to execute management tasks. This helps you manage resources for Windows and Linux operating systems running on [Amazon EC2](#) or on-premises.
- **Maintain Security and Compliance** - AWS Systems Manager helps maintain security and compliance by scanning your instances against your patch, configuration, and custom policies. You can define patch baselines, maintain up-to-date anti-virus definitions, and enforce firewall policies. You can also remotely manage your servers at scale without manually logging in to each server. Systems Manager also provides a centralized store to manage your configuration data, whether its plain text, such as database strings, or secrets, such as passwords. This allows you to separate your secrets and configuration data from code.

52.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Systems Manager. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

52.3 Service Constraints

Please see <https://aws.amazon.com/documentation/systems-manager/> for more information.

52.4 Technical Requirements

Use AWS Systems Manager to organize, monitor, and automate management tasks on your AWS resources.

- **User Guide** – Provides detailed descriptions of AWS Systems Manager concepts, and includes instructions for using both the console and the command-line interface to configure, manage, and automate tasks on groups of AWS instances and resources.

[HTML](#) | [PDF](#) | [Kindle](#)

- **AWS Systems Manager in the AWS CLI Reference** – Describes the AWS CLI commands that you can use to automate systems management tasks.
[HTML](#)
- **API Reference** - Describes the API operations for AWS Systems Manager in detail. In addition, the API Reference provides sample requests, responses, and errors for the supported web service protocols.

[HTML](#) | [PDF](#)

53.0 Service Definition – AWS CodeCommit

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

53.1 Service Overview

AWS CodeCommit is a fully managed source control service that makes it easy for companies to host secure and highly scalable private Git repositories. AWS CodeCommit eliminates the need to operate your own source control system or worry about scaling its infrastructure. You can use AWS CodeCommit to securely store anything from source code to binaries, and it works seamlessly with your existing Git tools.

Top benefits include:

- **Fully Managed** – AWS CodeCommit eliminates the need to host, maintain, back up, and scale your own source control servers. The service automatically scales to meet the growing needs of your project.
- **Secure** – AWS CodeCommit automatically encrypts your files in transit and at rest. AWS CodeCommit is integrated with AWS IAM, allowing you to assign user-specific permissions to your repositories.
- **High Availability** – AWS CodeCommit has a highly scalable, redundant, and durable architecture. The service is designed to keep your repositories highly available and accessible.
- **Store Anything** – AWS CodeCommit allows you to store any type of file, and there are no repository size limits. This allows you to store and version application assets, like images and libraries, along with your code.
- **Faster Development Lifecycle** – AWS CodeCommit keeps your repositories close to your build, staging, and production environments in the AWS Cloud. You can transfer incremental changes instead of the entire application. This allows you to increase the speed and frequency of your development lifecycle.
- **Use Your Existing Tools** – AWS CodeCommit supports all Git commands and works with your existing Git tools. You can keep using your preferred development environment plugins, continuous integration/continuous delivery systems, and graphical clients with AWS CodeCommit.

53.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS CodeCommit. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

53.3 Service Constraints

Please see <https://aws.amazon.com/codecommit/> for more information.

53.4 Technical Requirements

AWS CodeCommit is a version control service that enables you to privately store and manage Git repositories in the AWS Cloud.

- **User Guide** – Describes how to get started and store your assets using AWS CodeCommit.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS CodeCommit in detail. Also provides details of request and response syntax and errors for the supported web services protocols.
[HTML](#) | [PDF](#)
- **AWS CLI Reference for AWS CodeCommit** – Describes the AWS CLI commands that you can use to automate management of your assets.
[HTML](#)

54.0 Service Definition – AWS CodeDeploy

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

54.1 Service Overview

AWS CodeDeploy is a service that automates code deployments to any instance, including Amazon EC2 instances and instances running on-premises. AWS CodeDeploy makes it easier for you to rapidly release new features, helps you avoid downtime during application deployment, and handles the complexity of updating your applications. You can use AWS CodeDeploy to automate software deployments, eliminating the need for error-prone manual operations, and the service scales with your infrastructure so you can easily deploy to one instance or thousands.

Top benefits include:

- **Automated Deployments** – AWS CodeDeploy fully automates your code deployments, allowing you to deploy reliably and rapidly. You can consistently deploy your application across your development, test, and production environments on any instance, whether running in the cloud or on-premises. The service scales with your infrastructure so you can deploy to one instance or thousands.
- **Minimise Downtime** – AWS CodeDeploy helps maximise your application availability during the software deployment process. It performs rolling updates across your instances and tracks application health according to configurable rules. Code deployments can easily be stopped and rolled back if there are errors.

- **Centralised Control** – AWS CodeDeploy allows you to easily launch and track the status of your application deployments through the AWS Management Console or the AWS CLI. AWS CodeDeploy gives you a detailed report, allowing you to view when each application revision was deployed and to which instances. You can also create push notifications to receive live updates about your deployments.
- **Easy to Adopt** – AWS CodeDeploy is platform and language agnostic and works with any application. You can easily reuse your existing setup code. AWS CodeDeploy can also integrate with your existing software release process or continuous delivery toolchain (e.g., Jenkins).

54.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS CodeDeploy. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

54.3 Service Constraints

Please see <https://aws.amazon.com/codedeploy/> for more information.

54.4 Technical Requirements

AWS CodeDeploy is a deployment service that enables developers to automate the deployment of applications to instances and to update the applications as required.

- **User Guide** – Describes how to get started and deploy applications using AWS CodeDeploy.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS CodeDeploy. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)
- **AWS CLI Reference for AWS CodeDeploy** – Describes the AWS CLI commands that you can use to automate deployments.
[HTML](#)

55.0 Service Definition – AWS CodePipeline

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

55.1 Service Overview

AWS CodePipeline is a [continuous delivery](#) service for fast and reliable application updates. AWS CodePipeline builds, tests, and deploys your code every time there is a code change, based on the release process models you define. This enables you to rapidly and reliably deliver features and updates. You can easily build out an end-to-end solution by using our prebuilt plugins for popular third-party services like GitHub or integrating your own custom plugins into any stage of your release process. With AWS CodePipeline, you only pay for what you use. There are no up-front fees or long-term commitments.

Top benefits include:

- **Rapid Delivery** – AWS CodePipeline automates your software release process, allowing you to rapidly release new features to users. With AWS CodePipeline, you can quickly iterate on feedback and get new features to customers faster.
- **Improved Quality** – Automating your build, test, and release process allows you to easily test each code change and catch bugs while they are small and simple to fix. You can assure the quality of your code by running each change through your standardised release process.
- **Configurable Workflow** – AWS CodePipeline allows you to model the different stages of your software release process through a graphical user interface. You can specify the tests to run and the steps to deploy your application and its dependencies.
- **Get Started Fast** – With AWS CodePipeline, you can immediately begin to model your software release process. There are no servers to provision or set up. AWS CodePipeline is a fully managed continuous delivery service that connects to your existing tools and systems.
- **Easy to Integrate** – AWS CodePipeline can easily be extended to adapt to your specific needs. You can use our prebuilt plugins or your own custom plugins in any step of your release process. For example, you can pull your source code from GitHub, use your on-premises Jenkins build server, run load tests using a third-party service, or pass on deployment information to your custom operations dashboard.

55.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS CodePipeline. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

55.3 Service Constraints

Please see <https://aws.amazon.com/codepipeline/> for more information.

55.4 Technical Requirements

AWS CodePipeline is a continuous delivery service that enables you to model, visualise, and automate the steps required to release your software.

- **User Guide** – Walks through how to set up AWS CodePipeline and integrate it with other services.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS CodePipeline in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)
- **AWS CLI Reference for AWS CodePipeline** – Describes the AWS CLI commands that you can use to manage your assets.
[HTML](#)

56.0 Service Definition – AWS X-Ray

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

56.1 Service Overview

AWS X-Ray helps developers analyze and debug production, distributed applications, such as those built using a microservices architecture. With X-Ray, you can understand how your application and its underlying services are performing to identify and troubleshoot the root cause of performance issues and errors. X-Ray provides an end-to-end view of requests as they travel through your application, and shows a map of your application's underlying components. You can use X-Ray to analyze both applications in development and in production, from simple three-tier applications to complex microservices applications consisting of thousands of services.

Top benefits include:

- **Review Request Behavior** - AWS X-Ray traces user requests as they travel through your entire application. It aggregates the data generated by the individual services and resources that make up your application, providing you an end-to-end view of how your application is performing.

- **Discover Application Issues** - With AWS X-Ray, you can glean insights into how your application is performing and discover root causes. With X-Ray's tracing features, you can follow request paths to pinpoint where in your application and what is causing performance issues. X-Ray provides annotations so you can append metadata to traces, making it possible to tag and filter trace data so you can discover patterns and diagnose issues.
- **Improve Application Performance** - AWS X-Ray helps you identify performance bottlenecks. X-Ray's service maps let you see relationships between services and resources in your application in real time. You can quickly see where high latencies are occurring in your application and then drill down into the specific services and paths that are performing poorly.
- **Ready to Use with AWS** - AWS X-Ray works with [Amazon EC2](#), [Amazon EC2 Container Service](#) (Amazon ECS), and [AWS Elastic Beanstalk](#). You can use X-Ray with applications written in Java, Node.js, and .NET that are deployed on these services. Support for [AWS Lambda](#) is coming soon.
- **Designed for a Variety of Applications** - AWS X-Ray works for both simple and complex applications, either in development or in production. You can analyze simple asynchronous event calls, three-tier web applications, or complex microservices applications consisting of thousands of services. With X-Ray, you can trace requests made to applications that span multiple AWS accounts, AWS Regions, and Availability Zones.

56.2 Backup/Restore and Disaster Recovery

This requirement is AWS X-Ray. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

56.3 Service Constraints

Please see <https://aws.amazon.com/xray/> for more information.

56.4 Technical Requirements

AWS X-Ray makes it easy for developers to analyze the behavior of their distributed applications by providing request tracing, exception collection, and profiling capabilities.

- **Developer Guide** - Describes how to get started and instrument your applications using AWS X-Ray.
[HTML](#) | [PDF](#)
- **API Reference** - Describes all the API operations for AWS X-Ray in detail. Also provides details of request and response syntax and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

57.0 Service Definition – AWS Service Catalog

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

57.1 Service Overview

AWS Service Catalog allows organisations to create and manage catalogues of IT services that are approved for use on AWS. These IT services can include everything from virtual machine images, servers, software, and databases to complete multi-tier application architectures. AWS Service Catalog allows you to centrally manage commonly deployed IT services and helps you achieve consistent governance and meet your compliance requirements, while enabling users to quickly deploy only the approved IT services they need.

Top benefits include:

- **Ensure Compliance with Corporate Standards** – AWS Service Catalog provides a single location where organisations can centrally manage catalogues of IT services. With AWS Service Catalog, you can control which IT services and versions are available, the configuration of the available services, and permission access by individual, group, department, or cost centre.
- **Help Employees Quickly Find and Deploy Approved IT Services** – With AWS Service Catalog, you define your own catalogue of AWS Cloud services and AWS Marketplace software and make them available for your organisation. Then, end users can quickly discover and deploy IT services using a self-service portal.
- **Centrally Manage IT Service Lifecycle** – AWS Service Catalog enables you to add new versions of IT services as necessary, and end users are notified so they can easily keep abreast of the latest updates. With AWS Service Catalog you can control the use of IT services by specifying constraints, such as the AWS Region in which a product can be launched or allowed IP ranges.

57.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Service Catalog. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

57.3 Service Constraints

Please see <https://aws.amazon.com/servicecatalog/> for more information.

57.4 Technical Requirements

AWS Service Catalog allows IT administrators to create, manage, and distribute portfolios of approved products to end users, who can then access the products they need in a personalised portal. Typical products include servers, databases, websites, or applications that are deployed using AWS resources (for example, an Amazon EC2

instance or an Amazon RDS database). You can control which users have access to specific products to enforce compliance with organisational business standards, manage product lifecycles, and help users find and launch products with confidence.

- **Administrator Guide** – Provides a conceptual overview of AWS Service Catalog and includes detailed instructions for using the service as an administrator.
[HTML](#) | [PDF](#)
- **User Guide** – Provides a conceptual overview of AWS Service Catalog and includes detailed instructions for using the service as an end user.
[HTML](#) | [PDF](#)

58.0 CloudFormation

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

58.1 Service Overview

AWS CloudFormation gives developers and systems administrators an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.

You can use AWS CloudFormation's [sample templates](#) or create your own templates to describe the AWS resources, and any associated dependencies or runtime parameters, required to run your application. You don't need to figure out the order for provisioning AWS Cloud services or the subtleties of making those dependencies work. AWS CloudFormation takes care of this for you. After the AWS resources are deployed, you can modify and update them in a controlled and predictable way, in effect applying version control to your AWS infrastructure the same way you do with your software. You can also visualise your templates as diagrams and edit them using a drag-and-drop interface with the [AWS CloudFormation Designer](#).

You can deploy and update a template and its associated collection of resources (called a stack) by using the AWS Management Console, AWS CLI, or APIs. AWS CloudFormation is available at no additional charge, and you pay only for the AWS resources needed to run your applications.

Top benefits include:

- **Supports a Wide Range of AWS Resources** – AWS CloudFormation supports a wide range of AWS resources, allowing you to build a highly available, reliable, and scalable AWS infrastructure for your application needs.
- **Easy to Use** – AWS CloudFormation makes it easy to organise and deploy a collection of AWS resources and lets you describe any dependencies or special parameters to pass in at runtime. You can use one of the many CloudFormation [sample templates](#)—either verbatim or as a starting point.
- **Declarative and Flexible** – To create the infrastructure you want, you enumerate what AWS resources, configuration values, and interconnections you need in a template and then let AWS CloudFormation do the rest with a few simple clicks in the AWS Management Console, one command by using the AWS CLI, or a single request by calling the APIs. You won't have to recall the details of how to create and interconnect the respective AWS resources via their service APIs; AWS CloudFormation does this for you. You also don't need to write a template from scratch if you start with one of the many sample templates that come with AWS CloudFormation.
- **Infrastructure as Code** – A template can be used repeatedly to create identical copies of the same stack (or to use as a foundation to start a new stack). You can capture and control region-specific infrastructure variations such as Amazon EC2 AMIs, as well as Amazon EBS and Amazon RDS snapshot names. Templates are simple JSON-formatted text files that can be placed under your normal source control mechanisms, stored in private or public locations such as Amazon S3, and exchanged via email. With AWS CloudFormation, you can “open the hood,” to see exactly which AWS resources make up a stack. You retain full control and have the ability to modify any of the AWS resources created as part of a stack.
- **Customised Via Parameters** – You can use parameters to customise aspects of your template at run time when the stack is built. For example, you can pass the Amazon RDS database size, Amazon EC2 instance types, database, and web server port numbers to AWS CloudFormation when you create a stack. You can also use a parameterised template to create multiple stacks that may differ in a controlled way. For example, your Amazon EC2 instance types, Amazon CloudWatch alarm thresholds, and Amazon RDS read-replica settings may differ among AWS Regions if you receive more customer traffic in the US than in Europe. You can use template parameters to tune the settings and thresholds in each region separately and still be sure that the application is deployed consistently across the regions.
- **Visualise and Edit with Drag-and-Drop Interface** – [AWS CloudFormation Designer](#) provides a visual diagram of your template with icons representing your AWS resources and arrows showing their relationships. You can build and edit templates using the drag-and-drop interface, then edit the template details using the integrated JSON text editor. AWS CloudFormation Designer allows you to spend more time designing your AWS infrastructure and less time manually coding your templates.
- **Integration Ready** – You can integrate AWS CloudFormation with the development and management tools of your choice. AWS CloudFormation publishes progress events through the Amazon SNS. With Amazon SNS you can track stack creation and deletion progress via email and integrate with other processes programmatically.

58.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS CloudFormation. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

58.3 Service Constraints

Please see <https://aws.amazon.com/cloudformation/> for more information.

58.4 Technical Requirements

AWS CloudFormation enables you to create and provision AWS infrastructure deployments predictably and repeatedly. It helps you leverage AWS products such as Amazon EC2, Amazon EBS, Amazon SNS, Elastic Load Balancing, and Auto Scaling to build highly reliable, highly scalable, cost-effective applications in the cloud without worrying about creating and configuring the underlying AWS infrastructure. AWS CloudFormation enables you to use a template file to create and delete a collection of resources together as a single unit (a stack).

- **User Guide** – Provides a conceptual overview of AWS CloudFormation and includes instructions on using the various features with the AWS CLI.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS CloudFormation in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)
- **AWS CloudFormation in the AWS CLI Reference** – Describes the AWS CloudFormation commands that are available in the AWS CLI.
[HTML](#)

59.0 Service Definition – Amazon QuickSight

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

59.1 Service Overview

Amazon QuickSight is a very fast, cloud-powered Business Intelligence (BI) service that makes it easy for all employees to build visualizations, perform analysis, and quickly get business insights from their data. Amazon QuickSight uses a new, Super-fast, Parallel,

In-memory Calculation Engine (“SPICE”) to perform advanced calculations and render visualizations rapidly. Amazon QuickSight integrates automatically with AWS data services, enables organisations to scale to hundreds of thousands of users, and delivers fast and responsive query performance to them via SPICE’s query engine. At one tenth the cost of traditional solutions, Amazon QuickSight enables you to deliver rich BI functionality to everyone in your organisation.

Top benefits include:

- **Get Started Quickly** – No need to spend months building complex data models and invest in complex and costly software or hardware to generate the first report. Simply log in to Amazon QuickSight, point to a data source or upload a file, and begin visualizing your data with actionable insights in about a minute.
- **Access Data from Multiple Sources** – Easily connect Amazon QuickSight to AWS data services, including Amazon Redshift, Amazon RDS, Amazon Aurora, Amazon EMR, Amazon DynamoDB, Amazon S3, and Amazon Kinesis; upload CSV, TSV, and spreadsheet files; or connect to third-party data sources such as Salesforce.
- **Easy to Use** – Amazon QuickSight automatically infers data types and relationships and provides suggestions for the best possible visualizations, optimised for your data, to help you get quick, actionable business insights.
- **Get Answers Fast** – Amazon QuickSight uses [SPICE](#) to generate answers on large data sets.
- **Tell a Story with Your Data** – Securely [share your analysis](#) with others in your organisation by building interactive stories for collaboration using the StoryBoard and annotations. Recipients can further explore the data and respond back with their insights and knowledge, making the whole organisation efficient and effective.
- **Use AWS Partner BI Solutions with Amazon QuickSight** – Amazon QuickSight provides partners a simple SQL-like interface to query the data stored in SPICE so that you can continue using your existing BI tools from AWS BI partners while benefiting from the faster performance delivered by SPICE.

59.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon QuickSight. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

59.3 Service Constraints

Please see <https://aws.amazon.com/quicksight/> for more information.

59.4 Technical Requirements

Amazon QuickSight is a fast business analytics service to build visualizations, perform ad hoc analysis, and quickly get business insights from your data. QuickSight

seamlessly discovers AWS data sources, enables organisations to scale to hundreds of thousands of users, and delivers fast and responsive query performance by using a robust in-memory engine (SPICE).

- [User Guide](#)
Describes all Amazon QuickSight concepts and features, and provides instructions on using these features in the QuickSight web application.
[HTML](#) | [PDF](#) | [Kindle](#)

60.0 Service Definition – Amazon Machine Learning

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

60.1 Service Overview

Amazon Machine Learning is a service that makes it easy for developers of all skill levels to use ML technology. Amazon Machine Learning provides visualization tools and wizards that guide you through the process of creating ML models without having to learn complex ML algorithms and technology. Once your models are ready, Amazon Machine Learning makes it easy to obtain predictions for your application using simple APIs, without having to implement custom prediction generation code or manage any infrastructure.

Amazon Machine Learning is based on the same proven, highly scalable, ML technology used for years by Amazon's internal data scientist community. The service uses powerful algorithms to create ML models by finding patterns in your existing data. Then, Amazon Machine Learning uses these models to process new data and generate predictions for your application.

Amazon Machine Learning is highly scalable and can generate billions of predictions daily and serve those predictions in real time and at high throughput. With Amazon Machine Learning, there is no up-front hardware or software investment, and you pay as you go, so you can start small and scale as your application grows.

Top benefits include:

- **Easily Create ML Models** – Amazon Machine Learning APIs and wizards make it easy for any developer to create and fine-tune ML models from data stored in [Amazon S3](#), [Amazon Redshift](#),

or MySQL databases in [Amazon RDS](#), and query these models for predictions. The service's built-in data processors, scalable ML algorithms, interactive data and model visualization tools, and quality alerts help you build and refine your models quickly.

- **From Models to Predictions in Seconds** – Amazon Machine Learning is a managed service that provides end-to-end model creation, deployment, and monitoring. Once your model is ready, you can quickly and reliably generate predictions for your applications, eliminating the time and investment needed to build, scale, and maintain ML infrastructure.
- **Scalable, High-Performance Prediction Generation Service** – Amazon Machine Learning prediction APIs can be used to generate billions of predictions for your applications. You can request predictions for large numbers of data records all at once using the batch prediction API or use the real-time API to obtain predictions for individual data records and use them within interactive web, mobile, or desktop applications.
- **Low Cost and Efficient** – With Amazon Machine Learning there is no setup cost and you pay as you go, so you can start small and scale as your application grows.
- **Leverage Proven Technology** – Amazon Machine Learning is based on the same proven, highly scalable, ML technology used by Amazon to perform critical functions like supply chain management, fraudulent transaction identification, and catalogue organisation.

60.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Machine Learning. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

60.3 Service Constraints

Please see <https://aws.amazon.com/machine-learning/> for more information.

60.4 Technical Requirements

Amazon Machine Learning makes it easy for developers to build smart applications, including applications for fraud detection, demand forecasting, targeted marketing, and click prediction. The powerful algorithms of Amazon Machine Learning create ML models by finding patterns in your existing data. The service uses these models to process new data and generate predictions for your application.

- **Developer Guide** – Provides a conceptual overview of Amazon Machine Learning and includes detailed instructions for using the service. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for Amazon Machine Learning in detail. Also provides sample requests and responses for supported web

service

protocols.

[HTML](#) | [PDF](#)

61.0 Service Definition – Amazon Translate

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

61.1 Service Overview

Amazon Translate is a neural machine translation service that delivers fast, high-quality, and affordable language translation. Neural machine translation is a form of language translation automation that uses deep learning models to deliver more accurate and more natural sounding translation than traditional statistical and rule-based translation algorithms. Amazon Translate allows you to localize content - such as websites and applications - for international users, and to easily translate large volumes of text efficiently.

Top benefits include:

- **Highly Accurate & Continuously Learning** - Amazon Translate is a neural machine translation service. The translation engines are always learning from new and expanded datasets to produce more accurate and fluent translations for a wide range of use cases.
- **Easy to Integrate into Your Applications** - Amazon Translate removes the complexity of building real-time and batch translation capabilities into your applications with a simple API call. This makes it easy to localize an application or web site, or process multilingual data within your existing workflows.
- **Scalable** - Whether it's just a few words or large volumes of text, Amazon Translate scales easily as your translation needs grow. The service provides consistently fast and reliable translations, regardless of the volume of translation requests you make.

61.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Translate. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

61.3 Service Constraints

Please see <https://aws.amazon.com/translate/resources/> for more information.

61.4 Technical Requirements

Amazon Translate uses advanced machine learning technologies to provide high-quality translation on demand. Use it to translate unstructured text documents or to build applications that work in multiple languages. For additional resources and information, please see <https://aws.amazon.com/translate/resources/>.

- **Developer Guide** – Provides a conceptual overview of Amazon Translate, includes detailed instructions for using the various features, and provides a complete API reference for developers.

[HTML](#) | [PDF](#)

62.0 Service Definition – Amazon Transcribe

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

62.1 Service Overview

Amazon Transcribe is an automatic speech recognition (ASR) service that makes it easy for developers to add speech-to-text capability to their applications. Using the Amazon Transcribe API, you can analyze audio files stored in Amazon S3 and have the service return a text file of the transcribed speech.

Amazon Transcribe can be used for lots of common applications, including the transcription of customer service calls and generating subtitles on audio and video content. The service can transcribe audio files stored in common formats, like WAV and MP3, with time stamps for every word so that you can easily locate the audio in the original source by searching for the text. Amazon Transcribe is continually learning and improving to keep pace with the evolution of language.

Top benefits include:

- **Simple-to-Use API** - The Amazon Transcribe API makes it easy to convert speech to text. No complicated programming is required. Just call the API with a few lines of code, and Amazon Transcribe will return the text from your audio file stored in Amazon S3.
- **Support for a Wide Range of Use Cases** - Amazon Transcribe is designed to provide accurate and automated transcripts for a wide range of audio quality. You can generate subtitles for any

video or audio files, and even transcribe low quality telephony recordings such as customer service calls.

- **Easy-to-Read Transcriptions** - Most speech recognition systems output a string of text without punctuation. Amazon Transcribe uses deep learning to add punctuation and formatting automatically, so that the output is more intelligible and can be used without any further editing.
- **Custom Vocabulary** - Amazon Transcribe gives you the ability to expand and customize the speech recognition vocabulary. You can add new words to the base vocabulary and generate highly-accurate transcriptions specific to your use case, such as product names, domain-specific terminology, or names of individuals.
- **Timestamp Generation** - Amazon Transcribe returns a timestamp for each word, so that you can easily locate the audio in the original recording by searching for the text.
- **Recognize Multiple Speakers** - Amazon Transcribe is able to recognize when the speaker changes and attribute the transcribed text appropriately. This can significantly reduce the amount of work needed to transcribe audio with multiple speakers like telephone calls, meetings, and television shows.

62.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Transcribe. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

62.3 Service Constraints

Please see <https://aws.amazon.com/transcribe/> for more information.

62.4 Technical Requirements

For additional resources and information, please see <https://aws.amazon.com/transcribe/resources/>.

- **Developer Guide** – Provides a conceptual overview of Amazon Transcribe, includes detailed instructions for using the various features, and provides a complete API reference for developers.

[HTML](#) | [PDF](#)

63.0 Service Definition – AWS DeepLens

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

63.1 Service Overview

AWS DeepLens is the world's first deep-learning enabled video camera for developers of all skill levels to grow their machine learning skills through hands-on computer vision tutorials, example code, and pre-built models. Customers can run the included sample projects as is, connect them with other AWS Cloud services, train a model in Amazon SageMaker and deploy it to AWS DeepLens, or extend the functionality by triggering a lambda function when an action takes place. Customers can even apply more advanced analytics on the cloud using Amazon Kinesis Video Streams and Amazon Rekognition video. AWS DeepLens provides the building blocks for machine learning needs.

Top benefits include:

- **A new way to learn machine learning** - AWS DeepLens allows developers of all skill levels to get started with deep learning in less than 10 minutes by providing sample projects with practical, hands-on examples which can start running with a single click.
- **Custom built for deep learning** - AWS DeepLens was designed with deep learning in mind. With over 100 GFLOPS of compute power on the device, it can process deep learning predictions on HD video for real time.
- **Build custom models with Amazon SageMaker** - Models trained in [Amazon SageMaker](#) can be sent to AWS DeepLens with just a few clicks from the AWS Management Console.
- **Broad framework support** - AWS developers can run any deep learning framework, including [TensorFlow](#) and Caffe. AWS DeepLens comes pre-installed with a high performance, efficient, optimized inference engine for deep learning using [Apache MXNet](#).
- **Integrated with AWS** - AWS DeepLens integrates with [Amazon Rekognition](#) for advanced image analysis, [Amazon SageMaker](#) for training models, and with [Amazon Polly](#) to create speech-enabled projects. The device also connects securely to [AWS IoT](#), [Amazon SQS](#), [Amazon SNS](#), [Amazon S3](#), [Amazon DynamoDB](#), and more.
- **Fully programmable** - AWS DeepLens is easy to customize and is fully programmable using [AWS Lambda](#). The deep learning models in DeepLens

even run as part of an [AWS Lambda](#) function, providing a familiar programming environment to experiment with.

63.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS DeepLens. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

63.3 Service Constraints

Please see https://aws.amazon.com/documentation/deeplens/?nc1=h_ls for more information.

63.4 Technical Requirements

AWS DeepLens pairs a connected HD camera developer kit with a set of sample projects to help developers learn machine learning concepts using hands-on computer vision use cases. AWS DeepLens is preconfigured to work with multiple AWS services to provide a deep learning framework that is optimized for AWS DeepLens, so it is easy to build computer vision applications. AWS DeepLens is powerful enough for experts, but designed to help all developers get started quickly with little to no experience in deep learning.

- **Developer Guide** - Helps you understand the components and features that AWS DeepLens provides and how to use them. You'll learn how to access and use AWSs through a web-based GUI.

[HTML](#) | [PDF](#)

64.0 Service Definition – Amazon Transcribe

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

64.1 Service Overview

Amazon Transcribe is an automatic speech recognition (ASR) service that makes it easy for developers to add speech-to-text capability to their applications. Using the

Amazon Transcribe API, you can analyze audio files stored in Amazon S3 and have the service return a text file of the transcribed speech.

Amazon Transcribe can be used for lots of common applications, including the transcription of customer service calls and generating subtitles on audio and video content. The service can transcribe audio files stored in common formats, like WAV and MP3, with time stamps for every word so that you can easily locate the audio in the original source by searching for the text. Amazon Transcribe is continually learning and improving to keep pace with the evolution of language.

Top benefits include:

- **Simple-to-Use API** - The Amazon Transcribe API makes it easy to convert speech to text. No complicated programming is required. Just call the API with a few lines of code, and Amazon Transcribe will return the text from your audio file stored in Amazon S3.
- **Support for a Wide Range of Use Cases** - Amazon Transcribe is designed to provide accurate and automated transcripts for a wide range of audio quality. You can generate subtitles for any video or audio files, and even transcribe low quality telephony recordings such as customer service calls.
- **Easy-to-Read Transcriptions** - Most speech recognition systems output a string of text without punctuation. Amazon Transcribe uses deep learning to add punctuation and formatting automatically, so that the output is more intelligible and can be used without any further editing.
- **Custom Vocabulary** - Amazon Transcribe gives you the ability to expand and customize the speech recognition vocabulary. You can add new words to the base vocabulary and generate highly-accurate transcriptions specific to your use case, such as product names, domain-specific terminology, or names of individuals.
- **Timestamp Generation** - Amazon Transcribe returns a timestamp for each word, so that you can easily locate the audio in the original recording by searching for the text.
- **Recognize Multiple Speakers** - Amazon Transcribe is able to recognize when the speaker changes and attribute the transcribed text appropriately. This can significantly reduce the amount of work needed to transcribe audio with multiple speakers like telephone calls, meetings, and television shows.

64.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Transcribe. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

64.3 Service Constraints

Please see <https://aws.amazon.com/transcribe/> for more information.

64.4 Technical Requirements

For additional resources and information, please see <https://aws.amazon.com/transcribe/resources/>.

- **Developer Guide** – Provides a conceptual overview of Amazon Transcribe, includes detailed instructions for using the various features, and provides a complete API reference for developers.

[HTML](#) | [PDF](#)

65.0 Service Definition – Amazon Kinesis

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

65.1 Service Overview

Amazon Kinesis is a platform for [streaming data](#) on AWS, offering powerful services to make it easy to load and analyse streaming data and also providing the ability for you to build custom streaming data applications for specialised needs. Web applications, mobile devices, wearables, industrial sensors, and many software applications and services can generate staggering amounts of streaming data—sometimes TBs per hour—that need to be collected, stored, and processed continuously. Amazon Kinesis services enable you to do that simply and at a low cost. Amazon Kinesis consists of Amazon Kinesis Firehose and Amazon Kinesis Streams.

Amazon Kinesis Firehose is the easiest way to load [streaming data](#) into AWS. It can capture and automatically load streaming data into [Amazon S3](#), [Amazon Redshift](#), and [Amazon Elasticsearch Service](#), enabling near real-time analytics with existing business intelligence tools and dashboards you're already using today. It is a fully managed service that automatically scales to match the throughput of your data and requires no ongoing administration. It can also batch, compress, and encrypt the data before loading it, minimizing the amount of storage used at the destination and increasing security. You can easily create a Firehose delivery stream from the AWS Management

Console, configure it with a few clicks, and start sending data to the stream from hundreds of thousands of data sources to be loaded continuously to AWS—all in just a few minutes. With Amazon Kinesis Firehose, you only pay for the amount of data you transmit through the service. There is no minimum fee or setup cost.

Amazon Kinesis Streams enables you to build custom applications that process or analyse [streaming data](#) for specialised needs. Amazon Kinesis Streams can continuously capture and store terabytes of data per hour from hundreds of thousands of sources such as website clickstreams, financial transactions, social media feeds, IT logs, and location-tracking events. With Amazon Kinesis Client Library (KCL), you can build Amazon Kinesis Applications and use streaming data to power real-time dashboards, generate alerts, implement dynamic pricing and advertising, and more. You can also emit data from Amazon Kinesis Streams to other AWS Cloud services such as Amazon S3, Amazon Redshift, Amazon EMR, and AWS Lambda.

65.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Kinesis. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

65.3 Service Constraints

Please see <https://aws.amazon.com/kinesis/> for more information.

65.4 Technical Requirements

Amazon Kinesis is a platform for streaming data on AWS, making it easy to load and analyse streaming data, and also providing the ability for you to build custom streaming data applications for specialised needs. Use Amazon Kinesis Streams to collect and process large streams of data records in real time. Use Amazon Kinesis Firehose to deliver real-time streaming data to destinations such as Amazon S3 and Amazon Redshift.

- Amazon Kinesis Streams

- **User Guide** – Provides a conceptual overview of Amazon Kinesis Streams and includes detailed development instructions for using the various features.
[HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for Amazon Kinesis Streams in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)
- Amazon Kinesis Firehose
 - **Developer Guide** – Provides a conceptual overview of Amazon Kinesis Firehose and includes detailed instructions for using the service.
[HTML](#) | [PDF](#) | [Kindle](#)
 - **API Reference** – Describes all the API operations for Amazon Kinesis Firehose in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

66.0 Service Definition – Amazon Elastic Transcoder

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

66.1 Service Overview

Amazon Elastic Transcoder is media transcoding in the cloud. It is designed to be a highly scalable, easy-to-use, and cost-effective way for developers and businesses to convert (or “transcode”) media files from their source format into versions that will playback on devices like smartphones, tablets, and PCs.

Top benefits include:

- **Easy to Use** – Amazon Elastic Transcoder is designed to be easy to use. You can get started by using the AWS Management Console, the service API, or the SDKs. System transcoding presets make it easy to get transcoding settings right the first time. We provide predefined presets to create media files that will play on a wide range of devices (like smartphones or tablets), as well as presets to create media files that are optimised for playback on a specific device (like the

Amazon Kindle Fire HD or Apple iPod touch). You can also create segmented files and playlists for delivery using the HLS, Smooth, or MPEG-DASH protocol to compatible devices. Developers building applications that need transcoding can use the [AWS SDKs](#) for Java, .NET, Node.js, PHP, Python and Ruby, and the [AWS CLI](#).

- **Elastically Scalable** – Amazon Elastic Transcoder is designed to scale seamlessly with your media transcoding workload. Amazon Elastic Transcoder is architected to handle large volumes of media files and large file sizes. Transcoding pipelines enable you to perform multiple transcodes in parallel. Amazon Elastic Transcoder leverages other AWS Cloud services like [Amazon S3](#), [Amazon EC2](#), [Amazon DynamoDB](#), [Amazon SWF](#), and [Amazon SNS](#) to give scalability and reliability.
- **Cost Effective** – Amazon Elastic Transcoder uses a content duration-based pricing model: you pay based on the length of the output, in minutes, of the media you are transcoding. For example, if your video's transcoded output is 30 minutes in duration, you will pay for 30 minutes of transcoding. Similarly, if you create a 20 minute video clip from a 30 minute input file, you will pay for 20 minutes of transcoding. With Amazon Elastic Transcoder, there are no minimum transcoding volumes, monthly commitments, or long-term contracts.
- **Managed** – Amazon Elastic Transcoder enables you to focus on your content rather than on managing transcoding software in a distributed cloud-based environment. The service takes care of scaling and operating the system and manages the process of keeping codecs up to date. Combined with our service API and SDKs, this makes it easy for you to build media solutions that use Amazon Elastic Transcoder.
- **Secure** – Your content is under your control: your assets are in your own Amazon S3 buckets, to which you give us access through AWS IAM roles. This makes it simple to fit seamlessly into your current security and identity framework without giving up control. In building Amazon Elastic Transcoder, we made use of security best practices gained by building other AWS Cloud services. For more information about security on AWS, please refer to the [AWS Security Center](#). For more information on compliance, including MPAA best practices, please refer to [AWS Compliance](#).
- **Seamless Delivery** – Using Amazon Elastic Transcoder, Amazon S3, and Amazon CloudFront, you can store, transcode, and deliver your content. By setting the Amazon S3 permissions for your CloudFront distribution in Amazon Elastic Transcoder, it is now a simple, one-step process to transcode content with Amazon Elastic Transcoder and deliver the multiple output videos via progressive download or adaptive bitrate streaming (HLS, Smooth, or MPEG-DASH) with CloudFront.
- **AWS Integration** – Amazon Elastic Transcoder provides an important media building block for creating end-to-end media solutions on AWS. For example, you can use Amazon Glacier to store master content, Amazon Elastic Transcoder to transcode masters to renditions for distribution stored in Amazon S3, and stream these renditions at scale over the Internet using Amazon CloudFront.

66.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Elastic Transcoder. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

66.3 Service Constraints

Please see <https://aws.amazon.com/elastictranscoder/> for more information.

66.4 Technical Requirements

Amazon Elastic Transcoder lets you convert media files that you have stored in Amazon S3 into media files in the formats required by consumer playback devices. For example, you can convert large, high-quality digital media files into formats that users can play back on mobile devices, tablets, web browsers, and connected televisions.

- **Developer Guide** – Provides a conceptual overview of Amazon Elastic Transcoder and includes detailed instructions for using its various features.

[HTML](#) | [PDF](#) | [Kindle](#)

67.0 Service Definition – AWS Elemental MediaConvert

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

67.1 Service Overview

AWS Elemental MediaConvert is a file-based video transcoding service with broadcast-grade features. It allows you to easily create video-on-demand (VOD) content for broadcast and multiscreen delivery at scale. The service combines advanced video and audio capabilities with a simple web services interface and pay-as-you-go pricing. With AWS Elemental MediaConvert, you can focus on delivering compelling media experiences without having to worry about the complexity of building and operating your own video processing infrastructure.

Top benefits include:

- **Broadcast-grade capabilities** - AWS Elemental MediaConvert lets you use a wide range of internet and professional media formats to produce high-quality video outputs that look great on any device. With support for ultra-high definition resolutions, high dynamic range video, graphic overlays, advanced audio features, content protection, and closed captioning, AWS Elemental MediaConvert offers a full set of tools to deliver high-quality viewing experiences.
- **Reliable and easy to manage** - AWS Elemental MediaConvert does not require any set up, management, or maintenance of underlying infrastructure. Simply submit jobs with the video processing settings you want and get started without spending time or resources managing transcoding

- infrastructure. Not only does AWS Elemental MediaConvert provision the required resources to process your jobs, but it also monitors them automatically, so you don't need to worry about reliability.
- **Simple, predictable pricing** - AWS Elemental MediaConvert lets customers create high-quality, end-to-end video processing workflows in the cloud without upfront investment or capital expenditures for video processing infrastructure. You simply pay based on the duration of video that is processed and the features you use.

67.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elemental MediaConvert. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/mediaconvert/>.

67.3 Service Constraints

Please see <https://aws.amazon.com/mediaconvert/resources/> for more information.

67.4 Technical Requirements

- **User Guide** - Describes the components and features that AWS Elemental MediaConvert provides and how to use them.
- [HTML | PDF](#)
- **API Reference** - Describes basic operations of AWS Elemental MediaConvert. Provides schema structure for job settings and detailed descriptions of encoding settings. Includes sample job requests.
- [HTML | PDF](#)

68.0 Service Definition – AWS Elemental MediaLive

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

68.1 Service Overview

AWS Elemental MediaLive is a broadcast-grade live video processing service. It lets you create high-quality video streams for delivery to broadcast televisions and internet-connected multiscreen devices, like connected TVs, tablets, smart phones, and set-top boxes. The service works by encoding your live video streams in real-time, taking a larger-sized live video source and compressing it into smaller versions for distribution to

your viewers. With AWS Elemental MediaLive, you can easily set up streams for both live events and 24x7 channels with advanced broadcasting features, high availability, and pay-as-you-go pricing. AWS Elemental MediaLive lets you focus on creating compelling live video experiences for your viewers without the complexity of building and operating broadcast-grade video processing infrastructure.

Top benefits include:

- **Simple deployment and management** - AWS Elemental MediaLive automates the provisioning and management of all the infrastructure used for video encoding, letting you deploy a simple live channel in minutes. The service transparently provisions resources and manages all the scaling, failover, monitoring, and reporting needed to power a live video stream. This lets you focus on your live content, not your encoding infrastructure.
- **Broadcast-grade capabilities** - AWS Elemental MediaLive makes it easy for anyone to produce broadcast-quality live streaming video. The service includes support for advanced capabilities such as ad marker support, audio features including audio normalization and Dolby audio, and multiple caption standards. You can use channel presets to get started quickly, or adjust settings to precisely match your encoding requirements.
- **Highly available** - AWS Elemental MediaLive provides built-in reliability and resiliency. The service transparently manages resources across multiple Availability Zones, and automatically monitors their health, so that any potential issues are detected and resolved without disrupting live channels. With AWS Elemental MediaLive, you can exceed the reliability of infrastructure typically used for broadcast workloads with a straightforward pay-as-you-go model based on the hours of content processed.
- **Increased efficiency and reduced cost** - With AWS Elemental MediaLive, you only pay for the service as you use it, with no upfront investment in encoding infrastructure and no operational overhead devoted to managing physical resources. Pricing for AWS Elemental MediaLive uses a straightforward model based on the hours of content processed and delivered.

68.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elemental MediaLive. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/medialive/resources/>.

68.3 Service Constraints

Please see <https://aws.amazon.com/medialive/resources/> for more information.

68.4 Technical Requirements

- **User Guide** – Describes the components and features that AWS Elemental MediaLive provides and how to use them.

[HTML](#) | [PDF](#)

- **API Reference** – Describes all the AWS Elemental MediaLive API operations. Also provides sample requests, responses, and errors for the supported web services protocols.

[HTML](#) | [PDF](#)

69.0 Service Definition – AWS Elemental MediaPackage

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

69.1 Service Overview

AWS Elemental MediaPackage reliably prepares and protects your video for delivery over the Internet. From a single video input, AWS Elemental MediaPackage creates video streams formatted to play on connected TVs, mobile phones, computers, tablets, and game consoles. It makes it easy to implement popular video features for viewers (start-over, pause, rewind, etc.), like those commonly found on DVRs. AWS Elemental MediaPackage can also protect your content using Digital Rights Management (DRM). AWS Elemental MediaPackage scales automatically in response to load, so your viewers will always get a great experience without you having to accurately predict in advance the capacity you'll need.

Top benefits include:

- **Reach a wide range of connected devices** - AWS Elemental MediaPackage makes it easy to package and distribute your content to a broad range of video playback devices, including web players, smart phones, game consoles, tablets, and connected TVs.

- **Advanced video experiences and content protection** - AWS Elemental MediaPackage lets you configure a DVR-like experience for viewers of your live stream. It offers support for more than 20 digital rights management (DRM) providers, supports advanced audio features, and multi-language subtitle tracks.
- **Built-in scalability and reliability** - AWS Elemental MediaPackage scales automatically as the audience for your video grows. It automatically manages resources across multiple Availability Zones, and monitors their health, so that any potential issues are detected and resolved without disrupting your live video stream.
- **Easy integration with AWS cloud services** - AWS Elemental MediaPackage is built to work with Amazon CloudFront CDN for global distribution and with AWS Elemental MediaLive for live encoding to form a complete solution for live video processing and delivery. Integration with Amazon CloudWatch gives you real-time monitoring and notifications.

69.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elemental MediaPackage. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

69.3 Service Constraints

Please see <https://aws.amazon.com/mediapackage/resources/> for more information.

69.4 Technical Requirements

- **User Guide** - Describes the components and features that AWS Elemental MediaPackage provides and how to use them.
[HTML](#) | [PDF](#)
- **API Reference** - Describes all the AWS Elemental MediaPackage API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

70.0 Service Definition – AWS Elemental MediaStore

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

70.1 Service Overview

AWS Elemental MediaStore is an AWS storage service optimized for media. It gives you the performance, consistency, and low latency required to deliver live streaming video content. AWS Elemental MediaStore acts as the origin store in your video workflow. Its high performance capabilities meet the needs of the most demanding media delivery workloads, combined with long-term, cost-effective storage.

Top benefits include:

- **High performance, optimized for video** - AWS Elemental Media Store is optimized to deliver performance to meet the unique requirements of high-scale, high-quality media workloads (delivering low-latency reads and writes concurrently). This means you can deliver consistent quality-of-service to your viewers, lowering the risk of buffering video and reducing end-to-end latency.
- **Scale with your audience** - AWS Elemental MediaStore scales automatically with the volume of requests you receive. Increases in load don't negatively impact the quality of the viewing experience. This automatic scalability eliminates the expense and complexity of pre-provisioning capacity. If your audience size goes down, there is no need to de-scale - the service automatically sizes to varied capacity while you pay only for what you use.
- **Familiar management tools for access control** - AWS Elemental MediaStore integrates with AWS features for access control, using AWS Identity and Access Management (IAM) policies and roles, with support for resource policies, allowing you to specify granular access controls.

70.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elemental MediaStore. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/mediastore/resources/>.

70.3 Service Constraints

Please see <https://aws.amazon.com/mediastore/resources/> for more information.

70.4 Technical Requirements

- **User Guide** - Describes the components and features that AWS Elemental MediaStore provides and how to use them.
- [HTML | PDF](#)

- **API Reference** – Describes all the AWS Elemental MediaStore API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
- [HTML | PDF](#)

71.0 Service Definition – AWS Elemental MediaTailor

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

71.1 Service Overview

AWS Elemental MediaTailor lets video providers insert individually targeted advertising into their video streams without sacrificing broadcast-level quality-of-service. With AWS Elemental MediaTailor, viewers of your live or on-demand video each receive a stream that combines your content with ads personalized to them. But unlike other personalized ad solutions, with AWS Elemental MediaTailor your entire stream – video and ads – is delivered with broadcast-grade video quality to improve the experience for your viewers. AWS Elemental MediaTailor delivers automated reporting based on both client and server-side ad delivery metrics, making it easy to accurately measure ad impressions and viewer behavior. You can easily monetize unexpected high-demand viewing events with no up-front costs using AWS Elemental MediaTailor. It also improves ad delivery rates, helping you make more money from every video, and it works with a wider variety of content delivery networks, ad decision servers, and client devices.

Top benefits include:

- **Easily deliver targeted ads to any platform** - AWS Elemental MediaTailor makes it easy to personalize ads for users across different devices from a single ad monetization service. Targeting criteria can be set up to fetch personalized ad content from the ad decision server which makes the determination about which ads to send based on viewer, time, and content. Support for popular ad insertion protocols gives you the flexibility to serve individual ads or groups of ads based on your specific business needs.
- **Improve viewing experiences** - Inserting advertising into content streams as viewers play them allows you to match the quality and format of your advertising to your video content stream. Unlike other approaches to ad insertion that can cause ads to play with variable quality, AWS Elemental MediaTailor ensures the video quality of both your ads and of your core

- content are the same. This gives your viewers an experience that is superior to other streaming video ad solutions.
- **Increase the accuracy of ad view reporting** - AWS Elemental MediaTailor lets you increase the accuracy of ad tracking by capturing data directly from the viewing device, reducing the effects of ad blocking software and adhering to established advertising industry standards. Other approaches often fail to accurately capture viewing behavior due to ad blocking software, lack of support for certain devices, or dependence on inadequate third-party tools for tracking. With AWS Elemental MediaTailor, ad impressions and viewing behavior are accurately measured across web, iOS, Android, and other connected viewing devices, helping you more effectively measure the revenue impact of every ad delivered.

71.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Elemental MediaTailor. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/mediatailor/resources/>.

71.3 Service Constraints

Please see <https://aws.amazon.com/mediastore/resources/> for more information.

71.4 Technical Requirements

- **User Guide** – Describes the components and features that AWS Elemental MediaTailor provides and how to use them.
- [HTML | PDF](#)

72.0 Service Definition – Amazon Elasticsearch Service

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

72.1 Service Overview

Amazon Elasticsearch Service is a managed service that makes it easy to deploy, operate, and scale Elasticsearch in the AWS Cloud. Elasticsearch is a popular open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics. You can set up and configure your Amazon Elasticsearch cluster in minutes from the AWS Management Console. Amazon

Elasticsearch Service provisions all the resources for your cluster and launches it. The service automatically detects and replaces failed Elasticsearch nodes, reducing the overhead associated with self-managed infrastructure and Elasticsearch software. Amazon Elasticsearch Service allows you to easily scale your cluster via a single API call or a few clicks in the AWS Management Console. With Amazon Elasticsearch Service, you get direct access to the Elasticsearch open-source API so that code and applications you're already using with your existing Elasticsearch environments will work seamlessly.

Top benefits include:

- **Simple to Deploy** – Amazon Elasticsearch Service makes it easy to deploy Elasticsearch in the AWS Cloud. Use the AWS Management Console or simple API calls to access a production-ready Amazon Elasticsearch cluster in minutes without worrying about infrastructure provisioning or installing and maintaining Elasticsearch software.
- **Easy to Administer** – Amazon Elasticsearch Service simplifies time-consuming management tasks—such as ensuring high availability, patch management, failure detection and node replacement, backups, and monitoring—allowing you to pursue higher-value application development.
- **Scalable** – Amazon Elasticsearch Service enables you to monitor your cluster through Amazon CloudWatch metrics and resize your cluster up or down via a single API call or a few clicks on the AWS Management Console.
- **Integrated with Logstash or Kibana** – Logstash is an open-source data pipeline that helps you process logs and other event data and load them into Elasticsearch. Kibana is an open-source analytics and visualization platform that helps you get a better understanding of your data in Elasticsearch. Amazon Elasticsearch Service includes built-in integration with Kibana and supports integration with Logstash.
- **Cost Effective** – Amazon Elasticsearch Service saves you the administrative costs of setting up and managing Elasticsearch. You can scale up and scale down your cluster to deliver optimum performance as data and usage patterns change, paying only for the resources you actually consume. The on-demand pricing allows you to pay for resources by the hour with no long-term commitments and frees you from the costs and complexities of planning, purchasing, and maintaining hardware.
- **Secure** – You can control access to the Elasticsearch APIs using AWS IAM policies. Using IAM policies you can allow your applications to access your Amazon Elasticsearch clusters securely.

72.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Elasticsearch Service. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

72.3 Service Constraints

Please see <https://aws.amazon.com/elasticsearch-service/> for more information.

72.4 Technical Requirements

Amazon Elasticsearch Service is a managed service that makes it easy to deploy, operate, and scale Elasticsearch, a popular open-source search and analytics engine. Amazon Elasticsearch Service also offers security options, high availability, data durability, and direct access to the Elasticsearch API.

- **Developer Guide** – Describes how to get started with Amazon Elasticsearch Service, explains key concepts, and provides step-by-step instructions that show you how to use the features.

[HTML](#) | [PDF](#)

73.0 Service Definition – AWS Data Pipeline

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

73.1 Service Overview

AWS Data Pipeline is a web service that helps you reliably process and move data between different AWS compute and storage services, as well as on-premises data sources, at specified intervals. With AWS Data Pipeline, you can regularly access your data where it's stored, transform and process it at scale, and efficiently transfer the results to AWS Cloud services such as Amazon S3, Amazon RDS, Amazon DynamoDB, and Amazon EMR.

AWS Data Pipeline helps you easily create complex data processing workloads that are fault tolerant, repeatable, and highly available. You don't have to worry about ensuring resource availability, managing inter-task dependencies, retrying transient failures or timeouts in individual tasks, or creating a failure notification system. AWS Data Pipeline also allows you to move and process data that was previously locked up in on-premises data silos.

Top benefits include:

- **Reliable** – AWS Data Pipeline is built on a distributed, highly available infrastructure designed for fault-tolerant execution of your activities. If failures occur in your activity logic or data sources, AWS Data Pipeline automatically retries the activity. If the failure persists, AWS Data Pipeline sends you failure notifications via [Amazon SNS](#). You can configure your notifications for successful runs, delays in planned activities, or failures.
- **Simple** – Creating a pipeline is quick and easy via our drag-and-drop console. Common preconditions are built into the service, so you don't need to write any extra logic to use them. For example, you can check for the existence of an Amazon S3 file by simply providing the name of the Amazon S3 bucket and the path of the file that you want to check for, and AWS Data Pipeline does the rest. In addition to its easy visual pipeline creator, AWS Data Pipeline provides a library of pipeline templates. These templates make it simple to create pipelines for a number of more complex use cases, such as regularly processing your log files, archiving data to Amazon S3, or running periodic SQL queries.
- **Flexible** – AWS Data Pipeline allows you to take advantage of a variety of features such as scheduling, dependency tracking, and error handling. You can use activities and preconditions that AWS provides and/or write your own custom ones. This means that you can configure AWS Data Pipeline to take actions like run Amazon EMR jobs, execute SQL queries directly against databases, or execute custom applications running on Amazon EC2 or in your own data centre. This allows you to create powerful custom pipelines to analyse and process your data without having to deal with the complexities of reliably scheduling and executing your application logic.
- **Scalable** – AWS Data Pipeline makes it equally easy to dispatch work to one machine or many, in serial or parallel. With AWS Data Pipeline's flexible design, processing a million files is as easy as processing a single file.
- **Low Cost** – AWS Data Pipeline is inexpensive to use and is billed at a low monthly rate. You can try it for free under the AWS Free Usage Tier.
- **Transparent** – You have full control over the computational resources that execute your business logic, making it easy to enhance or debug your logic. Additionally, full execution logs are automatically delivered to Amazon S3, giving you a persistent, detailed record of what has happened in your pipeline.

73.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Data Pipeline. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

73.3 Service Constraints

Please see <https://aws.amazon.com/datapipeline/> for more information.

73.4 Technical Requirements

AWS Data Pipeline is a web service that you can use to automate the movement and transformation of data. With AWS Data Pipeline, you can define data-driven workflows, so that tasks can be dependent on the successful completion of previous tasks.

- **Developer Guide** – Provides a conceptual overview of AWS Data Pipeline and includes detailed development instructions for using the various features.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS Data Pipeline in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

74.0 Service Definition – AWS Glue

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

74.1 Service Overview

AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy for customers to prepare and load their data for analytics. You can create and run an ETL job with a few clicks in the AWS Management Console. You simply point AWS Glue to your data stored on AWS, and AWS Glue discovers your data and stores the associated metadata (e.g. table definition and schema) in the AWS Glue Data Catalog. Once cataloged, your data is immediately searchable, queryable, and available for ETL. AWS Glue generates the code to execute your data transformations and data loading processes.

AWS Glue generates code that is customizable, reusable, and portable. Once your ETL job is ready, you can schedule it to run on AWS Glue's fully managed, scale-out Apache Spark environment. AWS Glue provides a flexible scheduler with dependency resolution, job monitoring, and alerting.

AWS Glue is serverless, so there is no infrastructure to buy, set up, or manage. It automatically provisions the environment needed to complete the job, and customers pay only for the compute resources consumed while running ETL jobs. With AWS Glue, data can be available for analytics in minutes.

74.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Glue. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

74.3 Service Constraints

Please see <https://aws.amazon.com/documentation/glue/> for more information.

74.4 Technical Requirements

AWS Glue is a fully managed ETL (extract, transform, and load) service that makes it simple and cost-effective to categorize your data, clean it, enrich it, and move it reliably between various data stores. AWS Glue consists of a central data repository known as the AWS Glue Data Catalog, an ETL engine that automatically generates Python code, and a flexible scheduler that handles dependency resolution, job monitoring, and retries. AWS Glue is serverless, so there's no infrastructure to set up or manage. Use the AWS Glue console to discover your data, transform it, and make it available for search and querying. You can also use the AWS Glue API operations to interface with AWS Glue.

- **Developer Guide** – Provides a conceptual overview of AWS Glue, detailed instructions for using the various features, and a complete API reference for developers.

[HTML](#) | [PDF](#) | [Kindle](#)

75.0 Service Definition – Amazon Simple Email Service (Amazon SES)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

75.1 Service Overview

Amazon SES is a cost-effective email service built on the reliable and scalable infrastructure that Amazon.com developed to serve its own customer base. With Amazon SES, you can send transactional email, marketing messages, or any other type of high-quality content to your customers. You can also use Amazon SES to receive messages and deliver them to an Amazon S3 bucket, call your custom code via an AWS Lambda function, or publish notifications to Amazon SNS. With Amazon SES, you have no required minimum commitments—you pay as you go, and you only pay for what you use.

Top benefits include:

- **Inexpensive** – There are no up-front fees or fixed expenses with Amazon SES, and you benefit from the efficiencies of Amazon's scale. You pay low charges for the number of emails sent, number of emails received, data transfer fees, and attachments.
- **Reliable** – Amazon SES runs within Amazon's cloud network infrastructure and data centres. Multiple servers and data centres provide high availability and data durability.
- **Scalable** – Amazon SES is based on the scalable cloud-based email technology used by Amazon websites around the world to send billions of messages a year.
- **Highly Deliverable** – Amazon SES takes proactive steps to maximise the percentage of your emails that arrive in your recipients' inboxes.
- **Highly Configurable** – When you use Amazon SES to receive your email, you have fine-grained control over what happens to your incoming messages. You can set up rules that apply to your entire domain or to specific email addresses.
- **Designed for Use with Other AWS Cloud Services** – Amazon SES integrates with [Amazon SNS](#), [Amazon Route 53](#), [Amazon S3](#), [Amazon WorkMail](#), [AWS Lambda](#), [AWS KMS](#), and [Amazon CloudWatch](#). Additionally, emails you send from [Amazon EC2](#) and [AWS Elastic Beanstalk](#) qualify for the free usage tier.

75.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon SES. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

75.3 Service Constraints

Please see <https://aws.amazon.com/ses/> for more information.

75.4 Technical Requirements

Amazon SES is an email sending and receiving service that provides an easy, cost-effective way for you to send email.

- **Developer Guide** – Provides a conceptual overview of Amazon SES and includes detailed instructions for using its various features.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the Amazon SES API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

76.0 Service Definition – Amazon Simple Queue Service (Amazon SQS)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

76.1 Service Overview

Amazon SQS is a fast, reliable, scalable, fully managed message queuing service. Amazon SQS makes it simple and cost effective to decouple the components of a cloud application. You can use Amazon SQS to transmit any volume of data, at any level of throughput, without losing messages or requiring other services to be always available. With Amazon SQS, you can offload the administrative burden of operating and scaling a highly available messaging cluster, while paying a low price for only what you use.

Top benefits include:

- **Reliable** – Amazon SQS runs within Amazon’s high-availability data centres, so queues will be available whenever applications need them. To prevent messages from being lost or becoming unavailable, all messages are stored redundantly across multiple servers and data centres.
- **Simple** – Developers can get started with Amazon SQS by using only three APIs: SendMessage, ReceiveMessage, and DeleteMessage. Additional APIs are available to provide advanced functionality.
- **Scalable** – Amazon SQS was designed to enable an unlimited number of messaging services to read and write an unlimited number of messages at any time.
- **Secure** – Authentication mechanisms are provided to ensure that messages stored in Amazon SQS queues are secured against unauthorised access.
- **Inexpensive** – No up-front or fixed expenses. The only costs of sending messages through Amazon SQS are small per-request handling fees and data transfer fees.

76.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon SQS. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

76.3 Service Constraints

Please see <https://aws.amazon.com/sqs/> for more information.

76.4 Technical Requirements

Amazon SQS is a messaging queue service that handles message or workflows between other components in a system.

- **Getting Started Guide** – Introduces you to Amazon SQS, helps you set up an account, and walks you through a simple example to use Amazon SQS for the first time. Also provides tips and links to advanced product features and resources.

[HTML](#) | [PDF](#) | [Kindle](#)

- **Developer Guide** – Provides a conceptual overview of Amazon SQS and includes detailed development instructions for using the various features.

[HTML](#) | [PDF](#) | [Kindle](#)

- **API Reference** – Describes all the API operations for Amazon SQS in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

[HTML](#) | [PDF](#)

- **Amazon SQS section of AWS CLI Reference** – Describes the AWS CLI commands that you can use to automate queues.

[HTML](#)

77.0 Service Definition – AWS AppSync

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

77.1 Service Overview

AWS AppSync automatically updates the data in web and mobile applications in real time, and updates data for offline users as soon as they reconnect. AWS AppSync makes it easy to build collaborative mobile and web applications that deliver responsive, collaborative user experiences.

You can use AWS AppSync to build native mobile and web apps with iOS, Android, JavaScript and React Native. Get started by going to the AWS AppSync console, specify the data for your app with simple code statements, and AWS AppSync will manage everything needed to store, process, and retrieve the data for your application.

Top benefits include:

- **Rapid prototyping and development with GraphQL** - With AWS AppSync you can specify the data requirements of the application with simple code statements and iterate quickly during the prototyping and development process. AWS AppSync uses GraphQL, an open standard query language that makes it easy for applications to request data from the cloud.
- **Build real-time, collaborative mobile and web apps** - AWS AppSync makes it easy to build collaborative apps that keep shared data updated in real time. You just specify the data for your app with simple code statements and AWS AppSync manages everything needed to keep the app data updated in real time. AWS AppSync also automatically updates the data for offline users as soon as they reconnect. Only the data needed is fetched, reducing message traffic and improving application performance and battery life.
- **Automatically manages offline users' data** - AWS AppSync automatically manages all the data operations for offline users. The service supports an offline programming model where application data is not only available for users who are offline, but users can also add and update app data locally as well. This makes it easy to build apps that cache important data locally for offline use, and then synchronize with the cloud when the device reconnects.
- **Helps keep data secure** - AWS AppSync helps you keep users and app data secure. The service integrates with Amazon Cognito and AWS Identity and Access Management, so you can set fine-grained permissions on GraphQL operations that put strict controls on who can access the data.
- **Build apps with data from multiple sources** - AWS AppSync makes it easy to combine data from different sources. For example, you could build an app that uses AppSync to combine a user's activity data from their phone and web usage, fitness tracker activity, and social media postings and combine it with location data to summarize all the user's activities and locations for the day. With AppSync, your app can access data in Amazon DynamoDB, trigger AWS Lambda functions, or run Amazon Elasticsearch queries and combine data from these services to provide the exact data you need for your app.
- **Automatically resolves data conflicts** - AWS AppSync can automatically resolve conflicts on data that is shared between multiple users and devices. Conflicts can happen when two or more users make updates at the same time. The service has default rules to detect conflicts and ensure they are resolved, but also allows you to add your own business logic into the conflict resolution process with AWS Lambda functions.

77.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS AppSync. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

77.3 Service Constraints

Please see <https://aws.amazon.com/appsync/> for more information.

77.4 Technical Requirements

AWS AppSync is an enterprise level, fully managed GraphQL service with real-time data synchronization and offline programming features.

- **Developer Guide** - Create and interact with data sources using GraphQL from your application. You can build a new application or integrate existing data sources with AWS AppSync.

[HTML](#) | [PDF](#)

- **API Reference** – Describes all the API operations for AWS AppSync in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

[HTML](#) | [PDF](#)

78.0 Service Definition – Amazon Simple Workflow Service (Amazon SWF)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

78.1 Service Overview

Amazon SWF helps developers build, run, and scale background jobs that have parallel or sequential steps. You can think of Amazon SWF as a fully managed state tracker and task coordinator in the cloud. If your application's steps take more than 500 milliseconds to complete, you need to track the state of processing, and you need to recover or retry if a task fails. Amazon SWF can help you.

Top benefits include:

- **Logical Separation** – Amazon SWF promotes a separation between the control flow of your background job’s stepwise logic and the actual units of work that contain your unique business logic. This allows you to separately manage, maintain, and scale “state machinery” of your application from the core business logic that differentiates it. As your business requirements change, you can easily change application logic without having to worry about the underlying state machinery, task dispatch, and flow control.
- **Reliable** – Amazon SWF runs within Amazon’s high-availability data centres, so the state tracking and task processing engine is available whenever applications need them. Amazon SWF redundantly stores the tasks, reliably dispatches them to application components, tracks their progress, and keeps their latest state.
- **Simple** – Amazon SWF replaces the complexity of custom-coded workflow solutions and process automation software with a fully managed cloud workflow web service. This eliminates the need for developers to manage the infrastructure plumbing of process automation so they can focus their energy on the unique functionality of their application.
- **Scalable** – Amazon SWF seamlessly scales with your application’s usage. No manual administration of the workflow service is required as you add more cloud workflows to your application or increase the complexity of your workflows.
- **Flexible** – Amazon SWF lets you write your application components and coordination logic in any programming language and run them in the cloud or on-premises.

78.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon SWF. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

78.3 Service Constraints

Please see <https://aws.amazon.com/swf/> for more information.

78.4 Technical Requirements

Amazon SWF makes it easy to build applications that coordinate work across distributed components. In Amazon SWF, a task represents a logical unit of work that is performed by a component of your application. Coordinating tasks across the application involves managing intertask dependencies, scheduling, and concurrency in accordance with the logical flow of the application. Amazon SWF gives you full control over implementing tasks and coordinating them without worrying about underlying complexities such as tracking their progress and maintaining their state.

- **Developer Guide** – Provides a conceptual overview of Amazon SWF and includes detailed development instructions for using the various features.

[HTML](#) | [PDF](#) | [Kindle](#)

- **API Reference** – Describes all the API operations for Amazon SWF in detail. Also provides sample requests, responses, and errors for the supported web services protocols.

[HTML](#) | [PDF](#)

79.0 Service Definition – Alexa for Business

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

79.1 Service Overview

Alexa for Business allows organizations of all sizes to introduce Alexa to their workplace. With Alexa for Business, customers can use the Alexa they know as an intelligent assistant to stay organized and focus on the work that matters. Alexa helps workers be more productive as they move throughout their day, at home and at their desks as enrolled users with personal devices, and in conference rooms, copy rooms or other shared spaces with shared devices. Alexa for Business includes the tools and controls that administrators need to deploy and manage shared Alexa devices, skills, and users at scale.

Top benefits include:

- **Easily Provision and Manage Alexa Devices** - Alexa for Business allows you to provision and manage Alexa devices in your organization from a centralized console. With Alexa for Business, you can easily provision multiple Alexa devices at the same time, and automatically connect them to your Alexa for Business account. You can specify device locations, enable a set of skills that can be used, and prevent users from tampering with them. This saves time because you don't need to manage these devices individually.
- **Configure Conference Rooms** - Alexa for Business makes it easy for you to configure Alexa to control your conference rooms. Alexa for Business lets you specify the type of conferencing equipment you use and your preferred meeting applications, which allows Alexa to start most meetings, on most devices, in any room. You can use Alexa devices as audio conferencing devices in small conference rooms, or to control equipment in larger rooms. Alexa for Business is an open service, and the Alexa for Business APIs allow you to build skills so that Alexa can work with additional equipment or perform specific tasks in your conference rooms.
- **Manage Users** - Alexa for Business allows you to invite your end users to enroll their personal Alexa account with your Alexa for Business account. This lets them continue to use the Alexa features and skills they've already enabled in their personal Alexa account, as well as the work skills you provide, on any of their devices, at work or at home. Alexa for Business gives you the ability to make work skills available and provide access to your corporate calendar system so that they can use Alexa to manage their calendar.

- **Create Custom Skills** - Alexa for Business lets you build your own private custom skills for your workplace, your employees, or your customers to use. You can make these skills available only to your shared Alexa devices, and your enrolled users. Alexa for Business provides an additional set of APIs that provide information about device location, which lets you add context to your skills. For example, you could build a skill that lets a user report a printer problem to IT, and the skill could use the device location so that IT knows which printer is broken. Building custom skills is easy, and the Alexa Skills Kit provides tools, documentation, and code samples to help you get started.

79.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Alexa for Business. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

79.3 Service Constraints

Please see <https://aws.amazon.com/alexaforbusiness/> for more information.

79.4 Technical Requirements

Alexa for Business gives you the tools you need to manage Alexa devices, enroll users, and assign skills. You can build your own voice skills using the Alexa Skills Kit and the Alexa for Business API. You can also make them available as private skills for your organization.

- **Administration Guide** - Helps you perform several administrative tasks for Alexa for Business, such as getting set up and managing your devices, rooms, skills, and users.

[HTML | PDF](#).

- **API Reference** – Describes the API operations for Alexa for Business.
[HTML | PDF](#)

80.0 Service Definition – Amazon AppStream 2.0

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

80.1 Service Overview

Amazon AppStream 2.0 is a fully managed, secure application streaming service that allows customers to stream desktop applications from AWS to any device, without rewriting them. AppStream 2.0 provides users with instant-on access to the applications they need, and a responsive, fluid user experience running in an HTML5 web browser.

Top benefits include:

- **Run Desktop Applications in a Web Browser** – With Amazon AppStream 2.0, your desktop applications can run securely in an HTML 5 compatible web browser on Windows and Linux PCs, Macs, and Chromebooks. You can import your productivity and graphics applications without rewriting them, maintain a single version for all your users, and provide easy access to your users from anywhere.
- **Instant-On Access** - Amazon AppStream 2.0 can provide users instant-on access to their desktop applications in a browser on the desktop device of their choice. There are no delays, no large files to download, and no time-consuming installations. Users get a responsive, fluid experience that is indistinguishable from natively installed apps.
- **Secure Applications and Data** - With Amazon AppStream 2.0, applications and data remain on AWS – only encrypted pixels are streamed to end users. Applications run on a streaming instance dedicated to each user so that compute resources are not shared. Applications can run inside your own VPC, and you can use Amazon VPC security features to control access. This allows you to isolate your applications and deliver them in a secure way.
- **Easily integrate with your IT Environment** - Amazon AppStream 2.0 can integrate with your existing AWS services, and your on-premises environments. By running applications inside your VPCs, your users can access data and other resources that you're running on AWS, reducing the movement of data between AWS and your location and providing a faster user experience. Amazon AppStream 2.0 supports identity federation, which allows your users to access their applications using their existing credentials. You can also allow authenticated access to your IT resources from applications running on Amazon AppStream 2.0.
- **Fully Managed Service** - With Amazon AppStream 2.0, you don't need to plan, deploy, manage, or upgrade any application streaming infrastructure. Amazon AppStream 2.0 manages the AWS resources required to host and run your applications, scales automatically, and provides access to your end users on-demand.
- **Consistent, Scalable Performance** - Amazon AppStream 2.0 runs on AWS with access to compute capabilities not available on local devices, which means that your applications run with consistently high performance. You can instantly scale locally and globally, and ensure that your users always get a low-latency experience. Unlike on-premises solutions, you can quickly deploy your

applications to a new AWS region that is closest to your users, and start streaming with no incremental capital investment.

80.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon AppStream 2.0. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

80.3 Service Constraints

Please see <https://aws.amazon.com/appstream2/> for more information.

80.4 Technical Requirements

The Amazon AppStream web service deploys your application on AWS infrastructure and streams input and output between your application and devices such as personal computers, tablets, and mobile phones. Your application's processing occurs in the cloud, so it can scale to handle vast computational loads. Devices need only display output and return user input, so the client application on the device can be lightweight in terms of file size and processing requirements.

- **Developer Guide** – Helps you get started using Amazon AppStream 2.0. You will learn how to quickly and easily provision and maintain one or more WorkSpaces.

[HTML](#) | [PDF](#)

- **API Reference Guide** - Describes the API operations for Amazon AppStream 2.0.

[HTML](#) | [PDF](#)

81.0 Service Definition – AWS Greengrass

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

81.1 Service Overview

AWS Greengrass is software that lets you run local compute, messaging, data caching, sync, and ML inference capabilities for connected devices in a secure way. With AWS

Greengrass, connected devices can run [AWS Lambda](#) functions, keep device data in sync, and communicate with other devices securely – even when not connected to the Internet. Using AWS Lambda, Greengrass ensures your IoT devices can respond quickly to local events, use Lambda functions running on Greengrass Core to interact with local resources, operate with intermittent connections, stay updated with over the air updates, and minimize the cost of transmitting IoT data to the cloud.

[ML Inference](#) is a feature of AWS Greengrass that makes it easy to perform machine learning inference locally on Greengrass Core devices using models that are built and trained in the cloud.

AWS Greengrass seamlessly extends AWS to devices so they can act locally on the data they generate, while still using the cloud for management, analytics, and durable storage. With Greengrass, you can use familiar languages and programming models to create and test your device software in the cloud, and then deploy it to your devices. AWS Greengrass can be programmed to filter device data and only transmit necessary information back to the cloud. AWS Greengrass authenticates and encrypts device data at all points of connection using the security and access management capabilities of [AWS IoT Core](#). This way, data is never exchanged between devices when they communicate with each other and the cloud, without proven identity.

Top benefits include:

- **Respond to Local Events in Near Real-time** - AWS Greengrass devices can act locally on the data they generate, while still using the cloud for management, analytics, and durable storage.
- **Operate Offline** - AWS Greengrass lets connected devices operate even with intermittent connectivity to the cloud. Once the device reconnects, Greengrass synchronizes the data on the device with [AWS IoT](#), providing seamless functionality regardless of connectivity.
- **Simplified Device Programming with AWS Lambda** - AWS Greengrass uses the same AWS Lambda programming models you use in the cloud so you can create and test your device software in the cloud first, and then deploy it seamlessly to your devices. Greengrass lets you execute Lambda functions locally, reducing the complexity of developing embedded software.
- **Reduce the Cost of Running IoT Applications** - With AWS Greengrass you can program the device to filter device data locally and only transmit the data you need for your applications to cloud. This reduces the amount of raw data transmitted to the cloud and lowers cost, and increases the quality of the data you send to the cloud so you can achieve rich insight at a lower cost.

81.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Greengrass. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

81.3 Service Constraints

Please see <https://aws.amazon.com/greengrass/> for more information.

81.4 Technical Requirements

AWS Greengrass seamlessly extends AWS onto physical devices so they can act locally on the data they generate, while still using the cloud for management, analytics, and durable storage. AWS Greengrass ensures your devices can respond quickly to local events and operate with intermittent connectivity. AWS Greengrass minimizes the cost of transmitting data to the cloud by allowing you to author AWS Lambda functions that execute locally.

- **Developer Guide** - Walks through how to set up AWS Greengrass and integrate it with other services.
[HTML](#) | [PDF](#)
- **AWS Greengrass sections of the AWS CLI Reference** - Describes the AWS CLI commands that you can use to administer and manipulate AWS Greengrass.
[Greengrass](#)
- **API Reference** - Describes all the API operations for AWS Greengrass in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

82.0 Service Definition – Amazon Cognito

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

82.1 Service Overview

Amazon Cognito lets you easily add user sign-up and sign-in to your mobile and web apps. With Amazon Cognito, you can also authenticate users through social identity

providers such as Facebook, Twitter, or Amazon, or by **using your** own identity solution. In addition, Amazon Cognito enables you to save data locally on users' devices, allowing your applications to work even when the devices are offline. You can then synchronise data across users' devices so that their app experience remains consistent regardless of the device they use.

With Amazon Cognito, you can focus on creating great app experiences instead of worrying about building, securing, and scaling a solution to handle user management, authentication, and syncing across devices.

Top benefits include:

- **Add User Sign-Up and Sign-In to Your Mobile Web Apps** – With Amazon Cognito Identity, you can easily and securely add sign-up and sign-in functionality to your mobile and web apps. Amazon Cognito Identity is fully managed and can scale to support hundreds of millions of users. You can also implement enhanced security features, such as email and phone number verification and multi-factor authentication. In addition, Amazon Cognito Identity lets you customise workflows through AWS Lambda, such as adding app-specific logic to user registration for fraud detection and user validation.
- **Federate Identities and Provide Secure Access to AWS Resources** – Amazon Cognito Identity allows you to easily federate identities so that your users can sign in through social identity providers such as Facebook, Twitter, Google, and Amazon or using your own identity solution. When you use Amazon Cognito, your app is provided with temporary, limited-privilege credentials that it can use to access AWS resources. You can set granular access permissions on your AWS resources. For example, you can limit access to a folder within an Amazon S3 bucket to a particular app user or enable unauthenticated users to access a restricted set of resources. This means your app can access the resources it needs and that you can follow security best practices by not hardcoding credentials in your app. To learn more, visit our docs.
- **Store and Sync Across Devices** – When you use Amazon Cognito Sync, your app can read and write to a local data store. This means that your app can work in the same way regardless of whether the device is online or offline. You can also save user data, such as user preferences, sign-in, and game state, and then sync this data across a user's devices to make their experience consistent across their devices. Amazon Cognito Sync resolves conflicts and intermittent network connectivity, enabling you to focus on delivering great app experiences instead of creating and managing a user data sync solution.

82.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Cognito. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

82.3 Service Constraints

Please see <https://aws.amazon.com/cognito/> for more information.

82.4 Technical Requirements

Amazon Cognito lets you easily add user sign-up and sign-in to your mobile and web apps. With Amazon Cognito, you can also authenticate users through social identity providers such as Facebook, Twitter, or Amazon, or by using your own identity solution.

- **Developer Guide** – Describes how to get started with AWS WAF and AWS Shield Advanced, explains key concepts, and provides step-by-step instructions that show you how to use the features
[HTML](#)

83.0 Service Definition – AWS Device Farm

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

83.1 Service Overview

Improve the quality of your iOS, Android, and web applications by testing against real phones and tablets in the AWS Cloud.

Top benefits include:

- **Use the Same Devices Your Customers Use** – Run tests and interact with a large selection of physical devices. Unlike emulators, physical devices give you a more accurate understanding of the way users interact with your app by taking into account factors like memory, CPU usage, location, and modifications made by manufacturers and carriers to the firmware and software. We are always adding devices to the fleet.
- **Simulate Real-World Environments** – Fine-tune your test environment by configuring location, language, application data, and installing prerequisite apps to simulate real-world customer conditions.
- **Integrate with Your Development Workflow** – Use our service plugins and API to automatically initiate tests and get results from Integrated Development Environments (IDEs) and continuous integration environments like Android Studio and Jenkins.
- **Reproduce and Fix Issues Faster** – Manually reproduce issues and run tests in parallel. We collect videos, logs, and performance data so you can dive deep and solve problems quickly. For automated tests, we'll identify and group issues so you can focus on the most important problems first.
- **Choose the Tests That Work for You** – Run our built-in test suite (no scripting required) or customise your tests by selecting from open-source test frameworks like Appium, Calabash, and Espresso ([see supported frameworks](#)). You can also perform manual tests with Remote Access.
- **Test and Reproduce Issues with Confidence** – Each device is dedicated to you for the duration of your test or Remote Access session and completely isolated from other devices at the software and hardware level. After you are done, the device is wiped to remove your application and data.

83.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Device Farm. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

83.3 Service Constraints

Please see <https://aws.amazon.com/device-farm/> for more information.

83.4 Technical Requirements

AWS Device Farm is an app testing service that enables you to test your iOS, Android, and Fire OS apps on real, physical phones and tablets that are hosted by AWS. The service allows you to upload your own tests or use built-in, script-free compatibility tests.

- **Developer Guide** – Provides a conceptual overview of AWS Device Farm and includes detailed instructions for using the service.
[HTML](#) | [PDF](#)
- **API Reference** – Describes all the AWS Device Farm operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

84.0 Service Definition – AWS IoT

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

84.1 Service Overview

AWS IoT is a managed cloud platform that lets connected devices easily and securely interact with cloud applications and other devices. AWS IoT can support billions of devices and trillions of messages and can process and route those messages to AWS endpoints and to other devices reliably and securely. With AWS IoT, your applications can keep track of and communicate with all your devices, all the time, even when they aren't connected.

AWS IoT makes it easy to use AWS Cloud services like AWS Lambda, Amazon Kinesis, Amazon S3, Amazon Machine Learning, Amazon DynamoDB, Amazon CloudWatch, AWS CloudTrail, and Amazon Elasticsearch Service with built-in Kibana

integration to build IoT applications that gather, process, analyse and act on data generated by connected devices, without having to manage any infrastructure.

Top benefits include:

- **Connect and Manage Your Devices** – AWS IoT allows you to easily connect devices to the cloud and to other devices. AWS IoT supports HTTP, WebSockets, and MQTT, a lightweight communication protocol specifically designed to tolerate intermittent connections, minimise the code footprint on devices, and reduce network bandwidth requirements. AWS IoT also supports other industry-standard and custom protocols, and devices can communicate with each other even if they are using different protocols.
- **Secure Device Connections and Data** – AWS IoT provides authentication and end-to-end encryption throughout all points of connection, so that data is never exchanged between devices and AWS IoT without proven identity. In addition, you can secure access to your devices and applications by applying policies with granular permissions.
- **Process and Act Upon Device Data** – With AWS IoT, you can filter, transform, and act upon device data on the fly, based on business rules you define. You can update your rules to implement new device and application features at any time. AWS IoT makes it easy to use AWS Cloud services like AWS Lambda, Amazon Kinesis, Amazon S3, Amazon Machine Learning, Amazon DynamoDB, Amazon CloudWatch, and Amazon Elasticsearch Service for even more powerful IoT applications.
- **Read and Set Device State at Any Time** – AWS IoT stores the latest state of a device so that it can be read or set at anytime, making the device appear to your applications as if it were online all the time. This means that your application can read a device's state even when it is disconnected and also allows you to set a device state and have it implemented when the device reconnects.

84.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS IoT. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

84.3 Service Constraints

Please see <https://aws.amazon.com/iot/> for more information.

84.4 Technical Requirements

AWS IoT enables secure, bi-directional communication between Internet-connected things (such as sensors, actuators, embedded devices, or smart appliances) and the AWS Cloud over MQTT and HTTP.

- **Developer Guide** – Walks through how to set up AWS IoT and integrate it with other services. [HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for AWS IoT in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

- **AWS IoT sections of the AWS CLI Reference** – Describes the AWS CLI commands that you can use to administer and manipulate AWS IoT.
[IoT](#) | [IoT-Data](#)
- **Thing Shadow REST API** – Describes all of the REST API operations for working with Thing Shadows.
[HTML](#)

85.0 Service Definition – AWS IoT Core

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

85.1 Service Overview

AWS IoT Core is a managed cloud platform that lets connected devices easily and securely interact with cloud applications and other devices. AWS IoT Core can support billions of devices and trillions of messages, and can process and route those messages to AWS endpoints and to other devices reliably and securely. With AWS IoT Core, your applications can keep track of and communicate with all your devices, all the time, even when they aren't connected.

AWS IoT Core makes it easy to use AWS services like [AWS Lambda](#), [Amazon Kinesis](#), [Amazon S3](#), [Amazon Machine Learning](#), [Amazon DynamoDB](#), [Amazon CloudWatch](#), [AWS CloudTrail](#), and [Amazon Elasticsearch Service](#) with built-in [Kibana](#) integration, to build IoT applications that gather, process, analyze and act on data generated by connected devices, without having to manage any infrastructure.

Top benefits include:

- **Connect and manage your devices** - AWS IoT Core allows you to easily connect devices to the cloud and to other devices. AWS IoT Core supports HTTP, WebSockets, and MQTT, a lightweight communication protocol specifically designed to tolerate intermittent connections, minimize the code footprint on devices, and reduce network bandwidth requirements. AWS IoT Core also supports other industry-standard and custom protocols, and devices can communicate with each other even if they are using different protocols.
- **Process and act upon device data** - With AWS IoT Core, you can filter, transform, and act upon device data on the fly, based on business rules you define. You can update your rules to implement new device and application features at any time. AWS IoT Core makes it easy to use AWS services like AWS Lambda, Amazon Kinesis, Amazon S3, Amazon Machine Learning, Amazon DynamoDB, Amazon CloudWatch, and Amazon Elasticsearch Service for even more powerful IoT applications.
- **Read and set device state at any time** - AWS IoT Core stores the latest state of a device so that it can be read or set at anytime, making the device appear to your applications as if it were online all the time. This means that your application can read a device's state even when it is

disconnected, and also allows you to set a device state and have it implemented when the device reconnects.

85.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS IoT Core. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

85.3 Service Constraints

Please see <https://aws.amazon.com/iot-core/> for more information.

85.4 Technical Requirements

AWS IoT Core is a platform that enables you to connect devices to AWS Services and other devices, secure data and interactions, process and act upon device data, and enable applications to interact with devices even when they are offline.

- **Developer Guide**
[PDF Version or HTML Version](#)
- **API References**
[PDF Version or HTML Version](#)

86.0 Service Definition – Amazon FreeRTOS

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

86.1 Service Overview

Amazon FreeRTOS (a:FreeRTOS) is an operating system for microcontrollers that makes small, low-power edge devices easy to program, deploy, secure, connect, and manage. Amazon FreeRTOS is based on the FreeRTOS kernel, a popular open source operating system for microcontrollers, and extends it with software libraries that make it easy to securely connect your small, low-power devices to AWS cloud services like AWS IoT Core or to more powerful edge devices running AWS Greengrass.

A microcontroller (MCU) is a single chip containing a simple processor that can be found in many devices, including appliances, sensors, fitness trackers, industrial automation, and automobiles. Many of these small devices could benefit from

connecting to the cloud or locally to other devices. For example, smart electricity meters need to connect to the cloud to report on usage and building security systems need to communicate locally so that a door will unlock when you badge in. Microcontrollers have limited compute power and memory capacity and typically perform simple, functional tasks. Microcontrollers frequently run operating systems which do not have built in functionality to connect to local networks or the cloud, making IoT applications a challenge. Amazon FreeRTOS helps solve this problem by providing both the core operating system (to run the edge device) as well as software libraries that make it easy to securely connect to the cloud (or other edge devices) so you can collect data from them for IoT applications and take action.

Top benefits include:

- **Based on the FreeRTOS Kernel** - Amazon FreeRTOS is open source and is based on the FreeRTOS kernel, the most popular real-time operating system for microcontrollers. With Amazon FreeRTOS, you can take advantage of the large ecosystem of existing tools developed for the FreeRTOS kernel. If you are already using the FreeRTOS kernel on your existing devices, you can choose to use Amazon FreeRTOS and take advantage of the functionality provided by the additional libraries, or continue to use just the kernel operating system.
- **Easily Program, Deploy, and Manage Low Power Connected Devices** - Amazon FreeRTOS includes software libraries that make it easy to program commonly needed IoT capabilities into your device, including libraries that help configure devices to a local network using common connectivity options like Wi-Fi or Ethernet. Amazon FreeRTOS also includes an over-the-air (OTA) update feature, now in beta, to remotely update devices with feature enhancements or security patches.
- **Secure Data and Device Connections** - Amazon FreeRTOS comes with libraries to help secure device data and connections, including support for data encryption and key management. Amazon FreeRTOS includes support for Transport Layer Security (TLS v1.2) to help your devices connect securely to the cloud. Amazon FreeRTOS also has a code signing feature to ensure your device code is not compromised during deployment and OTA updates (now in beta).
- **Connect to the Cloud or to a Local Edge Device** - Amazon FreeRTOS devices can connect directly to cloud services like AWS IoT Core, or to a local edge device, like AWS Greengrass devices. By connecting Amazon FreeRTOS devices to a Greengrass Core device, you can continue to communicate with the device even if it loses connection to the cloud.
- **Broad Hardware and Technology Ecosystem** - Amazon FreeRTOS gives you the flexibility to easily build IoT solutions on a variety of chipsets and supports a variety of architectures, such as ARM and MIPS. For new devices, you can choose from a variety of qualified chipsets from trusted providers, including Espressif, Microchip, NXP Semiconductors, STMicroelectronics, Texas Instruments, and others from the Amazon FreeRTOS Qualification Program.

86.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon FreeRTOS. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

86.3 Service Constraints

Please see <https://aws.amazon.com/iot-core/> for more information.

86.4 Technical Requirements

Amazon FreeRTOS is an operating system that makes microcontroller-based edge devices easy to program, deploy, secure, and maintain. Amazon FreeRTOS extends the FreeRTOS kernel, a popular open source operating system for microcontrollers, with software libraries that can be used to securely connect devices (locally and to the cloud) and update them remotely.

- **User Guide** - Provides detailed information about the microcontroller operating system that makes small, low-powered edge devices easy to program, deploy, secure, and maintain.
[HTML](#) | [PDF](#)

Additional documentation is available at <https://aws.amazon.com/documentation/freertos/>.

87.0 Service Definition – AWS IoT 1-Click

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

87.1 Service Overview

AWS IoT 1-Click is a service that makes it easy for simple devices to trigger AWS Lambda functions that execute a specific action. Some examples of possible actions include calling technical support, reordering goods and services, or locking and unlocking doors and windows. Simple devices are cloud-connected, single-purpose devices like buttons, badge readers, asset trackers, and motion sensors. The AWS IoT Enterprise Button, based on the Amazon Dash Button hardware, is a simple device that you can buy and use with the AWS IoT 1-Click service. You can use the AWS IoT Enterprise Button to capture quick customer feedback without burdening customers with

time-consuming questionnaires. For example, in an airport lounge, patrons can press one of five buttons to rate their overall experience on a scale from one to five.

Creating Lambda triggers for any of your simple devices is easy. You just download the 1-Click mobile app from Google Play or the Apple App Store and choose from the list of devices that are already preconfigured to securely connect to AWS. After selecting your device, you choose the Lambda function you want to run when the device sends the trigger. You can write Lambda functions to start or stop a machine or remotely control your home appliances. Once you select the Lambda function you want to use, you can deploy it to your device with a single click. After your devices are deployed, AWS IoT 1-Click.

Top benefits include:

- **Easily Create Lambda Triggers for Any Device** - With AWS IoT 1-Click, you can choose the action for your device by selecting one of the predefined AWS Lambda functions for common actions like sending emails or SMS messages, or you can select from custom Lambda functions you have created. You can write code to do things like playing a prerecorded message, counting assets, or tracking completion of tasks on an assembly line.
- **Deploy with a Single Click** - Devices listed in the AWS IoT 1-Click mobile app are already pre-configured to connect securely to AWS IoT Core, so you don't have to create, install, or manage any credentials like certificates. To deploy your devices, you assign an AWS Lambda function to each device, and click a button in the AWS IoT 1-Click mobile app.
- **Easily Manage Simple Devices** - AWS IoT 1-Click provides simple forms for you to specify and update user ID, location, and device name for your devices and group your devices by these attributes, making it easy to organize and manage many different device types. You can also view reports in the AWS IoT 1-Click mobile app to track the utilization and status of your deployed devices, including information about how many Lambda invocations have failed or which devices have low battery life.

87.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon FreeRTOS. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

87.3 Service Constraints

Please see <https://aws.amazon.com/iot-core/> for more information.

87.4 Technical Requirements

Amazon FreeRTOS is an operating system that makes microcontroller-based edge devices easy to program, deploy, secure, and maintain. Amazon FreeRTOS extends the FreeRTOS kernel, a popular open source operating system for microcontrollers, with software libraries that can be used to securely connect devices (locally and to the cloud) and update them remotely.

- **User Guide** - Provides detailed information about the microcontroller operating system that makes small, low-powered edge devices easy to program, deploy, secure, and maintain. [HTML](#) | [PDF](#)

Additional documentation is available at <https://aws.amazon.com/documentation/freertos/>.

88.0 Service Definition – AWS IoT Analytics

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

88.1 Service Overview

AWS IoT Analytics is a fully-managed service that makes it easy to run sophisticated analytics on massive volumes of IoT data without having to worry about all the cost and complexity typically required to build your own IoT analytics platform. It is the easiest way to run analytics on IoT data and get insights to make better and more accurate decisions for IoT applications and machine learning use cases.

IoT data is highly unstructured which makes it difficult to analyze with traditional analytics and business intelligence tools that are designed to process structured data. IoT data comes from devices that often record fairly noisy processes (such as temperature, motion, or sound), and as a result the data from these devices can frequently have significant gaps, corrupted messages, and false readings that must be cleaned up before analysis can occur. Also, IoT data is often only meaningful in the context of other data from external sources. For example, to determine when to water their crops, vineyard irrigation systems often enrich humidity sensor data with data on rainfall at the vineyard, allowing them to be more efficient with water usage while maximizing their harvest yield.

AWS IoT Analytics automates each of the difficult steps that are required to analyze data from IoT devices. IoT Analytics filters, transforms, and enriches IoT data before storing it in a time-series data store for analysis. You can setup the service to collect only the data you need from your devices, apply mathematical transforms to process the data, and enrich the data with device-specific metadata such as device type and location before storing the processed data. Then, you can analyze your data by running ad hoc or scheduled queries using the built-in SQL query engine, or perform more complex analytics and machine learning inference. IoT Analytics makes it easy to get started with machine learning by including pre-built models for common IoT use cases so you can quickly answer questions like which devices are about to fail or which customers are at risk of abandoning their wearable devices.

AWS IoT Analytics is fully managed and scales automatically to support up to petabytes of IoT data. With IoT Analytics, you can analyze data from millions of devices and build fast, responsive IoT applications without managing hardware or infrastructure.

For more information, visit the [AWS IoT Analytics documentation](#) page.

Top benefits include:

- **Easily Run Queries on IoT Data** - With AWS IoT Analytics, you can run simple, ad-hoc queries using the built-in IoT Analytics SQL query engine. The service allows you to use standard SQL queries to extract data from the data store to answer questions like the average distance traveled for a fleet of connected vehicles or how many doors are locked after 7pm in a smart building. These queries can be re-used even if connected devices, fleet size, and analytic requirements change.
- **Run Time-Series Analytics** - AWS IoT Analytics also supports time-series analyses so you can analyze the performance of devices over time and understand how and where they are being used, continuously monitor device data to predict maintenance issues, and monitor sensors to predict and react to environmental conditions.
- **Data Storage Optimized for IoT** - AWS IoT Analytics stores the processed device data in a time-series data store that is optimized to deliver fast response times on IoT queries that typically include time as a criteria. The raw data is also automatically stored for later processing or to reprocess it for another use case.
- **Prepares Your IoT Data for Analysis** - AWS IoT Analytics includes data preparation techniques that make it easy to prepare and process your data for analysis. IoT Analytics is integrated with AWS IoT Core so it is easy to ingest device data directly from connected devices. It cleans false readings, fills gaps in the data, and performs mathematical transformations of message data. As the data is ingested, IoT Analytics can process it using conditional statements, filter data to collect just the data you want to analyze, and enrich it with information from the AWS IoT registry. You can also use AWS Lambda functions to enrich your device data from external sources like the Weather Service, HERE Maps, Salesforce, or Amazon DynamoDB. For example, you could

combine weather data and mapping information to create better information about a device's environment.

- **Tools for Machine Learning** - AWS IoT Analytics makes it easy to apply machine learning to your IoT data with hosted Jupyter notebooks. You can directly connect your IoT data to the notebook and build, train, and execute models right from the IoT Analytics console without having to manage any of the underlying infrastructure. Using AWS IoT Analytics, you can apply machine learning algorithms to your device data to produce a health score for each device in your fleet. For example, an auto manufacturer can detect which of their customers have worn brake pads and alert them to seek maintenance for their vehicles.
- **Automated Scaling with Pay-As-You Go Pricing** - AWS IoT Analytics is a fully managed and pay-as-you go service that scales automatically to support up to petabytes of IoT data. With IoT Analytics, you can analyze your entire fleet of connected devices without managing hardware or infrastructure. As your needs change, compute power and the data store automatically scale up or down so you always have the right capacity for your IoT applications and you only pay for the resources that you use.

88.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS IoT Analytics. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

88.3 Service Constraints

Please see <https://aws.amazon.com/documentation/iotanalytics/> for more information.

88.4 Technical Requirements

AWS IoT Analytics provides advanced data analysis for AWS IoT. You can collect large amounts of device data, process messages, and store them. You can then query the data and run sophisticated analytics to make accurate decisions in your IoT applications and machine learning use cases.

- **User Guide** - Walks through how to set up the service and integrate with other AWS services. [HTML](#) | [PDF](#)
- **User Guide** - Describes all the API operations for AWS IoT Analytics in detail. Also provides sample requests, responses, and errors for the supported web services protocols. [HTML](#) | [PDF](#)

89.0 Service Definition – AWS IoT Button

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

89.1 Service Overview

The AWS IoT Button is a programmable button based on the Amazon Dash Button hardware. This simple Wi-Fi device is easy to configure and designed for developers to get started with [AWS IoT Core](#), [AWS Lambda](#), [Amazon DynamoDB](#), [Amazon SNS](#), and many other Amazon Web Services without writing device-specific code.

You can code the button's logic in the cloud to configure button clicks to count or track items, call or alert someone, start or stop something, order services, or even provide feedback. For example, you can click the button to unlock or start a car, open your garage door, call a cab, call your spouse or a customer service representative, track the use of common household chores, medications or products, or remotely control your home appliances.

The button can be used as a remote control for Netflix, a switch for your Philips Hue light bulb, a check-in/check-out device for Airbnb guests, or a way to order your favorite pizza for delivery. You can integrate it with third-party APIs like Twitter, Facebook, Twilio, Slack or even your own company's applications. Connect it to things we haven't even thought of yet. We can't wait to see what you will build with the AWS IoT Button!

Top benefits include:

- Easy to configure
- Designed for developers
- Get started with AWS IoT Core, AWS Lambda, Amazon DynamoDB

For more information, please refer to <https://aws.amazon.com/iotbutton/>.

90.0 Service Definition – AWS IoT Device Defender (in Preview)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

90.1 Service Overview

AWS IoT Device Defender is a fully managed service that helps you secure your fleet of IoT devices. AWS IoT Device Defender continuously audits the security policies associated with your devices to make sure that they aren't deviating from security best

practices. A security policy is a set of technical controls that devices follow to help keep information secure when communicating with other devices and the cloud. AWS IoT Device Defender makes it easy to maintain and enforce security policies, such as ensuring device identity, authenticating and authorizing devices, and encrypting device data. AWS IoT Device Defender continuously audits the security policies on your devices against a set of predefined security best practices. AWS IoT Device Defender sends an alert if there are any gaps in your policies that might create a security risk, such as identity certificates being shared across multiple devices or a device with a revoked identity certificate trying to connect to [AWS IoT Core](#).

AWS IoT Device Defender also lets you monitor devices for behavior that deviates from what you have defined as appropriate behavior for each device. Then, if something doesn't look right, AWS IoT Device Defender sends out an alert so you can take action to remediate the issue. For example, traffic spikes in outbound traffic might indicate that a device is participating in a DDoS attack.

AWS IoT Device Defender can send alerts to the [AWS IoT Console](#), [Amazon CloudWatch](#), and [Amazon SNS](#). If you determine that you need to take an action based on an alert, you can use the [AWS IoT Device Management](#) service to take mitigating actions such as pushing security fixes.

Top benefits include:

- **Audit Device Configurations for Security Vulnerabilities** - AWS IoT Device Defender audits the security policies associated with your devices against a set of defined IoT security best practices so you know exactly where you have security gaps. You can run audits on a continuous or ad-hoc basis. AWS IoT Device Defender comes with security best practices that you can select and run as part of the audit. For example, you can create an audit to check for identity certificates that are inactive, revoked, expiring, or pending transfer in less than 7 days. Audits make it possible for you to continuously monitor security policies as device configurations change.
- **Continuously Monitor Device Behavior to Identify Anomalies** - AWS IoT Device Defender detects anomalies in device behavior that may indicate a compromised device by monitoring incoming device metrics and data and comparing them against expected device behavior that you define. For example, AWS IoT Device Defender lets you define how many ports are open on the device, who the device can talk to, where it is connecting from, and how much data it sends or receives. Then it monitors the device traffic and alerts you if something looks wrong, like traffic from devices to a known malicious IP or unauthorized endpoints.
- **Receive Alerts and Take Action** - AWS IoT Device Defender publishes security alerts to the AWS IoT Console, Amazon CloudWatch, and Amazon SNS when a security policy audit fails or when behavior anomalies are detected so you can investigate and determine the root cause. For

example, AWS IoT Device Defender can alert you when device identities haven't been used for a long time or when device identities are accessing sensitive APIs. From the AWS IoT Device Defender console, you can also view recommended actions you can take to minimize the impact of security issues such as revoking permissions, rebooting a device, resetting factory defaults, or pushing security fixes to any of your connected devices. Then you can use the AWS IoT Device Management service to take the desired action.

91.0 Service Definition – AWS IoT Device Management

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

91.1 Service Overview

AWS IoT Device Management makes it easy to securely onboard, organize, monitor, and remotely manage IoT devices at scale. IoT Device Management lets you register your devices individually or in bulk, and manage permissions so that devices remain secure. Then, you use the IoT Device Management console to organize your devices into groups, monitor and troubleshoot device functionality, and send remote updates to your devices. AWS IoT Device Management allows you to scale your device fleets and reduce the cost and effort of managing large IoT device deployments.

Top benefits include:

- **Fast Device Onboarding** - With AWS IoT Device Management, you can securely add device attributes like device name, type and manufacturing year, certificates and access policies to the AWS IoT Registry in bulk, assign them to devices, and put large fleets of devices into service quickly.
- **Simple Device Organization** - AWS IoT Device Management lets you organize your devices into groups and manage access policies for these groups. This makes it easy to track, operate, and manage your devices according to business and security requirements, such as deploying a firmware update for all devices in a building or defining how devices communicate with each other. You can create a hierarchy for your groups such as grouping multiple sensors within a single vehicle and grouping multiple vehicles in a fleet. Then, your devices will inherit access policies based on the group hierarchy.
- **Locate Devices Quickly** - AWS IoT Device Management lets you quickly search and find any device across your entire device fleet in near real-time. You can easily find devices based on a combination of attributes like device ID, device state and type, and quickly find specific devices so that you can take action or troubleshoot your devices.
- **Easy Remote Management** - AWS IoT Device Management makes it easy for you to maintain the health of your device fleet. With IoT Device Management, you can remotely update the software running on your devices after they have been deployed in the field – allowing you to ensure that devices are always running on the latest software. You can also remotely execute actions on the devices such as reboots, factory resets, software updates, and security patches..

91.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS IoT Device Management. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

91.3 Service Constraints

Please see <https://aws.amazon.com/documentation/iot-device-management/> for more information.

91.4 Technical Requirements

AWS IoT Device Management is a cloud-based device management service that makes it easy for customers to securely manage IoT devices throughout their lifecycle. Customers can use AWS IoT Device Management to onboard device information and configuration, organize their device inventory, monitor their fleet of devices, and remotely manage devices deployed across many locations. This remote management includes over-the-air (OTA) updates to device software.

- **User Guide** - Describes device provisioning, thing groups, and jobs that are sent to and executed on one or more devices connected to AWS IoT Core.
[HTML](#) | [PDF](#)
- **API Reference** - Describes all the API operations for AWS IoT Device Management in detail. Provides examples of request and response syntax.
[HTML](#) | [PDF](#)

92.0 Service Definition – Amazon WorkSpaces

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

92.1 Service Overview

Amazon WorkSpaces is a fully managed, secure desktop computing service that runs on the AWS Cloud. Amazon WorkSpaces allows you to easily provision cloud-based virtual desktops and provide your users access to the documents, applications, and resources they need from any supported device, including Windows and Mac computers, Chromebooks, iPads, Kindle Fire tablets, and Android tablets. With just a

few clicks in the AWS Management Console, you can deploy high-quality cloud desktops for any number of users at a cost that is competitive with traditional desktops and half the cost of most VDI solutions.

Top benefits include:

- **Simple to Use** – Amazon WorkSpaces makes it easy to manage your desktop computing infrastructure by eliminating the need for up-front investments and avoiding the complexity of maintaining, patching, and managing a large physical desktop environment or a complex VDI solution.
- **Keep Your Data Secure** – Amazon WorkSpaces provides each user with access to persistent storage in the AWS Cloud. When users access their desktops using Amazon WorkSpaces, a remote display protocol is used to compress, encrypt, and encode data so that only images are transmitted. Amazon WorkSpaces also integrates with the [AWS KMS](#), providing the ability to encrypt WorkSpaces storage volumes.
- **Deploy and Manage Applications** – Amazon WorkSpaces Application Manager (Amazon WAM) offers a fast, flexible, and secure way for you to package, deploy, and update your organisation's desktop applications for Amazon WorkSpaces.
- **Choose the Hardware and Software You Need** – Amazon WorkSpaces offers a choice of bundles providing different amounts of CPU, memory, and storage so you can match your Amazon WorkSpaces to your requirements. Amazon WorkSpaces offers preinstalled OS and applications (including Microsoft Office), or you can bring your own Windows desktop licenses or other licensed software.
- **Support Multiple Devices** – Because the desktops are in the cloud, users can access their Amazon WorkSpaces from any supported device including Windows and Mac computers, Chromebooks, iPads, Kindle Fire tablets, and Android tablets.
- **Integrate Your Corporate Directory** – Amazon WorkSpaces securely integrates with your corporate AD so that your users can continue using their existing enterprise credentials to seamlessly access company resources. This also makes it easy to manage your WorkSpaces using familiar systems management tools.

92.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon WorkSpaces. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

92.3 Service Constraints

Please see <https://aws.amazon.com/workspaces/> for more information.

92.4 Technical Requirements

Amazon WorkSpaces offers you an easy way to provide a cloud-based desktop experience to your end users. You simply select from a choice of bundles that offer a

range of different amounts of CPU, memory, storage, and a choice of applications. Users can connect from a PC, Mac desktop computer, iPad, Kindle, or Android tablet.

- **Administration Guide** – Helps you get started using Amazon WorkSpaces. You will learn how to quickly and easily provision and maintain one or more WorkSpaces.
[HTML](#) | [PDF](#) | [Kindle](#)
- **Developer Guide** – Describes the API operations for Amazon WorkSpaces.
[HTML](#) | [PDF](#)

93.0 Service Definition – Amazon WorkMail

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

93.1 Service Overview

Amazon WorkMail is a secure, managed business email and calendar service with support for existing desktop and mobile email clients. Amazon WorkMail gives users the ability to seamlessly access their email, contacts, and calendars using Microsoft Outlook, their web browser, or their native iOS and Android email applications. You can integrate Amazon WorkMail with your existing corporate directory and control both the keys that encrypt your data and the location in which your data is stored. You can get started for free with a 30-day trial of Amazon WorkMail for up to 25 users.

Top benefits include:

- **Managed Service** – Amazon WorkMail makes it easy to manage your corporate email infrastructure and eliminates the need for up-front investments to license and provision on-premises email servers. There is no complex software to install or maintain and no hardware to manage. Amazon WorkMail automatically handles all of the patches, back-ups, and upgrades.
- **Enterprise-Grade Security** – Amazon WorkMail automatically encrypts all of your data at rest with encryption keys you control, using the [AWS KMS](#). Amazon WorkMail also allows you to retain full control over data locality by choosing the AWS Region where all of your data is stored.
- **Outlook Compatible** – Amazon WorkMail provides native support for Microsoft Outlook on both Windows and Mac OS X, so users can continue to use the email client they are already using without needing to install any additional software, such as plug-ins for Microsoft Outlook.
- **Anywhere Access** – Users can synchronise their mailboxes with iOS, Android, Amazon Fire, and Windows Phone devices (using the Microsoft Exchange ActiveSync protocol). If you are migrating from an on-premises Microsoft Exchange server, your users' mobile devices can automatically connect to Amazon WorkMail with no end-user reconfiguration required and no change in user experience. A feature-rich web client is also available for users to access their email, calendar, and contacts.
- **Active Directory Integration** – Amazon WorkMail securely integrates with your existing Microsoft Active Directory so that users can access their mailbox using their existing credentials. This also makes it easy to manage users and groups with familiar systems management tools, such as Active Directory Users and Computers.

- **Low Cost** – Amazon WorkMail features simple, low, monthly per-user pricing and costs \$4 per user per month, which includes 50GB of storage per user.

93.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon WorkMail. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

93.3 Service Constraints

Please see <https://aws.amazon.com/workmail/> for more information.

93.4 Technical Requirements

Amazon WorkMail is a managed email and calendaring service that offers strong security controls and support for existing desktop and mobile clients.

- **User Guide** – Helps you configure your desktop email client or mobile device for Amazon WorkMail and use the Amazon WorkMail web application.
[HTML](#) | [PDF](#)
- **Administrator Guide** – Helps you perform administrative tasks for Amazon WorkMail, such as setting up Amazon WorkMail for your organisation, adding a domain, managing users, groups, and mobile devices.
[HTML](#) | [PDF](#)

94.0 Service Definition – Amazon WorkDocs

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

94.1 Service Overview

Amazon WorkDocs is a fully managed, secure enterprise storage and sharing service with strong administrative controls and feedback capabilities that improve user productivity. Users can comment on files, send them to others for feedback, and upload new versions without having to resort to emailing multiple versions of their files as attachments. Users can take advantage of these capabilities wherever they are, using the device of their choice, including PCs, Macs, tablets and phones. Amazon WorkDocs offers IT administrators the option of integrating with existing corporate directories, flexible sharing policies and control of the location where data is stored. Customers can

get started using Amazon WorkDocs with a 30-day free trial providing 1 TB of storage per user for up to 50 users.

Amazon WorkDocs offers an Administrative SDK. The Administrative SDK allows you to integrate your applications with Amazon WorkDocs by performing content and permissions updates, and managing users, programmatically.

Top benefits include:

- **Simple Document Feedback** – Users can comment on files, send them to others for feedback, and upload new versions without having to resort to emailing multiple versions of their files as attachments. Teammates can leave detailed comments on sections of a file as well as overall comments. Amazon WorkDocs notifies contributors about review activities and deadlines via email and offers versioning capabilities.
- **Central Hub** – Amazon WorkDocs provides users with a central location for both the documents and files they are reviewing as well as those they own and are soliciting feedback on. With all these files in one location, reviewers have access to all of the related feedback in a single web view, making reading or contributing comments as simple as a few clicks.
- **Access and Sync from Any Device** – Users can access files stored in Amazon WorkDocs and view and leave feedback on files anywhere, anytime, from the device of their choice. Users can leverage the Amazon WorkDocs apps for iPad, Kindle Fire, and Android tablets. The Amazon WorkDocs Sync client lets users have files saved on their computer automatically uploaded to Amazon WorkDocs over an encrypted connection and synced across their other devices.
- **Secure** – Amazon WorkDocs offers flexible security settings and access controls to manage data storage and sharing. You can use policies to control users' sharing behaviour and choose the AWS Region where your users' data is stored. All data is encrypted in transit and at rest, and users don't have to send documents as email attachments.
- **Integrate Your Corporate Directory** – Amazon WorkDocs can integrate with your existing Active Directory. This means that your users can easily access Amazon WorkDocs using their existing Active Directory credentials and allows you to control which users in your organisation are permitted to access the service.

94.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon WorkDocs. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

94.3 Service Constraints

Please see <https://aws.amazon.com/workdocs/> for more information.

94.4 Technical Requirements

Amazon WorkDocs is a fully managed, secure enterprise storage and sharing service with strong administrative controls and feedback capabilities that improve user productivity.

- **Administration Guide** – Helps you use Amazon WorkDocs to perform several administrative tasks, such as creating a new directory in the cloud, connecting to your on-premises directory, or setting user privileges and defaults.
[HTML](#) | [PDF](#)
- **User Guide** – Helps you use the Amazon WorkDocs collaboration and utility applications. File collaboration applications are available for desktop web browsers, as well as several tablets.
[HTML](#) | [PDF](#)

95.0 Service Definition – AWS Server Migration Service

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

95.1 Service Overview

AWS Server Migration Service is an agentless service which makes it easier and faster for you to migrate thousands of on-premises workloads to AWS. AWS Server Migration Service allows you to automate, schedule, and track incremental replications of live server volumes, making it easier for you to coordinate large-scale server migrations.

Top benefits include:

- **Easy to Get Started** – Start and manage server migration with a few clicks via the AWS Management Console. AWS Server Migration Service will automatically replicate live server volumes to AWS and create AMIs as needed.
- **Control** – Create and manage a customised replication schedule designed for large-scale migrations, and track the progress of each migration.
- **Agility** – Perform migrations faster while minimizing network bandwidth, by migrating only incremental changes made to on-premises servers.
- **Cost-Effective** – AWS Server Migration Service is free to use; pay only for the storage resources used during the migration process. Lea
- **Minimise Downtime** – Incremental server replication allows you to reduce server downtime significantly

95.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Server Migration Service. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

95.3 Service Constraints

Please see
http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_server_migration for more information.

95.4 Technical Requirements

AWS Server Migration Service combines data collection tools with automated server replication to speed the migration of on-premises servers to AWS.

- **User Guide**
Describes key concepts of AWS Server Migration Service and provides instructions for using the features of AWS Server Migration Service.
[HTML](#) | [PDF](#)
- **API Reference**
Documents the AWS Server Migration Service Query API.
[HTML](#) | [PDF](#)
- **AWS Server Migration Service section of AWS CLI Reference**
Documents the AWS Server Migration Service commands available in the AWS CLI.
[HTML](#)

96.0 Service Definition – Amazon Lightsail

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

96.1 Service Overview

Amazon Lightsail is designed to be the easiest way to launch and manage a virtual private server with AWS. Amazon Lightsail plans include everything you need to jumpstart your project—a virtual machine, SSD-based storage, data transfer, DNS management, and a static IP—for a low, predictable price. Amazon Lightsail virtual private servers run on the same highly available and reliable AWS Cloud infrastructure used by millions of customers.

Top benefits include:

- **Amazon Lightsail Virtual Private Server (VPS)** – Experience the power and reliability of AWS. Deploy in seconds and manage from the intuitive Lightsail management console or API.
- **Keeping servers simple** – Launch a virtual private server with just a few clicks
- **Built on the galaxy's leading cloud** – Give your application the power, reliability, and security of AWS.

- **You know code, we know infrastructure** – Let Amazon Lightsail do the heavy lifting. We make server management a breeze.
- **The cloud that grows with you** – Scale applications with access to dozens of AWS Cloud services.
- **Powerful API** – Use the simple and flexible Amazon Lightsail API to extend your application or integrate it with external applications.
- **Highly available storage** – Every Amazon Lightsail server comes with high-performing, persistent SSD-based block storage.
- **Speedy & secure networking** – Amazon Lightsail servers run at warp speed on the AWS network. Configure your network simply and securely, including your IP addresses, DNS, firewall, and more.
- **Snapshots** – Protect your data, clone your server, and more with Amazon Lightsail snapshots.
- **Access to AWS Services** – Extend the capabilities of your Amazon Lightsail server by connecting it to popular AWS Cloud services, including managed databases, CDN, and many others.

96.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Lightsail. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

96.3 Service Constraints

Please see <https://amazonlightsail.com/> for more information.

96.4 Technical Requirements

Amazon Lightsail is the easiest way to get started with AWS for developers who just need virtual private servers. Lightsail includes everything you need to launch your project quickly—a virtual machine, SSD-based storage, data transfer, DNS management, and a static IP—for a low, predictable price.

- **Administration Guide** – Describes how to get started with Amazon Lightsail instances to create a development environment or an application. Quickly configure the domain name and back it up using a snapshot.
[HTML](#)
- **API References** – Describes all the API operations for Amazon Lightsail in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

97.0 Service Definition – AWS Batch

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

97.1 Service Overview

AWS Batch enables developers, scientists, and engineers to easily and efficiently run hundreds of thousands of batch computing jobs on AWS. AWS Batch dynamically provisions the optimal quantity and type of compute resources (e.g., CPU or memory optimised instances) based on the volume and specific resource requirements of the batch jobs submitted. With AWS Batch, there is no need to install and manage batch computing software or server clusters that you use to run your jobs, allowing you to focus on analysing results and solving problems. AWS Batch plans, schedules, and executes your batch computing workloads across the full range of AWS compute services and features, such as Amazon EC2 and Spot Instances.

Top benefits include:

- **Fully Managed** – AWS Batch eliminates the need to operate third-party commercial or open source batch processing solutions. There is no batch software or servers to install or manage. AWS Batch manages all the infrastructure for you, avoiding the complexities of provisioning, managing, monitoring, and scaling your batch computing jobs.
- **Integrated with AWS** – AWS Batch is natively integrated with the AWS platform, allowing you to leverage the scaling, networking, and access management capabilities of AWS. This makes it easy to run jobs that safely and securely retrieve and write data to and from AWS data stores such as Amazon S3 or Amazon DynamoDB.
- **Cost-Optimised Resource Provisioning** – AWS Batch provisions compute resources and optimises the job distribution based on the volume and resource requirements of the submitted batch jobs. AWS Batch dynamically scales compute resources to any quantity required to run your batch jobs, freeing you from the constraints of fixed-capacity clusters. AWS Batch also dynamically bids for Spot Instances on your behalf, reducing the cost of running your batch jobs further.

97.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Batch. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

97.3 Service Constraints

Please see <https://aws.amazon.com/batch> for more information.

97.4 Technical Requirements

AWS Snowball is a service for customers who want to transport terabytes or petabytes of data to and from AWS, or who want to access the storage and compute power of the

AWS Cloud locally and cost effectively in places where connecting to the internet may not be an option.

- **User Guide** – Describes key concepts of AWS Batch and provides instructions for using the features of AWS Batch.
[HTML](#) | [PDF](#)
- **API References** – Describes all the API operations for AWS Batch in detail.
[HTML](#) | [PDF](#)

98.0 Service Definition – AWS Serverless Application Repository

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

98.1 Service Overview

The AWS Serverless Application Repository is a collection of serverless applications and serverless application components provided by AWS and other AWS partners and developers. With a growing selection of applications to choose from, the AWS Serverless Application Repository makes it easy to get started with the AWS serverless platform. The AWS Serverless Application Repository includes applications for Alexa Skills, chatbots, data processing, IoT, real time stream processing, web and mobile back-ends, social media trend analysis, image resizing, and more from publishers on AWS.

Top benefits include:

- **Deploy quickly** – quickly deploy code samples, components, and complete applications for common use cases such as web and mobile back-ends, event and data processing, logging, monitoring, IoT, and more.
- **Easily define AWS resources used** - each application is packaged with an [AWS Serverless Application Model \(SAM\)](#) template that defines the AWS resources used.
- **Publishing Applications to the Repository** - Publish applications to the AWS Serverless Application Repository to share solutions with developers or to help your customers quickly understand the value of products and services you sell and support.
- **Easy to share your published applications within your team**, across your organisation, or with the community at large. Publicly shared applications must include a link to the application's source code so others can view what the application does and how it works.

98.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Serverless Application Repository. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/documentation/>

98.3 Service Constraints

AWS Serverless Application Repository features and default service limits are found at <https://aws.amazon.com/serverless/serverlessrepo/faqs/>

98.4 Technical Requirements

With AWS Serverless Application Repository, you can quickly find and deploy serverless applications in the AWS Cloud. You can browse applications by category, or search for them by name, publisher, or event source. To use an application, you simply select it, configure required fields, and deploy it with a few clicks.

As a serverless application developer, you can also easily publish applications, sharing them across teams and organizations and with other AWS users. To publish serverless applications, you can use the AWS Management Console, AWS CLI, or AWS SDKs to upload the code, along with a simple manifest file.

- **Developer Guide**
Provides a conceptual overview of AWS Serverless Application Repository, includes detailed instructions for using features, and provides a complete API reference for developers.
[HTML](#) | [PDF](#)

99.0 Service Definition – AWS Organizations

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

99.1 Service Overview

AWS Organizations offers policy-based management for multiple AWS accounts. With Organizations, you can create groups of accounts and then apply policies to those groups. Organizations enables you to centrally manage policies across multiple accounts, without requiring custom scripts and manual processes.

Using AWS Organizations, you can create Service Control Policies (SCPs) that centrally

control AWS service use across multiple AWS accounts. You can also use Organizations to help automate the creation of new accounts through APIs. Organizations helps simplify the billing for multiple accounts by enabling you to setup a single payment method for all the accounts in your organization through consolidated billing. AWS Organizations is available to all AWS customers at no additional charge.

Top benefits include:

- **Centrally manage policies across multiple AWS accounts** - AWS Organizations helps you manage policies for multiple AWS accounts. With Organizations, you can create groups of accounts, and then attach policies to a group to ensure the correct policies are applied across the accounts. Organizations enables you to centrally manage policies across multiple accounts, without requiring custom scripts and manual processes.
- **Control access to AWS services** - With AWS Organizations, you can create Service Control Policies (SCPs) that centrally control AWS service use across multiple AWS accounts. SCPs put bounds around the permissions that AWS Identity and Access Management (IAM) policies can grant to entities in an account, such as IAM users and roles. For example, IAM policies for an account in your organization cannot grant access to AWS Direct Connect if access is not also allowed by the SCP for the account. Entities can only use the services allowed by both the SCP and the IAM policy for the account.
- **Automate AWS account creation and management** - You can use the AWS Organizations APIs to automate the creation and management of new AWS accounts. The Organizations APIs enable you to create new accounts programmatically, and to add the new accounts to a group. The policies attached to the group are automatically applied to the new account. For example, you can automate the creation of sandbox accounts for developers and grant entities in those accounts access only to the necessary AWS services.
- **Consolidate billing across multiple AWS accounts** - AWS Organizations enables you to set up a single payment method for all the AWS accounts in your organization through consolidated billing. With consolidated billing, you can see a combined view of charges incurred by all your accounts, as well as take advantage of pricing benefits from aggregated usage, such as volume discounts for Amazon EC2 and Amazon S3.

99.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Organizations. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

99.3 Service Constraints

Please see <https://aws.amazon.com/organizations/> for more information.

99.4 Technical Requirements

AWS Organizations is a service that enables Amazon Web Services (AWS) customers

to consolidate and centrally manage multiple AWS accounts. With AWS Organizations, you can create accounts and invite existing accounts to join your organization. You can organize those accounts into groups and attach policy-based controls. If you already have a Consolidated Billing family of accounts, those accounts automatically become part of your organization.

- **User Guide** - Introduces you to AWS Organizations, helps you set up an organization by inviting other accounts to join, and shows you how to organize your accounts into groups to control access to your AWS resources by applying service control policies.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** - Describes all the API operations for AWS Organizations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)
- **AWS Organizations section of AWS CLI Reference** - Describes the AWS CLI commands that you can use to administer AWS Organizations. Provides syntax, options, and usage examples for each command.
[HTML](#).

100.0 Service Definition – AWS Shield

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

100.1 Service Overview

AWS Shield is a managed DDoS protection service that safeguards web applications running on AWS. AWS Shield provides always-on detection and automatic inline mitigations that minimise application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection. There are two tiers of AWS Shield – Standard and Advanced. All AWS customers benefit from the automatic protections of AWS Shield Standard, at no additional charge. AWS Shield Standard defends against most common, frequently occurring network and transport layer DDoS attacks that target your website or applications.

For higher levels of protection against attacks targeting your web applications running on Elastic Load Balancing, Amazon CloudFront, and Amazon Route 53 resources, you

can subscribe to AWS Shield Advanced. In addition to the common network and transport layer protections that come with Standard, AWS Shield Advanced provides additional detection and mitigation against large and sophisticated DDoS attacks, near real-time visibility into attacks, and integration with AWS WAF, a web application firewall. AWS Shield Advanced also gives you access to the AWS DRT and protection against DDoS related spikes in your Elastic Load Balancing, Amazon CloudFront, or Amazon Route 53 charges.

AWS Shield Advanced is available globally on all Amazon CloudFront and Amazon Route 53 edge locations. You can protect your web applications hosted anywhere in the world by deploying Amazon CloudFront in front of your application. Your origin servers can be Amazon S3, Amazon EC2, Elastic Load Balancing, or a custom server outside of AWS. You can also enable AWS Shield Advanced directly on Elastic Load Balancing in the following AWS Regions: Northern Virginia, Oregon, Ireland, and Tokyo.

Top benefits include:

- **Seamless Integration and Deployment** – With AWS Shield your AWS resources are automatically protected from common, most frequently occurring network and transport layer DDoS attacks. You can achieve a higher level of defence by simply enabling AWS Shield Advanced protection for Elastic Load Balancing, Amazon CloudFront or Amazon Route 53 resources you want to protect using the management console or APIs.
- **Customisable Protection** – With AWS Shield Advanced, you have the flexibility to write customised rules to mitigate sophisticated application layer attacks. These customisable rules can be deployed instantly, allowing you to quickly mitigate attacks. You can set up proactive rules to automatically block bad traffic, or respond to incidents as they occur. You can also engage the 24x7 AWS DRT, who can write rules on your behalf to mitigate application layer DDoS attacks.
- **Cost Efficient** – As an AWS customer, you automatically get network layer protection against some of the most common DDoS attacks with AWS Shield. This protection does not require additional cost, resources, or time to initiate. With AWS Shield Advanced, you get “DDoS cost protection”, a feature that protects your AWS bill from Elastic Load Balancing, Amazon CloudFront, and Amazon Route 53 usage spikes as a result of a DDoS attack.

100.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Shield Advanced. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

100.3 Service Constraints

Please see <https://aws.amazon.com/shield/> for more information.

100.4 Technical Requirements

For additional protection against DDoS attacks, AWS also offers AWS Shield Advanced. AWS Shield Advanced provides expanded DDoS attack protection for your Amazon CloudFront distributions, Amazon Route 53 hosted zones, and Elastic Load Balancing load balancers. AWS Shield Advanced incurs additional charges.

- **Developer Guide** – Describes how to get started with AWS WAF and AWS Shield Advanced, explains key concepts, and provides step-by-step instructions that show you how to use the features.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes all the API operations for AWS Shield Advanced in detail.
[HTML](#) | [PDF](#)

101.0 Service Definition – AWS Application Discovery Service

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

101.1 Service Overview

AWS Application Discovery Service helps customers quickly and reliably plan application migration projects by automatically identifying applications running in on-premises data centres, their associated dependencies, and their performance profile. Planning data centre migrations can involve thousands of workloads that are often deeply interdependent. Application discovery and dependency mapping are important early first steps in the migration process, but difficult to perform at scale due to the lack of automated tools.

AWS Application Discovery Service automatically collects configuration and usage data from servers, storage and networking equipment to develop a list of applications, how they perform, and how they are interdependent. This information is retained in encrypted format in an AWS Application Discovery Service database which you can export as a CSV or XML file into your preferred visualization tool or cloud migration solution to help reduce the complexity and time in planning your cloud migration.

Top benefits include:

- **Simplify the Cloud Migration Planning Process** – AWS Application Discovery Service automatically identifies installed applications, network dependencies, and system performance data. Data from the service can be used to find multi-tier applications that should be migrated as a group. Start and manage discovery process with a few clicks via the AWS Management Console.
- **Easily Integrate with Cloud Migration Tools** – AWS Application Discovery Service provides a set of public APIs and an open data format. This approach allows you to collect data from other discovery solutions in the Application Discovery database, and use this data in a cloud migration solution or within a cloud migration framework provided by System Integration partners.
- **Protect Data with End-to-End Encryption** – AWS Application Discovery Service helps ensure that collected application configuration data is protected end-to-end by encrypting it on-premises, maintaining encryption in transit to AWS, and storing in an encrypted format in the Application Discovery database.
- **Consult with AWS Professionals and APN Partners** – AWS Application Discovery Service is available through engagements with AWS Professional Services or through one of our APN Partners. These trained professionals have helped hundreds of enterprise customers successfully complete their migration to the cloud.

101.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Application Discovery Service. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

101.3 Service Constraints

Please see <https://aws.amazon.com/documentation/application-discovery/> for more information.

101.4 Technical Requirements

The AWS Application Discovery Service helps customers quickly and reliably plan application migration projects by automatically identifying applications running in on-premises data centres, their associated dependencies, and their performance profile.

- **User Guide** – Walks through how to set up the AWS Application Discovery Service and integrate it with other services.
[HTML](#) | [PDF](#)
- **API Reference** – Describes all the API operations for the AWS Application Discovery Service in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

102.0 Service Definition – AWS CodeBuild

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

102.1 Service Overview

AWS CodeBuild is a fully managed build service that compiles source code, runs tests, and produces software packages that are ready to deploy. With AWS CodeBuild, you don't need to provision, manage, and scale your own build servers. AWS CodeBuild scales continuously and processes multiple builds concurrently, so your builds are not left waiting in a queue. You can get started quickly by using pre-packaged build environments, or you can create custom build environments that use your own build tools. With AWS CodeBuild, you are charged by the minute for the compute resources you use.

Top benefits include:

- **Fully Managed Build Service** – AWS CodeBuild eliminates the need to set up, patch, update, and manage your own build servers and software. There is no software to install or manage.
- **Continuous Scaling** – AWS CodeBuild scales automatically to meet your build volume. It immediately processes each build you submit and can run separate builds concurrently, which means your builds are not left waiting in a queue.
- **Pay as You Go** – With AWS CodeBuild, you are charged based on the number of minutes it takes to complete your build. This means you no longer have to worry about paying for idle build server capacity.
- **Extensible** – You can bring your own build tools and programming runtimes to use with AWS CodeBuild by creating customised build environments in addition to the pre-packaged build tools and runtimes supported by AWS CodeBuild.
- **Enables Continuous Integration and Delivery** – AWS CodeBuild belongs to a family of [AWS Code Services](#), which you can use to create complete, automated software release workflows for [continuous integration/continuous delivery](#) (CI/CD). You can also integrate AWS CodeBuild into your existing CI/CD workflow. For example, you can use AWS CodeBuild as a worker node for your existing Jenkins server setup for distributed builds.
- **Secure** – with AWS CodeBuild, your build artefacts are encrypted with customer-specific keys that are managed by the [AWS KMS](#). CodeBuild is integrated with [AWS IAM](#), so you can assign user-specific permissions to your build projects.

102.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS CodeBuild. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

102.3 Service Constraints

Please see <https://aws.amazon.com/codebuild/> for more information.

102.4 Technical Requirements

AWS CodeBuild is a fully managed build service that compiles your source code, runs unit tests, and produces artefacts that are ready to deploy.

- **User Guide** – Describes how to get started building your source code using AWS CodeBuild.
[HTML](#) | [PDF](#) | [Kindle](#)
- **API Reference** – Describes the API operations for AWS CodeBuild. Also provides details of related request and response syntax and errors.
[HTML](#) | [PDF](#)
- **AWS CLI Reference for AWS CodeBuild** – Describes the AWS CLI commands that you can use to automate building your source code.
[HTML](#)

103.0 Service Definition – AWS CodeStar

The following subsections provide service definition information.

103.1 Service Overview

AWS CodeStar enables you to quickly develop, build, and deploy applications on AWS. AWS CodeStar provides a unified user interface, enabling you to easily manage your software development activities in one place. With AWS CodeStar, you can set up your entire [continuous delivery](#) toolchain in minutes, allowing you to start releasing code faster. AWS CodeStar makes it easy for your whole team to work together securely, allowing you to easily manage access and add owners, contributors, and viewers to your projects. Each AWS CodeStar project comes with a project management dashboard, including an integrated issue tracking capability powered by Atlassian JIRA Software. With the AWS CodeStar project dashboard, you can easily track progress across your entire software development process, from your backlog of work items to teams' recent code deployments.

There is no additional charge for using AWS CodeStar. You only pay for the AWS resources that you provision for developing and running your application (for example, Amazon EC2 instances).

Top benefits include:

- **Start developing on AWS in minutes** - AWS CodeStar makes it easy for you to set up your entire development and continuous delivery toolchain for coding, building, testing, and deploying your application code. To start a project, you can choose from a variety of AWS CodeStar templates for [Amazon EC2](#), [AWS Lambda](#), and [AWS Elastic Beanstalk](#). You have the option to choose [AWS CodeCommit](#) or GitHub to use as your project's source control. You also have the option to edit your source code using one of several options including [AWS Cloud9](#), Microsoft Visual Studio, or Eclipse. After you make your selections the underlying AWS services are provisioned in minutes, allowing you to quickly start coding and deploying your applications.
- **Manage software delivery in one place** - AWS CodeStar provides an easy way to coordinate your day-to-day development activities through a unified user interface, reducing the need to switch between various service consoles. AWS CodeStar's project dashboard lets you monitor application activity, and track progress across all stages of your software development process, including code commits, builds, tests, and deployments, from a central place. AWS CodeStar integrates [Atlassian JIRA Software](#), a third-party issue tracking and project management tool, allowing you to easily manage JIRA issues directly in the AWS CodeStar dashboard.
- **Work across your team securely** - AWS CodeStar enables you to collaborate on projects across your team in a secure manner. You can easily manage access for project owners, contributors, and viewers without needing to manually configure your own policy for each service. AWS CodeStar simplifies the process of setting up project access for teams by providing built-in role-based policies that follow AWS [Identity and Access Management](#) best practices.
- **Choose from a variety of project templates** - With AWS CodeStar project templates, you can easily develop a variety of applications such as websites, web applications, web services, and Alexa skills. AWS CodeStar project templates include the code for getting started on supported programming languages including Java, JavaScript, PHP, Ruby, C#, and Python.

103.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS CodeStar. For additional information beyond what is described herein, please refer to <https://docs.aws.amazon.com/codestar/latest/userguide/welcome.html>.

103.3 Service Constraints

Please see <https://aws.amazon.com/codestar/faqs/> for more information.

103.4 Technical Requirements

AWS CodeStar lets you quickly develop, build, and deploy applications on AWS.

- **User Guide**

Provides conceptual overviews of AWS CodeStar and explains how to use it to develop software applications on AWS.

[HTML](#) | [PDF](#)

- **API**

Reference

Describes all the API operations for AWS CodeStar. Provides examples of request and response syntax.

[HTML](#) | [PDF](#)

104.0 Service Definition – AWS Step Functions

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

104.1 Service Overview

AWS Step Functions makes it easy to coordinate the components of distributed applications and microservices using visual workflows. Building applications from individual components that each perform a discrete function lets you scale and change applications quickly. AWS Step Functions is a reliable way to coordinate components and step through the functions of your application. AWS Step Functions provides a graphical console to arrange and visualise the components of your application as a series of steps. This makes it simple to build and run multi-step applications. AWS Step Functions automatically triggers and tracks each step, and retries when there are errors, so your application executes in order and as expected. AWS Step Functions logs the state of each step, so when things do go wrong, you can diagnose and debug problems quickly. You can change and add steps without even writing code, so you can easily evolve your application and innovate faster.

AWS Step Functions manages the operations and underlying infrastructure for you to help ensure your application is available at any scale

Top benefits include:

- **Productivity: Build Applications Quickly** – AWS Step Functions includes a visual console and blueprints for commonly-used workflows that make it easy to coordinate the components of distributed applications into parallel and/or sequential steps. You can build applications in a matter of minutes, and then visualise and track the execution of each step to help ensure the application is operating as intended
- **Resilience: Scale and Recover Reliably** – AWS Step Functions automatically triggers each step so your application executes in order and as expected. It can handle millions of steps simultaneously to help ensure your application is available as demand increases. AWS Step Functions tracks the state of each step and handles errors with built-in retry and fallback, whether the step takes seconds or months to complete.
- **Agility: Evolve Applications Easily** – AWS Step Functions makes it easy to change workflows and edit the sequence of steps without revising the entire application. You can re-use components and steps without even changing their code to experiment and innovate faster. Your workflow can support thousands of individual components and steps, so you can freely build increasingly complex applications.

104.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for AWS Step Functions. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

104.3 Service Constraints

Please see <https://aws.amazon.com/step-functions/> for more information.

104.4 Technical Requirements

AWS Step Functions makes it easy to coordinate the components of distributed applications as a series of steps in a visual workflow. You can quickly build and run state machines to execute the steps of your application in a reliable and scalable fashion.

- **Developer Guide** – Describes key concepts of AWS Step Functions and provides instructions for using the features of AWS Step Functions.
[HTML](#) | [PDF](#)
- **API Reference** – Documents the AWS Step Functions API.
[HTML](#) | [PDF](#)

105.0 Service Definition – Amazon Athena

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

105.1 Service Overview

Amazon Athena is an interactive query service that makes it easy to analyse data in Amazon S3 using standard SQL. Amazon Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run.

Amazon Athena is easy to use. Simply point to your data in Amazon S3, define the schema, and start querying using standard SQL. Most results are delivered within seconds. With Amazon Athena, there's no need for complex ETL jobs to prepare your data for analysis. This makes it easy for anyone with SQL skills to quickly analyse large-scale datasets.

Top benefits include:

- **Start Querying Instantly: Serverless. No ETL** – Amazon Athena is serverless. You can quickly query your data without having to setup and manage any servers or data warehouses. Just point to your data in Amazon S3, define the schema, and start querying using the built-in query editor. Amazon Athena allows you to tap into all your data in S3 without the need to set up complex processes to extract, transform, and load the data (ETL).
- **Pay Per Query: Only pay for data scanned** – With Amazon Athena, you pay only for the queries that you run. You can save from 30% to 90% on your per-query costs and get better performance by compressing, partitioning, and converting your data into columnar formats. Amazon Athena queries data directly in Amazon S3. There are no additional storage charges beyond S3.
- **Open. Powerful. Standard: Built on Presto. Runs standard SQL** – Amazon Athena uses Presto with ANSI SQL support and works with a variety of standard data formats, including CSV, JSON, ORC, Avro, and Parquet. Amazon Athena is ideal for quick, ad-hoc querying but it can also handle complex analysis, including large joins, window functions, and arrays. Amazon Athena is highly available; and executes queries using compute resources across multiple facilities and multiple devices in each facility. Amazon Athena uses Amazon S3 as its underlying data store, making your data highly available and durable.
- **Fast. Really Fast: Interactive performance even for large datasets** – With Amazon Athena, you don't have to worry about having enough compute resources to get fast, interactive query performance. Amazon Athena automatically executes queries in parallel, so most results come back within seconds.

105.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Athena. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

105.3 Service Constraints

Please see <https://aws.amazon.com/athena/> for more information.

105.4 Technical Requirements

Amazon Athena is an interactive query service that makes it easy to analyse data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to setup or manage, and you pay only for the queries you run. To get started, simply point to your data in S3, define the schema, and start querying using standard SQL.

- **User Guide** – Learn how to use Athena to query data stored in Amazon S3.
[HTML](#) | [PDF](#)

106.0 Service Definition – Amazon Pinpoint

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

106.1 Service Overview

Amazon Pinpoint makes it easy to run targeted campaigns to drive user engagement in mobile apps. Amazon Pinpoint helps you understand user behaviour, define which users to target, determine which messages to send, schedule the best time to deliver the messages, and then track the results of your campaign.

Targeted push notifications based on app usage trends and user behaviour have become a popular approach for mobile app user engagement because response rates are often several times higher than traditional email marketing campaigns. By using targeted push notifications, you can increase message relevance and effectiveness, measure engagement, and continually improve your campaigns.

Getting started with Amazon Pinpoint is easy. First, AWS Mobile Hub guides you through the process to integrate the AWS Mobile SDK with your app. Next, you define your target segments, campaign message, and specify the delivery schedule. Once your campaign is running, Amazon Pinpoint provides metrics so you can run analytics and track the impact of your campaign.

With Amazon Pinpoint, there are no upfront setup costs and no fixed monthly cost. You only pay for the number of users your campaign targets, the messages you send, and the events you collect, so you can start small and scale as your application grows.

Top benefits include:

- **Understand User Behaviour** – Amazon Pinpoint gives you a clear view of how users are engaging with your mobile app. Amazon Pinpoint includes real-time analytics with dashboards for analysing user acquisition, user engagement, monetization, user demographics, custom events, and campaign funnels so you can understand how users engage with your application. With Pinpoint, you can filter user data by entire segments, segmentation attributes, or time, and then view and drill down on the data in the Pinpoint dashboard.
- **Create Targeted Campaigns** – Amazon Pinpoint lets you create targeted campaigns that deliver relevant and personalised push notification messages to targeted segments of your app's user base. Segmentation for the audience of your campaign can be defined by common metrics like the number of times your app is opened, or custom metrics like usage of specific features in your app. You can use data from a variety of different sources to define your target segments in Amazon Pinpoint. Pinpoint can automatically collect user data from your mobile app, or import data collected in other AWS Cloud services such as Amazon S3 and Amazon Redshift. Integration with data in third party data sources is easy. Just export your data to Amazon S3 and use the pre-built integration to import your data into Pinpoint. Once you define your target segments you can run push notification campaigns that can be customised and integrated directly into the user experiences of your mobile app.
- **Measure Results** – Amazon Pinpoint provides reporting on app usage activity to track the impact of your campaign on user engagement, including the number of times the app was opened as a result of the campaign, and revenue generated from campaigns. You can also export the resulting event data and run custom analytics using your existing analytics tools. Pinpoint can also help you A/B test different messages, track results, and then send the best message to your target segment.

106.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Pinpoint. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

106.3 Service Constraints

Please see <https://aws.amazon.com/pinpoint/> for more information.

106.4 Technical Requirements

Amazon Pinpoint is a campaign management web service for mobile apps. Use Amazon Pinpoint to understand app user behaviour and define which users to engage. Amazon Pinpoint helps you determine which push notifications to send, decide when to deliver the notifications, and track the results of your campaign.

- **User Guide** – Describes key concepts for Amazon Pinpoint and provides instructions for using the Amazon Pinpoint console.
[HTML](#) | [PDF](#)

- **Developer Guide** – Describes how to integrate Amazon Pinpoint functionality into your app and includes development instructions for its features.
[HTML](#) | [PDF](#)
- **API Reference** – Describes all the Amazon Pinpoint REST API actions.
[HTML](#) | [PDF](#)

107.0 Service Definition – Amazon SageMaker

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

107.1 Service Overview

Amazon SageMaker is a fully-managed platform that enables developers and data scientists to quickly and easily build, train, and deploy machine learning models at any scale. Amazon SageMaker removes all the barriers that typically slow down developers who want to use machine learning.

Machine learning often feels a lot harder than it should be to most developers because the process to build and train models, and then deploy them into production is too complicated and too slow. First, you need to collect and prepare your training data to discover which elements of your data set are important. Then, you need to select which algorithm and framework you'll use. After deciding on your approach, you need to teach the model how to make predictions by training, which requires a lot of compute. Then, you need to tune the model so it delivers the best possible predictions, which is often a tedious and manual effort. After you've developed a fully trained model, you need to integrate the model with your application and deploy this application on infrastructure that will scale. All of this takes a lot of specialized expertise, access to large amounts of compute and storage, and a lot of time to experiment and optimize every part of the process. In the end, it's not a surprise that the whole thing feels out of reach for most developers.

Amazon SageMaker removes the complexity that holds back developer success with each of these steps. Amazon SageMaker includes modules that can be used together or independently to build, train, and deploy your machine learning models.

Top benefits include:

- **Get to Production with Machine Learning Quickly** - Amazon SageMaker significantly reduces the amount of time needed to train, tune, and deploy machine learning models. Amazon SageMaker manages and automates all the sophisticated training and tuning techniques so you can get models into production quickly.
- **Choose Any Framework or Algorithm** - Amazon SageMaker supports all machine algorithms and frameworks so you can use the technology you are already familiar with. Apache MXNet and TensorFlow are pre-installed, and Amazon SageMaker offers a range of built-in, high performance machine learning algorithms. If you want to train with an alternative framework or algorithm, you can bring your own in a Docker container.
- **One-Click Training and Deployment** - Amazon SageMaker lets you begin training your model with a single click in the console or with a simple API call. When the training is complete, and you're ready to deploy your model, you can launch it with a single click in the Amazon SageMaker console.
- **Easily Integrate With Your Existing Workflow** - Amazon SageMaker is designed in three modules that can be used together or independently as part of any existing ML workflow you might already have in place.
- **Easy Access to Trained Models** - Amazon SageMaker makes it easy to integrate machine learning models into your applications by providing an HTTPS endpoint that can be called from any application.
- **Optimized for Speed** - Amazon SageMaker is pre-configured with the latest versions of TensorFlow and Apache MXNet, with CUDA9 library support for maximum performance with NVIDIA GPUs. With Amazon SageMaker P3 instances running NVIDIA Volta V100 GPUs, Amazon SageMaker lets you train deep learning models with unparalleled speed.

107.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon SageMaker. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

107.3 Service Constraints

Please see <https://aws.amazon.com/sagemaker/> for more information.

107.4 Technical Requirements

Amazon SageMaker is a fully managed machine learning service. With Amazon SageMaker, data scientists and developers can quickly and easily build and train machine learning models, and then directly deploy them into a production-ready hosted environment. It provides an integrated Jupyter authoring notebook instance for easy access to your data sources for exploration and analysis, so you don't have to manage servers. It also provides common machine learning algorithms that are optimized to run efficiently against extremely large data in a distributed environment. With native support for bring-your-own-algorithms and frameworks, Amazon SageMaker provides flexible distributed training options that adjust to your specific workflows.

- **Developer Guide** - Provides a conceptual overview of Amazon SageMaker, offers step-by-step instructions for training and deploying models, and includes a complete API reference for developers.
[HTML](#) | [PDF](#)

108.0 Service Definition – Amazon Comprehend

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

108.1 Service Overview

Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to find insights and relationships in text. Amazon Comprehend identifies the language of the text; extracts key phrases, places, people, brands, or events; understands how positive or negative the text is; and automatically organizes a collection of text files by topic.

You can use the Amazon Comprehend APIs to analyze text and use the results in a wide range of applications including voice of customer analysis, intelligent document search, and content personalization for web applications.

The service constantly learns and improves from a variety of information sources, including Amazon.com product descriptions and consumer reviews – one of the largest natural language data sets in the world - to keep pace with the evolution of language..

Top benefits include:

- **Get better answers from your text** - Amazon Comprehend can discover the meaning and relationships in text from customer support incidents, product reviews, social media feeds, news articles, documents, and other sources. For example, you can identify the feature that's most often mentioned when customers are happy or unhappy about your product.
- **Identify topics in a collection of text** - Amazon Comprehend can analyze a collection of documents and other text files (such as social media posts) and automatically organize them by relevant terms or topics. You can then use the topics to deliver personalized content to your customers or to provide richer search and navigation. For example, if you have an extensive collection of news articles, you can automatically group them by subject matter to enable your site to suggest new articles to visitors based on what they've read previously.
- **Continually learning, always improving** - Amazon Comprehend is always learning from new data and continually gains a deeper understanding of entities, key phrases, sentiments, and topics. The learning data covers a wide range of domains including Finance, Health, Media, Telecom, Education, Government, Advertising, and many others. The service also learns from

Amazon.com product descriptions and customer reviews, which is one of the largest collections of natural language in the world.

- **Integrated with Amazon S3 and AWS Glue** - Amazon Comprehend is integrated with Amazon S3 and AWS Glue. You can use Amazon S3 to store documents and texts, and directly access them from Amazon Comprehend. In addition, you can use AWS Glue to load documents and texts from various AWS data stores (Amazon Redshift, Amazon RDS, Amazon DynamoDB, etc.) to Amazon Comprehend, and subsequently load the results from Amazon Comprehend back into these AWS data stores for downstream analytics.

108.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Comprehend. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

108.3 Service Constraints

Please see <https://aws.amazon.com/comprehend/> for more information.

108.4 Technical Requirements

Amazon Comprehend uses natural language processing (NLP) to extract insights about the content of documents without the need of any special preprocessing. Amazon Comprehend processes any text files in UTF-8 format. It develops insights by recognizing the entities, key phrases, language, sentiments, and other common elements in a document. Use Amazon Comprehend to create new products based on understanding the structure of documents. With Amazon Comprehend you can search social networking feeds for mentions of products, scan an entire document repository for key phrases, or determine the topics contained in a set of documents.

- **Developer Guide** - Provides a conceptual overview of Amazon Comprehend, includes detailed instructions for using the various features, and provides a complete API reference for developers. [HTML](#) | [PDF](#) | [Kindle](#)

109.0 Service Definition – Amazon Lex

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

109.1 Service Overview

Amazon Lex is a service for building conversational interfaces into any application using voice and text. Amazon Lex provides the advanced deep learning functionalities of automatic speech recognition (ASR) for converting speech to text, and natural language

understanding (NLU) to recognize the intent of the text, to enable you to build applications with highly engaging user experiences and lifelike conversational interactions. With Amazon Lex, the same deep learning technologies that power Amazon Alexa are now available to any developer, enabling you to quickly and easily build sophisticated, natural language, conversational bots (“[chatbots](#)”).

Speech recognition and natural language understanding are some of the most challenging problems to solve in computer science, requiring sophisticated deep learning algorithms to be trained on massive amounts of data and infrastructure. Amazon Lex democratizes these deep learning technologies by putting the power of Amazon Alexa within reach of all developers. Harnessing these technologies, Amazon Lex enables you to define entirely new categories of products made possible through conversational interfaces.

As a fully managed service, Amazon Lex scales automatically, so you don’t need to worry about managing infrastructure. With Amazon Lex, you pay only for what you use. There are no upfront commitments or minimum fees.

Top benefits include:

- **Easy to Use** - Amazon Lex provides an easy-to-use console to guide you through the process of creating your own chatbot in minutes, building conversational interfaces into your applications. You supply just a few example phrases and Lex builds a complete natural language model through which your user can interact using voice and text, to ask questions, get answers, and complete sophisticated tasks.
- **Seamlessly Deploy and Scale** - With Amazon Lex, you can build, test, and deploy your chatbots directly from the Amazon Lex console. Lex enables you to easily publish your voice or text chatbots to mobile devices, web apps, and chat services such as Facebook Messenger (with Slack and Twilio integration coming soon). Once published, your Amazon Lex bot processes voice or text input in conversation with your end-users. Amazon Lex is a fully managed service so as your user engagement increases, you don’t need to worry about provisioning hardware and managing infrastructure to power your bot experience.
- **Built-in Integration with the AWS Platform** - Amazon Lex provides built-in integration with AWS Lambda, AWS MobileHub and Amazon CloudWatch and you can easily integrate with many other services on the AWS platform including Amazon Cognito, and Amazon DynamoDB. You can take advantage of the power of the AWS platform for security, monitoring, user authentication, business logic, storage and mobile app development.
- **Cost Effective** - With Amazon Lex, there are no upfront costs or minimum fees. You are only charged for the text or speech requests that are made. Amazon Lex’ pay-as-you-go pricing and low cost per request make it a cost-effective way to build conversational interfaces anywhere. With the Amazon Lex free tier, you can easily try Lex without any initial investment.

109.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Lex. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

109.3 Service Constraints

Please see <https://aws.amazon.com/lex/> for more information.

109.4 Technical Requirements

Amazon Lex is an AWS service for building conversational interfaces into applications using voice and text. With Amazon Lex, the same deep learning engine that powers Amazon Alexa is now available to any developer, enabling you to build sophisticated, natural language chatbots into your new and existing applications. Amazon Lex provides the deep functionality and flexibility of natural language understanding (NLU) and automatic speech recognition (ASR) to enable you to build highly engaging user experiences with lifelike, conversational interactions and create new categories of products.

- **Developer Guide** - Provides a conceptual overview of Amazon Lex, includes detailed instructions for using the various features, and provides a complete API reference for developers. [HTML](#) | [PDF](#)

110.0 Service Definition – Amazon Polly

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

110.1 Service Overview

Amazon Polly is a service that turns text into lifelike speech. Amazon Polly lets you create applications that talk, enabling you to build entirely new categories of speech-enabled products. Amazon Polly is an Amazon AI service that uses advanced deep learning technologies to synthesise speech that sounds like a human voice. Amazon Polly includes 47 lifelike voices spread across 24 languages, so you can select the ideal voice and build speech-enabled applications that work in many different countries.

Amazon Polly delivers the consistently fast response times required to support real-time, interactive dialog. You can cache and save Amazon Polly's speech audio to replay offline or redistribute. And Amazon Polly is easy to use—you simply send the text you want converted into speech to the Amazon Polly API, and Polly immediately returns the audio stream to your application so your application can play it directly or store it in a standard audio file format, such as MP3.

With Amazon Polly, you only pay for the number of characters you convert to speech, and you can save and replay Amazon Polly's generated speech. Amazon Polly's low cost per character converted, and lack of restrictions on storage and reuse of voice output, make it a cost-effective way to enable Text-to-Speech everywhere.

Top benefits include:

- **Natural Sounding Voices** – Amazon Polly provides 47 lifelike voices and supports 24 languages, including a wide range of male and female voices with a variety of accents. Amazon Polly's fluid pronunciation of text in multiple languages enables you to deliver high-quality voice output and create applications for global users.
- **Store and Redistribute Speech** – Unlike other solutions that require a royalty or charge a fee every time you replay previously generated audio, Amazon Polly allows for unlimited replays without any additional fees. These free replays extend to offline use as well. You can create speech files in a variety of standard formats, such as MP3 and OGG, and store these on devices such as a mobile phones or IoT devices for offline playback.
- **Fast Response** – Delivering lifelike voices and conversational user experiences requires consistently fast response times. Voice-enabled applications need to play synthesised speech without delay. Consider apps that provide spoken directions for navigation, eLearning applications that provide verbal instruction to students, and apps that engage the user through real time dialog. These apps are most effective when responses can start without perceived delays in the conversational flow. Even when you send lengthy text to Amazon Polly's API, it returns the audio to your application as a stream so you can play the voices immediately. These kinds of dynamic, spoken responses require access to a much larger quantity of speech audio than is typically available to store on users' devices. Amazon Polly is in the cloud, so you have access to a wide variety of synthesised speech. With Amazon Polly, your application can provide even more valuable responses that include real-time data.
- **Easy Integration** – Amazon Polly makes it easy to add voice to your website, mobile app, or device. With Polly, you just write the text you want converted to speech to the Amazon Polly API and it immediately returns the audio stream. Unlike other solutions that require a lengthy approval process, Polly doesn't require you to describe how you will use Amazon Polly's speech in your application, and there are no distribution agreements to sign, so you can start right away.
- **Low Cost** – Amazon Polly's pay-as-you-go pricing, low cost per character converted, and unlimited replays make it a cost-effective way to enable speech synthesis in virtually any application.

110.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Polly. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

110.3 Service Constraints

Please see <https://aws.amazon.com/polly/details/> for more information.

110.4 Technical Requirements

Amazon Polly is a Text-to-Speech cloud service that converts text into lifelike speech. You can use Amazon Polly to develop applications that increase engagement and accessibility. Amazon Polly supports multiple languages and includes a variety of lifelike voices, so you can build speech-enabled applications that work in multiple locations and use the ideal voice for your customers.

- **Developer Guide** – Provides a conceptual overview of Amazon Polly, including detailed instructions for using the various features.
[HTML](#) | [PDF](#)

111.0 Service Definition – Amazon Rekognition

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

111.1 Service Overview

Amazon Rekognition is a service that makes it easy to add image analysis to your applications. With Amazon Rekognition, you can detect objects, scenes, and faces in images. You can also search and compare faces. Amazon Rekognition's API enables you to quickly add sophisticated deep learning-based visual search and image classification to your applications.

Amazon Rekognition is based on the same proven, highly scalable, deep learning technology developed by Amazon's computer vision scientists to analyse billions of images daily for Prime Photos. Amazon Rekognition uses deep neural network models

to detect and label thousands of objects and scenes in your images, and we are continually adding new labels and facial recognition features to the service.

Amazon Rekognition's API lets you easily build powerful visual search and discovery into your applications. With Amazon Rekognition, you only pay for the images you analyse and the face metadata you store. There are no minimum fees and there are no upfront commitments.

Top benefits include:

- **Easily Integrate Powerful Image Recognition into Your App** – Amazon Rekognition removes the complexity of building image recognition capabilities into your applications by making powerful and accurate image analysis available with a simple API. You don't need computer vision or deep learning expertise to take advantage of Amazon Rekognition's reliable image analysis. With Rekognition's API, you can easily and quickly build image analysis into any web, mobile or connected device application.
- **Artificial Intelligence at the Core** – Amazon Rekognition is built on proven deep learning technology, already in use by Amazon. We continue to add support for new objects and to improve facial analysis. Amazon Rekognition's breadth and accuracy grow as we train it to take on new challenges.
- **Scalable Image Analysis** – Amazon Rekognition has been used to analyse billions of images per day. The service provides consistent response times regardless of the volume of analysis requests you make. Your application latency remains consistent, even as your request volume increases. There's no need to provision additional capacity when your app is successful and you suddenly have millions of users.
- **Integrated with Popular AWS Cloud Services** – Amazon Rekognition is designed to work seamlessly with popular AWS Cloud services like Amazon S3 and AWS Lambda. Amazon Rekognition's API can be called directly from AWS Lambda in response to Amazon S3 events. Since Amazon S3 and AWS Lambda scale automatically in response to your application's demand, you can build scalable, affordable, and reliable image analysis applications. For example, each time a person arrives at your residence, your door camera can upload a photo of the visitor to Amazon S3, triggering a AWS Lambda function that uses Amazon Rekognition APIs to identify your guest. You can run analysis directly on images stored in Amazon S3 without having to load or move the data. Support for AWS IAM makes it easy to securely control access to Amazon Rekognition APIs. Using AWS IAM, you can create and manage AWS users and groups to grant the appropriate access to your developers and end users.
- **Low Cost** – With Amazon Rekognition, you only pay for the number of images you analyse and the face metadata you store. There are no minimum fees or upfront commitments. Get started for free, and save more as you grow with Rekognition's tiered pricing model.

111.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Rekognition. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

111.3 Service Constraints

Please see <https://aws.amazon.com/rekognition/> for more information.

111.4 Technical Requirements

- **Developer Guide** – Provides a conceptual overview of Amazon Rekognition, including detailed instructions for using the various features.
[HTML](#) | [PDF](#)

112.0 Service Definition – Amazon Chime

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

112.1 Service Overview

Amazon Chime is a high-quality communications service that transforms online meetings with an easy-to-use app that works seamlessly across all your devices. With Amazon Chime, you can schedule and attend online meetings and video conferences, and chat, call, and collaborate, inside and outside your organisation, all with a single app. Now you can work productively from wherever you are.

Frustration free online meetings that start on time. High-definition video conferencing. A single, easy-to-use communications application that you can trust.

Top benefits include:

- **Frustration-Free Online Meetings** – Make your online meetings effortless with Amazon Chime. Meetings start on time, and they're easy to join. Meetings call you and with a single tap or click, you can join or notify participants that you're running late. A visual roster shows who is in the meeting and who is running late, and allows anyone to mute background noise. Amazon Chime automatically reconnects you if you get dropped. Crystal clear audio and HD video keep you focused on the discussion
- **Be Mobile** – Take your meetings anywhere with a rich, easy-to-use application available for Android, iOS, Mac, and Windows. Meetings and chats are always synchronised, you can join meetings from any device, and you can switch seamlessly between devices anytime – even in the middle of a meeting.
- **One Application for All Your Communications** – Use one application to conduct online meetings, video conferencing, chat, screen sharing, and file sharing, both inside and outside your organisation.

112.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Chime. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

112.3 Service Constraints

Amazon Chime features and default service limits are found at <https://chime.aws/faq/>.

112.4 Technical Requirements

Amazon Chime is a secure, real-time, unified communications service that transforms meetings by making them more efficient and easier to conduct.

- **Administrator Guide**
Helps you use Amazon Chime to perform several administrative tasks, such as creating an Amazon Chime account, inviting users, and managing licenses.
[HTML](#) | [PDF](#)
- **User Guide**
Helps you use Amazon Chime, including joining, attending, and scheduling meetings for your organisation.
[HTML](#) | [PDF](#)

113.0 Service Definition – Amazon Connect

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

113.1 Service Overview

Amazon Connect is a self-service, cloud-based contact centre service that makes it easy for any business to deliver better customer service at lower cost. Amazon Connect is based on the same contact centre technology used by Amazon customer service associates around the world to power millions of customer conversations. The self-service graphical interface in Amazon Connect makes it easy for non-technical users to design contact flows, manage agents, and track performance metrics—no specialised skills required. There are no up-front payments or long-term commitments and no infrastructure to manage with Amazon Connect; customers pay by the minute for Amazon Connect usage plus any associated telephony services.

Top benefits include:

- **Easy to set up and manage** – Setting up Amazon Connect is easy. With only a few clicks in the AWS Management Console, agents can begin taking calls within minutes. The drag and drop Contact Flow Editor allows you to create contact flows that define effective and efficient customer interactions without any coding.
- **Scalable and elastic** – There is no infrastructure to deploy or manage, so you can scale your Amazon Connect contact centre up or down, on-boarding up to tens of thousands of agents in

response to business cycles and paying only for the time callers are interacting with Amazon Connect plus any associated telephony charges.

- **Pay as you go** – As an on-demand service, you pay for Amazon Connect usage by the minute with no long-term commitments or upfront charges. There is no minimum monthly fee. You are charged based on the number of minutes you use Amazon Connect to engage with your end customers, at the specified per minute rate. Pricing is not based on capacity, agent seats, or maintenance.
- **Reliable** – Amazon Connect runs on AWS-proven [infrastructure](#) operating 42 Availability Zones within 16 geographic regions around the world. This makes Amazon Connect more highly available, fault tolerant, and scalable than would be possible if a contact centre solution was run from a single data centre.
- **Open platform** – Amazon Connect is an open platform so it is easy to integrate with other systems such as Customer Relationship Management (CRM) solutions or the AWS platform. For example, you can use [AWS Lambda](#) to run code for a serverless application or backend service to build contact flow experiences that adapt to your customer needs in real time. And [Amazon S3](#) can be used to store call recordings, where you can control access and lifecycle of the data.

113.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Connect. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

113.3 Service Constraints

Amazon Connect features and default service limits are found at <https://aws.amazon.com/connect/faqs/>.

113.4 Technical Requirements

Amazon Connect is a Contact Centre as a Service (CCaS) solution that offers easy, self-service configuration and enables dynamic, personal, and natural customer engagement at any scale.

- **Administrator Guide**
Describes an Amazon Connect virtual contact centre and helps administrators add users, choose telephony options, manage data storage, and configure encryption settings.
[HTML](#) | [PDF](#)
- **User Guide**
Describes key concepts of Amazon Connect and provides instructions for managing, measuring, and scaling a virtual contact centre.
[HTML](#) | [PDF](#)

114.0 Service Definition – Amazon Lumberyard

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

114.1 Service Overview

Amazon Lumberyard is the only game engine that gives you a combination of free, feature-rich development technology, native integration to the AWS Cloud to make it easier to create live and multiplayer games, and native integration of Twitch features that help you connect games to the world's leading social video platform and community for gamers.

By starting game projects with Lumberyard, you are able to spend more of your time creating differentiating gameplay and building communities of fans, and less time on the undifferentiated heavy lifting of building a game engine and managing server infrastructure.

Top benefits include:

- **Create the Highest-Quality Games:** You can use Amazon Lumberyard to build rich and engaging games with the highest ceiling of quality through its comprehensive and [proven](#) toolset, AAA performance, and modular design. With Lumberyard, you get a full-featured editor that lets you go from editing to playing your game in one keystroke (Ctrl-G), native code performance, stunning visuals, and hundreds of other features like performant networking, character and animation editors, particle editor, cloth physics, UI editor, audio tools, weather effects, flocking AI, perception handling, camera frameworks, path finding, and more. You also have full and free access to Lumberyard's source code, making it easy to customize the technology to create differentiated gameplay
- **Build Live, Multiplayer Features in Minutes**
- **Reach and Engage Fans on Twitch**

114.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Lumberyard. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

114.3 Service Constraints

Please see <https://aws.amazon.com/documentation/lumberyard/> for more information.

114.4 Technical Requirements

Lumberyard is a free AAA game engine deeply integrated with AWS and Twitch—with full source. Lumberyard provides a growing set of tools to help you create the highest quality games, engage massive communities of fans, and connect games to the vast compute and storage of the cloud. Participate in [the forums](#) or learn about new changes on [our blog](#).

- **User Guide** - Offers conceptual overviews and in-depth information for game developers and game engineers on how to use Lumberyard Editor and Lumberyard systems and features.
[HTML](#) | [PDF](#)
- **Legacy Reference** - Provides information on legacy Lumberyard systems and features like Flow Graph, Geppetto, and Mannequin.
[HTML](#) | [PDF](#)
- **Getting Started Guide** - Provides a detailed introduction to Lumberyard and includes a tutorial that walks you through an example use of Lumberyard Editor and other tools.
[HTML](#) | [PDF](#)
- **C++ API Reference Guide** - Describes the fundamental C++ API operations of the Lumberyard component entity system. This includes API operations for the component, entity, EBus, and reflection base classes.
[HTML](#)
- **Release Notes** - Presents a summary of features, improvements, resolved issues, and known issues in Lumberyard.
[HTML](#) | [PDF](#)

115.0 Service Definition – Amazon Sumerian (in Preview)

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

115.1 Service Overview

Amazon Sumerian lets you create and run virtual reality (VR), augmented reality (AR), and 3D applications quickly and easily without requiring any specialized programming or 3D graphics expertise. With Sumerian, you can build highly immersive and interactive scenes that run on popular hardware such as Oculus Rift, and HTC Vive as well as Android and iOS mobile devices. For example, you can build a virtual classroom that lets you train new employees around the world, or you can build a virtual environment that enables people to tour a building remotely. Sumerian makes it easy to create all the building blocks needed to build highly immersive and interactive 3D experiences including adding objects (e.g. characters, furniture, and landscape), and designing,

animating, and scripting environments. Sumerian does not require specialized expertise and you can design scenes directly from your browser

Top benefits include:

- **No VR/AR expertise needed** - Amazon Sumerian lets anyone create VR/AR applications. You don't need specialized programming or 3D graphics expertise to get started. You can create, edit, and publish applications with a web browser using the Sumerian editor.
- **Lifelike characters** - Sumerian allows you to easily create 3D characters, called Hosts, which narrate and guide users through your scenes. You can create lifelike, animated Hosts that speak a variety of languages and can converse with your users.
- **Design immersive environments** - Sumerian makes it easy to design scenes that immerse your users in 3D, lifelike surroundings that are populated with 3D objects and animated characters. You can import 3D objects and select from preloaded scene templates. You can also script the logic that controls how the objects in your scenes behave.
- **Run on multiple platforms** - Sumerian is platform-agnostic, so that once you're done building your scene, you can easily run it on popular hardware including Oculus Rift, HTC Vive, and mobile devices. Sumerian is based on the WebGL JavaScript API for graphic rendering and WebVR JavaScript API for virtual reality devices. It also supports the ARKit framework for augmented reality on iOS devices as well as ARCore for Android devices.

115.2 Backup/Restore and Disaster Recovery

This requirement is not applicable for Amazon Sumerian. For additional information beyond what is described herein, please refer to <http://aws.amazon.com/documentation/>.

115.3 Service Constraints

Please see <https://aws.amazon.com/greengrass/> for more information.

115.4 Technical Requirements

Amazon Sumerian is a set of tools for creating high-quality virtual reality (VR) experiences on the web. With Sumerian, you can construct an interactive 3D scene without any programming experience, test it in the browser, and publish it as a website that is immediately available to users.

- **User Guide** - Use Amazon Sumerian to create immersive virtual reality (VR) and augmented reality (AR) experiences with no programming required.

[HTML](#) | [PDF](#)

- **Tutorials and Articles** - Visit the Learn Amazon Sumerian website for developer articles and tutorials that introduce Sumerian features.

[HTML](#)

116.0 Service Definition – Amazon Gamelift

The following subsections provide service definition information in accordance with the requirements identified in the ITT.

116.1 Service Overview

Amazon GameLift, a managed service for deploying, operating, and scaling session-based multiplayer games, reduces the time required to build a multiplayer backend from thousands of hours to just minutes. Available for developers using [Amazon Lumberyard](#), Amazon GameLift is built on AWS's highly available cloud infrastructure and allows customers to quickly scale high-performance game servers up and down to meet player demand—without any additional engineering effort or upfront costs.

Top benefits include:

- **Server Management** - Amazon GameLift provisions instances, deploys your game servers onto running instances, load-balances traffic across fleets of game servers, monitors instance and game server health, and replaces unhealthy instances without your intervention.
- **Reliable Infrastructure** - Amazon GameLift runs within Amazon's proven network infrastructure and data centers.
- **Worldwide Game Placement** - Amazon GameLift continually scans for available game servers around the world. If low-latency game servers are not available, you can configure the service to automatically add more capacity nearby your players. Amazon GameLift maintains a queue of waiting players until new games start or new instances launch, then places those players into the lowest latency game.
- **Player Matchmaking** - Amazon GameLift's FlexMatch capabilities let you match players together based on rules you define. Whether you choose to build your matchmaking based on player skill, latency, or custom criteria, FlexMatch's simple but powerful rules language makes it easy for anyone to quickly create robust player matchmaking. FlexMatch results automatically route into a queue, which can then place the match in the AWS region that best balances player latency and server utilization.
- **Real-Time Insights** - The Amazon GameLift management console includes dashboards for monitoring the performance of your game. You can see real-time data on player demand, server capacity and health, CPU and memory utilization, and create operational alarms using Amazon CloudWatch. Amazon GameLift also lets you automatically collect and store game server logs for retrieval and inspection.

116.2 Backup/Restore and Disaster Recovery

Amazon GameLift is designed to safeguard your game servers from frequently occurring network and transport layer distributed denial of service (DDoS) attacks. You can also configure the service to autoscale server capacity to absorb DDoS attacks

without taking your players offline. For additional information beyond what is described herein, please refer to <https://aws.amazon.com/blogs/gametech/protect-multiplayer-game-servers-from-ddos-attacks-using-amazon-gamelift-2/>

116.3 Service Constraints

Please see <https://aws.amazon.com/gamelift/> for more information.

116.4 Technical Requirements

Amazon GameLift is a fully managed service for deploying, operating, and scaling session-based multiplayer game servers in the cloud. Built on the AWS computing environment, GameLift lets you scale high performance game servers up and down to meet player demand.

- **User Guide** - Provides a conceptual overview of Amazon GameLift and includes detailed instructions for the various features. Also guides you in using GameLift with the AWS Management Console.
[HTML](#) | [PDF](#)
- **API Reference** - Describes the Amazon GameLift API operations in detail. Also provides sample requests, responses, and errors for the supported web services protocols.
[HTML](#) | [PDF](#)

117.0 Cross Services

The following service definition topics are applicable to all AWS Service Offerings and are detailed once in a cross-service manner below.

117.1 Availability

The AWS Cloud infrastructure is built around regions and Availability Zones. A region is a physical location in the world where AWS has multiple Availability Zones. Availability Zones consist of one or more discrete data centres, each with redundant power, networking, and connectivity and housed in separate facilities. These Availability Zones offer you the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible with a single data centre. AWS currently has 18 regions, one Local Region around the world and 54 Availability Zones throughout the world: US East (Northern Virginia), US East (Ohio), US West (Oregon), US West (Northern California), AWS GovCloud (US) (Oregon), Canada (Central), EU (Ireland), EU (Frankfurt), EU (London), EU (Paris), South America (Sao Paulo), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), Asia Pacific (Seoul), Asia Pacific (Mumbai), Asia Pacific (Osaka), China (Ningxia) and China (Beijing). Information on each region can be found at the [AWS Global Infrastructure](#) web page.

Figure 1 depicts the current AWS Regions and Availability Zones, along with two new regions that are coming online throughout the next year. The AWS products and services that are available in each region are listed at the [Region Table](#) web page.



Region & Number of Availability Zones

US East
N. Virginia (6),
Ohio (3)

US West
N. California (3),
Oregon (3)

Asia Pacific
Mumbai (2),
Seoul (2),
Singapore (3),
Sydney (3),
Tokyo (4),
Osaka-Local (1)¹

Canada
Central (2)

China
Beijing (2),
Ningxia (2)

Europe
Frankfurt (3),
Ireland (3),
London (3),
Paris (3)

South America
São Paulo (3)

AWS
GovCloud (US-West) (3)



New Region (coming soon)

Bahrain
Hong Kong
SAR, China
Sweden
AWS GovCloud (US-East)



Figure 1 – Global Map of AWS Regions and Availability Zones

Figure 2 illustrates the relationship between regions and Availability Zones.

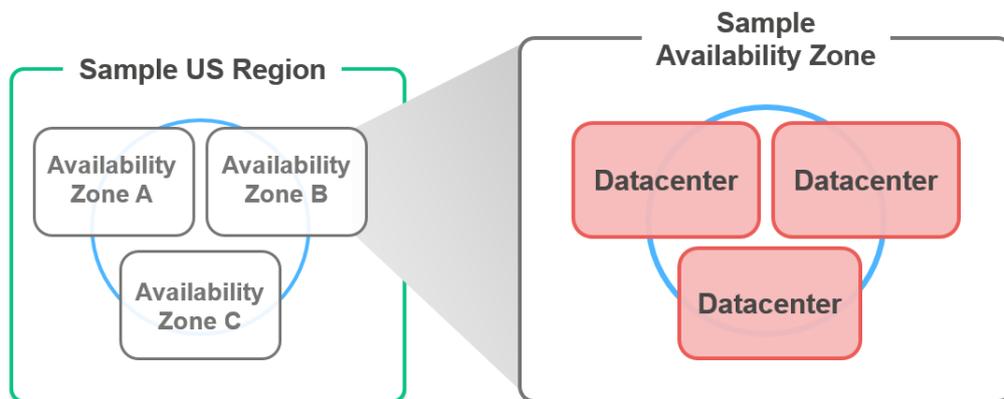


Figure 2 – Regions and Availability Zones

117.1.1 Region Availability

The AWS products and services that are available in the EU (Ireland), EU (London), EU (Paris) and EU (Frankfurt) regions are listed below. Exact service availability depends on a range of factors and choices made by customers when they architect and implement their solution; however, we currently have customers who are able to deliver service availability in these regions with very high SLAs using the tools and services provided by AWS.

The Services Offerings will be delivered from the AWS Region selected by the customer upon opening an AWS account. The customer may specify the AWS Region in which customer content will be stored. It is the customer's responsibility to select the relevant AWS Region in order to comply with its own security and governance requirements. AWS will not access or use customer content except as necessary to maintain or provide the Service Offerings. Customers acknowledge that AWS does not limit customers to any particular AWS Region. Note that not all AWS

services are available in every AWS Region; however, we are steadily expanding our service availability across AWS's global regions.

Table 1 – AWS Services in EMEA

| Services Offered: | Ireland | Frankfurt | London | Paris |
|---|---------|-----------|--------|-------|
| Alexa for Business | | | | |
| Amazon API Gateway | ✓ | ✓ | ✓ | ✓ |
| Amazon AppStream 2.0 | ✓ | ✓ | | |
| Amazon Athena | ✓ | ✓ | ✓ | |
| Amazon Aurora - MySQL-compatible | ✓ | ✓ | ✓ | ✓ |
| Amazon Aurora - PostgreSQL-compatible | ✓ | ✓ | ✓ | ✓ |
| Amazon Chime | | | | |
| Amazon Cloud Directory | ✓ | | ✓ | ✓ |
| Amazon CloudSearch | ✓ | ✓ | | |
| Amazon CloudWatch | ✓ | ✓ | ✓ | ✓ |
| Amazon CloudWatch Events | ✓ | ✓ | ✓ | ✓ |
| Amazon CloudWatch Logs | ✓ | ✓ | ✓ | ✓ |
| Amazon Cognito | ✓ | ✓ | ✓ | |
| Amazon Comprehend | ✓ | | | |
| Amazon Connect | | ✓ | | |
| Amazon DeepLens | ✓ | ✓ | ✓ | |
| Amazon DynamoDB | ✓ | ✓ | ✓ | ✓ |
| Amazon DynamoDB Accelerator (DAX) | ✓ | | | |
| Amazon Elastic Container Registry (ECR) | ✓ | ✓ | ✓ | ✓ |
| Amazon DynamoDB Accelerator (DAX) | ✓ | | | |
| Amazon Elastic Container Registry (ECR) | ✓ | ✓ | ✓ | ✓ |
| Amazon Elastic Container Service (ECS) | ✓ | ✓ | ✓ | ✓ |
| Amazon ElastiCache | ✓ | ✓ | ✓ | ✓ |
| Amazon Elastic Block Store (EBS) | ✓ | ✓ | ✓ | ✓ |
| Amazon Elastic Compute Cloud (EC2) | ✓ | ✓ | ✓ | ✓ |
| Amazon Elastic Compute Cloud (EC2) Auto Scaling | ✓ | ✓ | ✓ | ✓ |
| Amazon Elastic File System (EFS) | ✓ | ✓ | | |

| Services Offered: | Ireland | Frankfurt | London | Paris |
|--|---------|-----------|--------|-------|
| Amazon Elastic MapReduce | ✓ | ✓ | ✓ | ✓ |
| Amazon Elasticsearch Service | ✓ | ✓ | ✓ | ✓ |
| Amazon Elastic Transcoder | ✓ | | | |
| Amazon FreeRTOS | | | | |
| Amazon GameLift | ✓ | ✓ | ✓ | |
| Amazon Glacier | ✓ | ✓ | ✓ | ✓ |
| Amazon GuardDuty | ✓ | ✓ | ✓ | ✓ |
| Amazon Inspector | ✓ | ✓ | | |
| Amazon Kinesis Data Analytics | ✓ | | | |
| Amazon Kinesis Data Firehose | ✓ | ✓ | | |
| Amazon Kinesis Data Streams | ✓ | ✓ | ✓ | ✓ |
| Amazon Kinesis Video Streams | ✓ | ✓ | | |
| Amazon Lex | ✓ | | | |
| Amazon Lightsail | ✓ | ✓ | ✓ | |
| Amazon Machine Learning | ✓ | | | |
| Amazon Macie | | | | |
| Amazon Mobile Analytics | | | | |
| Amazon MQ | ✓ | ✓ | | |
| Amazon Neptune | | | | |
| Amazon Pinpoint | | | | |
| Amazon Polly | ✓ | ✓ | ✓ | ✓ |
| Amazon QuickSight | ✓ | | | |
| Amazon Redshift | ✓ | ✓ | ✓ | ✓ |
| Amazon Rekognition Image | ✓ | | | |
| Amazon Rekognition Video | ✓ | | | |
| Amazon Relational Database Service (RDS) | ✓ | ✓ | ✓ | ✓ |
| Amazon Route 53 Auto Naming | ✓ | | | |
| Amazon Sagemaker | ✓ | | | |
| Amazon SimpleDB | ✓ | | | |
| Amazon Simple Email Service (SES) | ✓ | | | |

| | Ireland | Frankfurt | London | Paris |
|--|----------------|------------------|---------------|--------------|
| Amazon Simple Notification Service (SNS) | ✓ | ✓ | ✓ | ✓ |
| Amazon Simple Queue Service (SQS) | ✓ | ✓ | ✓ | ✓ |
| Amazon Simple Storage Service (S3) | ✓ | ✓ | ✓ | ✓ |
| Amazon Simple Workflow Service (SWF) | ✓ | ✓ | ✓ | ✓ |
| Amazon Transcribe | ✓ | | | |
| Amazon Translate | ✓ | | | |
| Amazon Virtual Private Cloud (VPC) | ✓ | ✓ | ✓ | ✓ |
| Amazon WorkDocs | ✓ | | | |
| Amazon WorkMail | ✓ | | | |
| Amazon WorkSpaces | ✓ | ✓ | ✓ | |
| Amazon WorkSpaces Application Manager | ✓ | | | |
| AWS Auto Scaling | ✓ | | | |
| AWS Application Discovery Service | | | | |
| AWS AppSync | ✓ | | | |
| AWS Batch | ✓ | ✓ | ✓ | |
| AWS Certificate Manager | ✓ | ✓ | ✓ | ✓ |
| AWS Cloud9 | ✓ | | | |
| AWS CloudFormation | ✓ | ✓ | ✓ | ✓ |
| AWS CloudHSM | ✓ | ✓ | | |
| Services Offered: | Ireland | Frankfurt | London | Paris |

| | | | | |
|----------------------|---|---|---|---|
| AWS CloudHSM Classic | ✓ | ✓ | | |
| AWS CloudTrail | ✓ | ✓ | ✓ | ✓ |
| AWS CodeBuild | ✓ | ✓ | ✓ | ✓ |
| AWS CodeCommit | ✓ | ✓ | ✓ | ✓ |
| AWS CodeDeploy | ✓ | ✓ | ✓ | ✓ |
| AWS CodePipeline | ✓ | ✓ | ✓ | ✓ |

| Services Offered: | Ireland | Frankfurt | London | Paris |
|------------------------------------|---------|-----------|--------|-------|
| AWS CodeStar | ✓ | ✓ | ✓ | |
| AWS Config | ✓ | ✓ | ✓ | ✓ |
| AWS Database Migration Service | ✓ | ✓ | ✓ | ✓ |
| AWS Data Pipeline | ✓ | | | |
| AWS Device Farm | | | | |
| AWS Direct Connect | ✓ | ✓ | ✓ | ✓ |
| AWS Directory Service | ✓ | ✓ | ✓ | |
| AWS Elastic Beanstalk | ✓ | ✓ | ✓ | ✓ |
| AWS Elemental MediaConvert | ✓ | ✓ | ✓ | |
| AWS Elemental MediaLive | ✓ | | | |
| AWS Elemental MediaPackage | ✓ | ✓ | | ✓ |
| AWS Elemental MediaStore | ✓ | ✓ | | |
| AWS Elemental MediaTailor | ✓ | | | |
| AWS Fargate | ✓ | | | |
| AWS Glue | ✓ | ✓ | | |
| AWS Greengrass | | ✓ | | |
| AWS IoT Core | ✓ | ✓ | ✓ | |
| AWS 1-Click | ✓ | ✓ | ✓ | |
| AWS IoT Analytics | ✓ | | | |
| AWS IoT Device Management | ✓ | ✓ | ✓ | |
| AWS Key Management Service | ✓ | ✓ | ✓ | ✓ |
| AWS Lambda | ✓ | ✓ | ✓ | ✓ |
| AWS Managed Services | ✓ | | ✓ | |
| AWS Marketplace | ✓ | ✓ | ✓ | ✓ |
| AWS Migration Hub | | | | |
| AWS Mobile Hub | ✓ | ✓ | ✓ | |
| AWS OpsWorks Stacks | ✓ | ✓ | ✓ | ✓ |
| AWS OpsWorks for Chef Automate | ✓ | ✓ | | |
| AWS OpsWorks for Puppet Enterprise | ✓ | ✓ | | |
| AWS Personal Health Dashboard | ✓ | ✓ | ✓ | ✓ |

| Services Offered: | Ireland | Frankfurt | London | Paris |
|------------------------------|---------|-----------|--------|-------|
| AWS Secrets Manager | ✓ | ✓ | ✓ | |
| AWS Server Migration Service | ✓ | ✓ | ✓ | ✓ |
| AWS Service Catalog | ✓ | ✓ | ✓ | ✓ |
| AWS Shield Standard | ✓ | ✓ | ✓ | ✓ |
| AWS Shield Advanced | ✓ | ✓ | | |
| AWS Snowball | ✓ | ✓ | ✓ | ✓ |
| AWS Snowball Edge | ✓ | ✓ | ✓ | ✓ |
| AWS Snowmobile | ✓ | ✓ | ✓ | ✓ |
| AWS Step Functions | ✓ | ✓ | ✓ | |
| AWS Storage Gateway | ✓ | ✓ | ✓ | ✓ |
| AWS Sumerian | ✓ | | ✓ | |
| AWS Support | ✓ | ✓ | ✓ | ✓ |
| AWS Trusted Advisor | ✓ | ✓ | ✓ | ✓ |
| AWS WAF | ✓ | ✓ | | |
| AWS Systems Manager | ✓ | ✓ | ✓ | |
| AWS X-Ray | ✓ | ✓ | ✓ | |
| Elastic Load Balancing | ✓ | ✓ | ✓ | ✓ |
| VM Import/Export | ✓ | ✓ | ✓ | ✓ |

117.2 On-Boarding/Off-Boarding Processes and Service Migration

AWS maintains a cadre of Getting Started Guides and schedules regular webinars. These guides and webinars cover a variety of topics, including the two discussed in the following bullets. See <http://aws.amazon.com/documentation/gettingstarted/> for more details.

- **Getting Started with AWS** – This guide provides an introduction to AWS, examples of what you can do with AWS, basic information that you need to know to get started, and links to resources and documentation that will help you learn more no matter what your use case is. The guide is available in [HTML](#), [PDF](#), and [Kindle](#) formats.
- **AWS Management Console** – This guide provides an overview of the AWS Management Console and instructions on starting various services, steps to customising the navigation bar, and tips on accessing region settings, security credentials, and billing updates. The guide is available in [HTML format](#).
- **The G-Cloud 10 On-Boarding Guide** – In addition to the Getting Started Guides, AWS has created a guide exclusively for new Public Sector customers purchasing Services through G-Cloud. This guide provides an overview of the process of opening an account and getting setup, tailored specifically with Public Sector G-Cloud customers in mind. The guide is available in PDF format from your AWS Account Manager upon request.

AWS customers retain control and ownership of their data. Unless we terminate your use of the Service Offerings pursuant to the G-Cloud 10 Framework Agreement or Section 7.2(b) of the [Supplier Terms Document](#), AWS will not erase customer data for 30 days following an account termination. This allows customers to retrieve content from AWS Cloud services so long as the customer has paid any charges for any post-termination use of the service offerings and all other amounts due.

AWS allows customers to move data as needed off AWS storage using the public Internet or AWS Cloud services such as AWS Direct Connect, AWS Import/Export, and more.

With AWS, you can provision compute power, storage, and other resources, gaining access to a suite of elastic IT infrastructure services as your business demands them. With minimal cost and effort, you can move your application to the AWS Cloud and reduce capital expenses, minimise support and administrative costs, and retain

the performance, security, and reliability requirements your business demands. To see a step-by-step migration strategy, refer to the [Migrating Your Existing Applications to the AWS Cloud](#) whitepaper.

117.3 Service Management Details

All AWS Cloud services are driven by robust APIs that allow for a wide variety of monitoring and management tools to integrate easily with your AWS Cloud resources. Common tools from vendors such as Microsoft, VMware, BMC Software, Okta, RightScale, Eucalyptus, CA, Xceedium, Symantec, Racemi, and Dell, to name just a few, already support AWS.

117.3.1 AWS Management Console

The AWS Management Console is a single destination for managing all of your AWS resources, from Amazon EC2 instances to Amazon DynamoDB tables. Use the console to perform any number of tasks, from deploying new applications to monitoring the health of your application. The console enables you to manage all aspects of your AWS account, including accessing your monthly spending by service, managing security credentials, or even setting up new IAM users. The console supports all AWS Regions and lets you provision resources across multiple regions.

117.3.2 AWS Command Line Interface (CLI)

The AWS CLI is a unified tool used to manage your AWS Cloud services. With just one tool to download and configure, you can control multiple AWS Cloud services from the command line and automate them through scripts.

117.3.3 Use Your Existing Management Tools

It is likely that many of the tools that your organisation is using to manage your on-premises environments can be integrated with AWS. Integrating your AWS environment can provide a simpler and quicker path for cloud adoption because your operations team does not need to learn new tools or develop completely new processes.

117.4 Service Levels and Service Credits

AWS currently provides Service Level Agreements (SLAs), with a corresponding Service Credit regime, for several products. Due to the rapidly evolving nature of AWS's product offerings, SLAs are best reviewed directly on our website via the links below:

- Amazon EC2 SLA: <http://aws.amazon.com/ec2-sla/>
- AWS Shield Advanced: <https://aws.amazon.com/shield/sla/>
- Amazon S3 SLA: <http://aws.amazon.com/s3-sla/>
- Amazon CloudFront SLA: <http://aws.amazon.com/cloudfront/sla/>
- Amazon Route 53 SLA: <http://aws.amazon.com/route53/sla/>
- Amazon RDS SLA: <http://aws.amazon.com/rds-sla/>

See the Supplier Terms document affiliated with this framework catalogue for additional information.

117.1 Trial Service Details

The AWS Free Tier is designed to enable you to get hands-on experience with AWS at no charge for 12 months after you sign up. After creating your AWS account, you can use products and services listed at <http://aws.amazon.com/free/> for free within certain usage limits.

117.2 Termination Terms

The following subsection provides information relating to termination.

117.2.1 Process by Consumers

Please see the AWS UK G-Cloud 10 Supplier Terms Document in the Digital Marketplace.

117.3 Data Restoration/Service Migration

Traditional enterprise backup and recovery strategies typically take an agent-based approach whereby the entire contents of a server are backed up over either the LAN or the SAN. Conventional architectures have required this approach because replacing failed components is complex, time consuming, and operationally intensive. This has, in turn, created a backup environment that is complex to manage

and resource-intensive to operate, requiring technologies such as data de-duplication and virtual tape libraries to cope with ever-increasing workloads.

The AWS platform enables a far more lightweight approach to backup and recovery due, in part, to the following characteristics:

- Computers are now virtual abstract resources instantiated via code rather than being hardware-based.
- Capacity is available at incremental cost rather than up-front cost.
- Resource provisioning takes place in minutes, lending itself to real-time configuration.
- Server images are available on demand, can be maintained by an organisation, and can be activated immediately.

These characteristics offer you opportunities to recover deleted or corrupted data with less infrastructure overhead.

117.3.1 Protecting Configurations Rather Than Servers

Amazon EC2 enables the backup and recovery of a standard server, such as a web server or application server, so that you can focus on protecting configuration and stateful data rather than on the server itself. This set of data is much smaller than the aggregate set of server data, which typically includes various application files, operating system files, temporary files, and so on. This change of approach means that regular nightly incremental or weekly full backups can take far less time and consume less storage space.

When a compute instance is started in Amazon EC2, it is based upon an AMI and can also connect to existing storage volumes (e.g., Amazon EBS). In addition, when launching a new instance, it is possible to pass user data to the instance so that it can be accessed internally as dynamic configuration parameters.

A sample workflow would include the following steps:

1. Launch a new instance of a web server, passing it the “identity” of the web server and any security credentials required for initial setup. The instance is based upon a pre-built AMI that contains the operating system and relevant web server application (e.g., Apache, Internet Information Services [IIS]).
2. Upon startup, a boot script accesses a designated and secured Amazon S3 bucket that contains the specified configuration file(s).

3. The configuration file(s) contains various instructions for setting up the server (e.g., web server parameters, locations of related servers, additional software to install, patch updates).
4. The server executes the specified configuration and is ready for service. An open-source tool for performing this process, called cloud-init, is already installed on Amazon Linux AMIs and is also available for a number of other Linux distributions.

Figure 3 depicts a traditional backup approach, and Figure 4 depicts an Amazon EC2 backup approach.

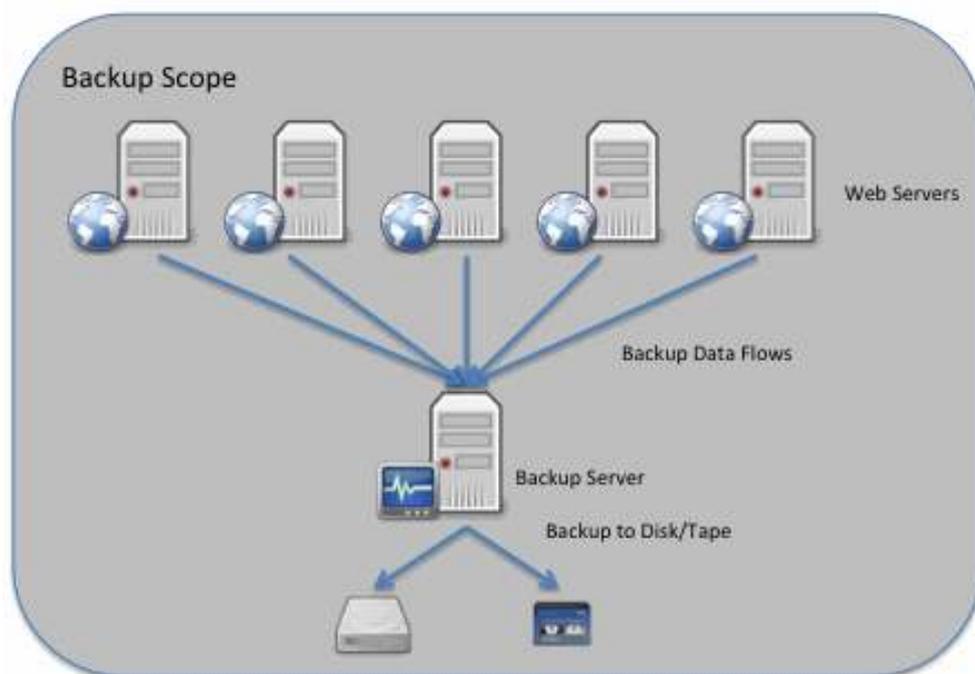


Figure 3 – Traditional Backup Approach

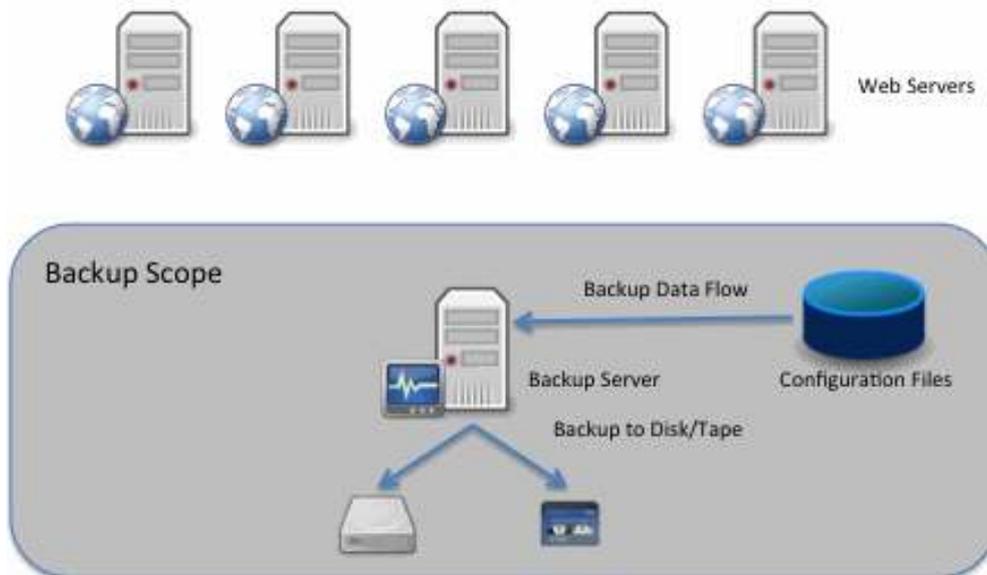


Figure 4 – Amazon EC2 Backup Approach

In the Amazon EC2 backup approach, there is no need to back up the server itself. The relevant configuration is contained in the combination of the AMI and the configuration file(s). Therefore, the only components requiring backup and recovery are the AMI and configuration file(s).

117.3.2 Amazon Machine Images (AMIs)

AMIs that you register are automatically stored in your account using Amazon EBS snapshots. These snapshots reside in Amazon S3 and are highly durable. This means that the underlying storage mechanism for the AMIs is protected from multiple failure scenarios.

It is also possible to share AMIs between separate AWS accounts. Consequently, you can create totally independent copies of the AMI by:

- Sharing the original AMI to another specified AWS account that you control.
- Starting a new instance based upon the shared AMI.
- Creating a new AMI from that running instance.

The new AMI is then stored in the second account and is an independent copy of the original AMI. You also have the option to create multiple copies of the AMI within the same account.

117.3.3 Configuration Files

Customers use a variety of version management approaches for configuration files, and you can follow the same regime for the files used to configure your Amazon EC2 instances. For example, you could store different versions of configuration files in designated locations and securely control them like any other code. You could then back up these code repositories using the appropriate backup cycle (e.g., daily, weekly, monthly) and snapshots to protected locations. Furthermore, you could use Amazon S3 to store your configuration files, taking advantage of the durability of the service in addition to backing up the files to an alternate location on a regular basis.

117.3.4 Database and File Servers

Backing up data for database and file servers differs from the web and application layers. In general, database and file servers contain larger amounts of business data (tens of GB to multiple TBs) that must be retained and protected at all times. In these cases, you can leverage efficient data movement techniques such as snapshots to create backups that are fast, reliable, and space-efficient.

For databases that are built upon Redundant Array of Independent Disks (RAID) sets of Amazon EBS volumes (and have total storage less than 1 TB), an alternative backup approach is to asynchronously replicate data to another database instance built using a single Amazon EBS volume. While the destination Amazon EBS volume will have slower performance, it is not being used for data access and can be easily snapshotted to Amazon S3 using the Amazon EBS snapshot capability.

More information on AWS backup and recovery can be found at <http://aws.amazon.com/backup-storage/>. Refer to the [Storage Options on the AWS Cloud](#) whitepaper for additional information.

117.3.5 Using AWS for Disaster Recovery

AWS provides the features and services that you can leverage for your disaster recovery processes. Using these AWS features, customers (oftentimes assisted by AWS partners) can build and operate agile and cost-effective disaster recovery solutions.

Traditional disaster recovery approaches involve the duplication of infrastructure to ensure the availability of spare capacity in a disaster scenario. This infrastructure needs to be procured, installed, and maintained so that it is ready to deal with the anticipated capacity requirements. Under normal operational circumstances, this infrastructure would typically be underutilised or over-provisioned.

With AWS, you can scale up your infrastructure on an as-needed basis, enabling faster disaster recovery of critical IT systems without incurring the infrastructure expense of a second physical site. This allows greater agility to change and optimise resources during a disaster recovery scenario. It also results in significant cost savings because you only pay for what you use when using the highly scalable, reliable, secure, fast, and inexpensive AWS Cloud infrastructure.

The AWS Cloud supports many popular disaster recovery architectures, from “pilot light” environments that are ready to scale up at a moment’s notice to “hot standby” environments that enable rapid failover. Benefits of using AWS for disaster recovery include:

- **Fast Performance** – Readily access fast, disk-based storage and file retrieval.
- **No Tape** – Eliminate costs associated with transporting, storing, and retrieving tape media and associated tape backup software.
- **Compliance** – Avoid fines for missing compliance deadlines with fast retrieval of files.
- **Elasticity** – Quickly add any amount of data and easily expire and delete data without handling physical media.
- **Secure** – Trust our secure and durable technology platform that maintains industry-recognised certifications and audits.
- **Partners** – Connect with AWS solution providers and system integration partners to help with your deployment.

A business typically decides on an acceptable Recovery Time Objective (RTO) and Recovery Point Objective (RPO) based on the financial impact to the business when systems are unavailable. AWS can work with customers to plan disaster recovery solutions so that they cost-effectively provide system recovery based on the RPO within the timeline and service level established by the RTO. Additional information on using AWS for disaster recovery can be found at <http://aws.amazon.com/disaster-recovery/>. We also recommend reviewing the [Using AWS for Disaster Recovery](#) whitepaper for additional information.

117.3.6 Service Migration

AWS migration services include:

- **AWS Application Discovery Service** – A service that helps you quickly and reliably plan application migration projects by automatically identifying applications running in on-premises data centres, their associated dependencies, and their performance profile.
- **AWS Database Migration Service** – A service that helps you migrate databases to AWS easily and securely. The source database remains fully operational during the migration, minimizing downtime to applications that rely on the database. AWS Database Migration Service can migrate your data to and from most widely used commercial and open-source databases.
- **AWS Server Migration Service** – An agentless service which makes it easier and faster for you to migrate thousands of on-premises workloads to AWS. AWS Server Migration Service allows you to automate, schedule, and track incremental replications of live server volumes, making it easier for you to coordinate large-scale server migrations.
- **AWS Snowball** – A petabyte-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of the AWS Cloud. Using Snowball addresses common challenges with large-scale data transfers including high network costs, long transfer times, and security concerns.

Transferring data with Snowball is simple, fast, secure, and can be as little as one-fifth the cost of high-speed Internet.

- **AWS Snowball Edge** – A 100TB data transfer device with on-board storage and compute capabilities. You can use Snowball Edge to move large amounts of data into and out of AWS, as a temporary storage tier for large local datasets, or to support local workloads in remote or offline locations.

117.4 Customer Responsibilities

As cloud computing customers are building systems on top of cloud infrastructure, the security and compliance responsibilities are shared between the Cloud Service Providers (CSP) and cloud customers. In an Infrastructure as a Service (IaaS) model, customers control how they architect and secure their applications and data put on the infrastructure, while CSPs are responsible for providing services on a highly secure and controlled platform and providing a wide array of additional security features. The level of CSP and customer responsibilities in this shared responsibility model depends on the cloud deployment model (see the [NIST Definition of Cloud Computing](#) models). Customers should be clear as to their responsibilities in each model. AWS's shared responsibility/security model is depicted in **Figure 5** below.

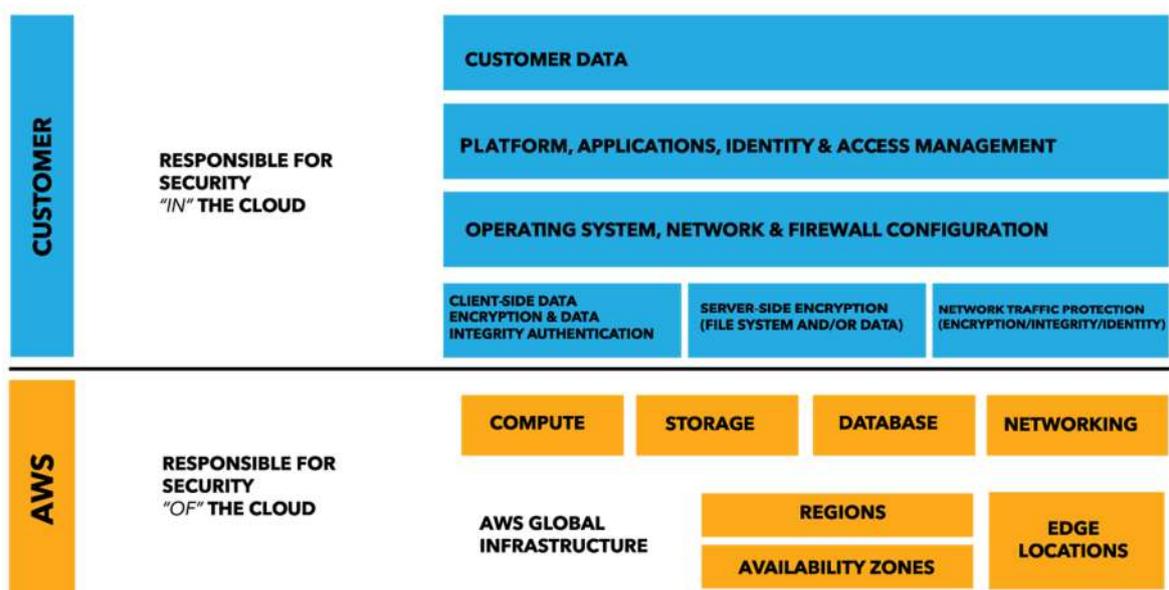


Figure 5 – AWS Shared Responsibility Model

- **AWS Responsibility:** AWS operates, manages, and controls the infrastructure components, from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate.
- **Customer/Partner Responsibility:** Customers/partners assume responsibility and management of the guest operating system (including updates and security patches), other associated application software, configuration of the AWS-provided security group firewalls, and other security, change management, and logging features.

AWS's shared responsibility model is further explained on the [AWS Compliance](#) web page. AWS does not access customer data, and customers are given the choice as to how they store, manage, and protect their data. There are four important basics regarding data ownership and management in the shared responsibility model:

- 1) Customers continue to own their data.
- 2) Customers choose the geographic location(s) in which to store their data—it does not move unless the customer decides to move it.
- 3) Customers can download or delete their data whenever they like.

Customers should consider the sensitivity of their data and decide if and how to encrypt the data while it is in transit

The Server Labs

The Server Labs (TSL) is a 100% privately founded IT Consultancy and Software Development Company with headquarters in the UK and offices in Germany and Spain and now established as a leader in Cloud Computing services. The Server Labs focuses on the design and implementation of IT architectures and advanced software engineering projects working with the most advanced technologies to provide its clients cost-effective, scalable and high performance solutions. The Server Labs has been using the Cloud since 2006 and working with its customers in the cloud since 2008 and was one of the first European partner's of Amazon Web Services.

The Server Labs has clients in many different industry areas such as space, finance and telecoms. We collaborate with our clients to obtain success, committed to innovation, enjoying what we do every day and growing with every challenge.

The Company's mission is:

- To provide **expert services** in the field of IT architectures and advanced software engineering
- To **improve radically the software development process**
- **To help organisations achieve better business results through the correct use of latest technologies**
- To have **100% satisfied clients**
- To create high quality **innovative software solutions**, providing added value to our customers

The specific value, experience and expertise that The Server Labs can provide for e-LfH are:

1. **Technical excellence** and capability to act as lead on architectural decisions and as technology expert in software and system subjects.
2. **Architecture experience at software and system level.**
3. **Proven experience in HPC and Big Data Projects**
4. **Real cloud computing experience, at IaaS, PaaS and SaaS levels, and for both compute power as well as storage solutions in different clouds.**
5. **Quality control based on ISO9001 for all software systems developed**
6. **Technological excellence**, especially in the main technologies required for the project, including HPC, web services and security technologies.



All our architects and engineers are experts with an average of 10 years' experience in the planning, design and development of complex software systems. Our multinational team has been a pioneer in Java technologies, Object Oriented Analysis and Design and distributed architectures, and has the required hands-on experience in many state of the art technologies. In the last few years, The Server Labs has positioned itself as a leader in Cloud computing services, helping organisations move to the Cloud at all levels. For more information on current projects being undertaken by The Server Labs in Cloud computing, please see Appendix B.

Our experience working across several industries has given us a good understanding of the different requirements so we are able to provide the solution that best suits each particular business and reuse the lessons learnt in the other industry sectors when applicable. Our clients span organisations such as Banks (BNP Paribas, BBVA, Caja Madrid) , the European Space Agency (ESA), Madrid Underground (Metro de Madrid), ICCAT, Amadeus, TRAGSA, TIBCO (a leader in Messaging and Service Bus architecture systems), O2 and Telefonica, Vodafone, ORACLE, several, Sun Microsystems, TUI and Marsans travel, etc.

118.0 Selected Customer References

BNP Paribas

Implementation of the Enterprise Development and QA platform

BNP Paribas' department for Risk Systems Development undertook a strategic change programme merging the technical architecture of a number of its key applications to streamline and unify the services they provide internally.

In order to consolidate tools and best practices we customised and implemented our own Enterprise Development Platform which covers the entire application lifecycle from development to ops with its integrated framework.

The implementation of the integrated platform enabled a substantial productivity increase based on a fully controlled, standardised development environment

How we helped

- Integration and customization /Implementation of The Server Labs development and quality control platform for BNP Paribas Risk Systems
- Consolidation of languages and tools
- Controlling and assuring the quality and standardization of development practices.
- Automation of quality and reporting processes.
- Automation of projects office.

Results

- Fully integrated and controlled development environment
- Increased project productivity and quality

Helix Nebula

European Science Cloud

The Server Labs forms part of an exclusive European partnership led by CERN, EMBL and ESA with the aim to establish a federated, sustainable and secure high-performance cloud computing platform.

Supported by industrial partners it will provide stable computing capacities and services that elastically meet demand.

How we helped

- We provide our technical expertise based on the successful development of complex Cloud architectures as well as Grid processing in the Cloud both in science and industry environments
- Development of the EC2 bridge in the Helix Nebula BB

Participants

- Consortium formed by: CERN, EMBL and ESA, Atos, Capgemini, CloudSigma, Interoute, Logica, Orange Business Services, SAP, SixSq, Telefonica, Terradue, Thales, The Server Labs and T-Systems, along with the Cloud Security Alliance, the OpenNebula Project and the European Grid Infrastructure (EGI.eu).
- Flagship use cases for testing and deployment of the science Cloud:
 - Cern: HPC Processing of LHC data
 - ESA: Supersites Earthquake data
 - EMBL: Genomic processing

Results of the initiative

- Successful deployment of the science Cloud for flagship projects

ESA

Gaia project

Moving the Gaia project's data processing to the cloud provided ESA with savings of around 50% compared to using in-house hardware. At the same time it provides a level of scalability which means that the work will be notably accelerated.

ESA's ambitious Gaia project aims to create a three-dimensional map of unprecedented size and precision charting the composition, formation and evolution of over one billion stars (around 1% of our Galaxy)

We have provided state of the art architectural solutions for different areas of the project including data management, High Performance Computing, creating a development platform for all the international teams involved and project management.

How we helped

- Design of a Cloud-based solution for the data processing reducing its cost by approximately 50% whilst increasing the flexibility of the testing activities
- Collaboration in the implementation of a distributed computing framework to improve performance.
- Deployment of a development platform that allows the different teams to work in a distributed environment enabling release management, automated testing, continuous integration and quality assurance.
- Database tuning.

Results

- Substantial cost reduction: the developed solution provides an estimated reduction of 50% of Total cost of ownership (TCO)
- Higher performance: Improved performance in the distributed framework being 6x, 7x times faster (in different areas of the processing)
- Productivity boost due to our development platform.

"The Gaia AGIS Cloud experiment has been very successful for us. It indicates that bringing the data processing to The Cloud can provide us with savings of up 50% compared to using in house hardware. An additional advantage is that it gives us the ability to scale to far more processors that we could have in house which means essentially that we can finish the job sooner"

William O'Mullane, Gaia Science Operations Development Manager, European Space Agency

Eumetsat

Technology assessment for the MTG programme

When Eumetsat was planning for a near real-time data processing function for its Meteosat Third generation programme (MTG) it became apparent that this would require an unprecedented level of computing power. Its on-demand re-processing of batch data pushed compute requirements to a speed up to 30 times faster than real time with much higher data volumes than seen in the existing systems.

We have led the feasibility study and the design of prototypes to evaluate grid, cloud and supercomputing architectures to determine the most suitable option and right technologies for the MTG programme to thus enable a long-term reduction of costs.

How we helped

- Technical assessment and cost evaluation of current technologies for Eumetsat's HPC/Big Data needs
- Development of the prototypes
- Evaluation of results and recommendation

Results

- Recommendation to create an architecture based on a hybrid cloud

International Commission for the Conservation of Atlantic Tunas (ICCAT)

Electronic Bluefin Tuna Catch Documentation

The development of an Electronic Bluefin Tuna Catch Documentation system (eBCD) enables the complete tracking from catch to end-market distribution of all Atlantic bluefin tuna operations.

The new electronic system provides greater accuracy and traceability via real-time information helping to more effectively detect and prevent illegal, unreported and unregulated fishing.

Together with Tragsa we are part of the project team to design, develop and implement an Electronic Bluefin Tuna Catch Documentation system (eBCD) for the International Commission for the Conservation of Atlantic Tunas (ICCAT) replacing the previously used manual and paper-based process with a real-time information system.

The solution will enable the complete electronic operability of the catch documentation programme, tracking from catch to end-market distribution all Atlantic bluefin tuna operations.

It will maximize use and performance for end users, providing maximum security and reliability to the system.

How we helped

- Architecture, design and implementation of the eBCD
- Application quality assurance
- Cloud hosting and infrastructure
- User training and maintenance support

Results

- Management of 5000 eBCD/year
- Support for more than 100 concurrent users
- High speed connectivity between service provider and ICCAT Secretariat

Amadeus

Implementation of the TSL Continuous Integration Platform

Amadeus' Internal Information Systems department decided to implement the TSL Continuous Integration platform into their development process in order to optimise the software development cycles and improve development quality, control and automation of the delivered applications.

The Internal Information Systems department works with many developers on a variety of projects. Therefore a standardised platform was vital in order to efficiently reassign team members to activities according to their availability ensuring that neither results nor schedules are compromised. The development platform allows constant tracking of project development health and quality ensuring that simultaneous tasks and projects are running smoothly.

How we helped

- Implementation of the TSL Continuous Integration platform
- Creation of a common framework providing clear mechanisms, standards, tools and conventions enabling the effective execution of both internal as well as external outsourced projects
- Optimization of the current development lifecycle as well as overall development quality, control and automation of the delivered applications
- Seamless transition to Agile development technologies

Results

- Reduction of error prone conditions
- Constant measurability of project health

“With the support of *The Server Labs (TSL)* we have recently undertaken a development platform and processes optimisation project for the development of our internal applications in Madrid. The project outcome has been more than satisfactory.

In addition to implementing effective quality control, *TSL* has helped us to create a framework where best practices are formalised and enforced through a Continuous Integration and Release Management platform.

Because we work with many developers on a variety of projects it is crucial to have a standardised platform in place so team members can be reassigned to activities according to their availability. This way results and schedules are not compromised. We track project development health and quality constantly which is key to keeping simultaneous tasks and projects running smoothly.”

Luis-Fernando Gonzalez

Web Solutions Manager, Group Internal Information Systems

Amadeus IT Group SA

CNIO

Migrating the data processing to the Cloud

The Spanish national cancer research organisation CNIO experienced that its vast computing requirements increasingly led to processing bottlenecks causing delays in sequencing projects and research programmes.

The Server Labs have worked with CNIO to help them overcome the bottleneck of data processing in its research programmes.

In an initial feasibility study, The Server Labs established how to transfer CNIO's computing requirements to the Cloud, aiming to make the processing more agile whilst substantially bringing down cost and reducing the need for in-house infrastructure.

How we helped

- Assessment to transfer CNIO's computing requirements to the Cloud enabling agile processing whilst reducing the need for in-house infrastructure
- Created a framework which incorporates all of the strengths of the cloud, in particular data durability, publishing mechanisms and audit trails to make results of experiment reproducible
- Designed and delivered data transfer

Results

- Development of a robust cloud-based platform to perform on-demand genomic processing tasks (at the same time enabling experiment results to be more easily reproduced, stored and published.)

IT Cost reduction of ca. 30% through Cloud-based processing

BBVA

1. Cloud Security

BBVA needed to develop a new secure cloud model able to keep confidential data safe whilst providing a highly productive platform based on Openstack and SDN for the development of internal applications. The new private cloud environment uses continuous delivery with automated security mechanisms.

In addition to substantially reducing the risk of security breaches based on web services the automation of the most important security operations tasks are now based on SDN solutions allowing a reduction of development time by over 20-30%.

How we helped

- Design of the project architecture and definition of the security platform and API services
- Deployment of the security platform and all operational processes
- Development of security policies
- Defining Best Practices for security

Results

- Automation of the most important security tasks based on SDN solutions
- Security policy is already applied by default decreasing develop time by over 20-30%
- Quick response time when facing changes in the integration processes
- Reducing the risk of security incidents based on web services

2. Corporate Security Application

The Server Labs (TSL) has developed and implemented an entirely new version of BBVA's corporate security application Faro. The application which has won a national award for the best application in the area of security has been unified and streamlined throughout all countries in which BBVA operates making its management and maintenance more flexible, transparent and cost efficient.

The new flexible, efficient and scalable application supports the management of all security functions of the BBVA group including for example dashboards, financial control, incident management, inventories and security personnel.

How we helped

- Development and implementation of a new and revised version of BBVA's global corporate security application Faro
- Deployment
- Designing the modular, multi country architecture
- Coordinating the multiple actors (Client, data centre, development team) for the go-live

- Deployment of TEST and development environments in the cloud

Results

- Flexible, transparent and cost efficient application management and maintenance
- Highly agile application
- Cost reduction - estimated 20% decrease of FTE in the corporate security area

Working with The Server Labs has been a very positive experience. This has also been demonstrated by the success and public recognition of our latest programme, the development of a unique technology platform for our global corporate security management.

The Server Labs has been instrumental in the design and development of its scalable and streamlined architecture enabling the worldwide rollout.

The Server Labs really stood out with its superior technical knowledge on all levels, its responsiveness and agility, and acting as our right hand throughout the entire project.

Inés Díaz Ochagavía
Director Production Corporate Security
BBVA

3. Continuous Integration Platform

With over 50 systems exchanging financial information using more than 2200 interfaces, maintenance costs and operational risk are high. BBVA wanted to significantly reduce the number of interfaces, unify the data format and thus be able to lower maintenance costs, quickly respond to changes, and centralize monitoring.

BBVA designed its new continuous integration and quality environment employing The Server Labs Enterprise Development Platform in the Cloud.

The development platform does not only control quality and improve productivity but is also instrumental enabling homogenized and standardised developments for all of BBVA's development projects, and as such providing a corporate reference architecture.

How we helped

- Supporting project organization and management
- Simplification of existing processes
- Development of integration processes, using a common framework.
- Design and development of common components, allowing later reutilization in other projects
- Design and development of auxiliary tools (maintenance, monitoring of existing processes)

Results

- Fast response to changes in the integration processes.
- Reduction of incidents in the integration software in operational environments.
- Standardization of developments.

4. Technical audit of Tugestionline

The Server Labs provided a technical audit driving the design and execution of a series of tests to boost the performance of BBVA's TuGestionLine. The web application, constantly growing in users, provides online accounting counselling for businesses and individuals.

How we helped

- Full analysis of the system architecture
- Design and execution of a series of tests to provide an accurate diagnostic about the web application's performance, quality and reliability
- Identification and quantification of possible improvement points to optimize performance, system stability and ability to deal with load peaks
- Load and capacity tests

Results

- Delivered a full performance analysis with a complete set of code fixes, reconfigurations and recommendations in to improve functionalities
- Web load time was substantially reduced tangibly improving the user experience

Webtools

Migrating e-encuesta.com to the Cloud

The growing business of Webtools' leading online survey application e-encuesta.com with its pronounced peak-times was increasingly causing server overloads limiting the number of clients that could be served simultaneously.

In order to give the application the ability to flexibly react to changing demand it was necessary to change its data centre-based system to a Cloud-based structure.

We successfully adapted the architectural base and executed the smooth migration of e-encuesta.com from a traditional data processing centre to the public Cloud using The Server Labs Rapid Cloud Transition as its central tool.

How we helped

- Designing a robust and efficient cloud architecture specific for the client's project.
- Security upgrades and enhancements to the administration of application.
- Development of a self-healing architecture, allowing that when problems are detected the architecture is able to react and reset the environment
- Provide a pre-emptive support to the client in order to make sure the application is evolving properly in time and the system is up to date.

Results

- Substantial reduction of operating cost
- Service now fully aligned to e-encuesta's needs
- Resources elastically aligned with e-encuesta's usage
- Improved performance and user experience: much higher response time and uptime
- Almost unlimited scalability supports its business in its continuous growth
- Increased reliability and stability of the environment

The Server Labs

Founded in 2004, THE SERVER LABS is specialized in the design and agile deployment of technology solutions in the area of enterprise integration, software engineering, cloud computing architectures and Big Data. Through our offices in the UK, Spain and Germany we offer our customers cutting-edge, cost-effective, scalable and high performance solutions enabling them to stay competitive with the most suitable and innovative IT technologies and solutions.

The list of our customers includes renowned organisations such as the European Space Agency, leading banks in Europe and Government bodies. Our services span from early strategic technology consulting and architecture services to the full implementation and agile deployment solutions.

More information about The Server Labs is available on www.theserverlabs.com

Metro de Madrid

1. Enhancement of real-time driver information system

Metro de Madrid required a set of new functionalities and developments to enhance the efficiency and reliability of ARCO, the real-time information system for its train drivers.

The improved design of ARCO brings scalability to the system through a new component, allowing growth in both lines and functionalities, and increasing the system's reliability. At the same time it helps resolve capacity issues, bringing down use of CPU for the exchanged components, freeing up valuable compute power for other systems.

Delivered on-time and successfully installed on a live environment, we provided an iterative delivery solution, based on Agile methodologies, developing and releasing the new functionalities on a progressive basis.

We aimed to improve the overall software quality with each new developed component, as well as by refactoring the ones to be modified.

How we helped

- Development of an iterative delivery solution based on Agile methodologies for progressive development and release of the new functionalities
- Improvement of the overall software quality through new development components and refactoring of those to be modified

Results

- Successfully delivered a new version of ARCO, packaging items of several software layers
- Notable improvement of the overall project management initiating the use of collaborative tools like Jira

2. Code Audits

We performed a number of code audits on applications developed by Metro in order to provide visibility of the quality of the code thus enabling informed decisions about the future courses of the respective projects.

How we helped

- Comprehensive analysis of the applications in question and quality of their code
- Flagging up potential problems and areas of improvement

Results

- Metro gained a complete overview of the code quality and was able to make information-based decisions on the further project development
- In some instance the analysis enabled a strong position for renegotiation with previous providers the improvement of code quality
- Documentation of alternative solutions
- Overall reduction of code decreasing the potential for errors

Unit 4

Scalability certification of the Unit4 Health application

In order to present their Health application to medium and large hospitals Unit4 wanted to validate its performance and sustainability with a substantial load of users.

After creating a virtual cloud environment we ran tests simulating a real life load with up to 4,000 concurrent users injecting 4 new users every second.

The simulations confirmed that Unit4's Health application was performing well above average results achieved by reference hospitals.

How we helped

- Creating a virtual environment using Amazon Web Services
- Planning, set up and execution of simulations
- Application and Database fine tuning, and recommendations for optimal configuration.

Results

- Various simulations were successfully concluded confirming the application's ability to handle a large number of concurrent users surpassing the results of reference hospitals.