MURF.AI

Security at Murf

AICPA SOC aicpa.org/soc4so SOC for Service Organizations | Service Organizations

ISO 27001:2013 BQSR IAS ACCREDITED CERTIFIED

GDPR COMPLIANT

At Murf, we value the trust our customers place in us. We recognise the critical importance of data protection and leave no stone unturned to ensure the safety and privacy of our users' information. We believe in being transparent about our data management practices and are committed to maintaining the highest standards of data security at all times.

Hosting

We host Murf services in Amazon Web Services (AWS) servers in the us-east-2 region located in Ohio in the United States.

AWS data centres are designed to be highly resilient and reliable. AWS has a number of measures in place to ensure the availability and durability of its services, such as:

1. Redundancy: Multiple layers of redundancy built in to help ensure that services remain available even in the event of a failure.
2. Environmental controls: Equipped with controls such as temperature and humidity control to help ensure the optimal operating conditions for the servers.
3. Network infrastructure: Have a highly redundant network infrastructure that is designed to be resilient to failures.

AWS takes security seriously and provides a number of features and services to help ensure the security of your data and systems. Some of the measures that AWS takes to ensure security include:

1. Physical security of data centres
2. Encryption of data in transit and at rest
3. Secure network architecture
4. Identity and access management (IAM) controls
5. Compliance with various industry standards and regulations. If you’d like to learn more about AWS security practices, please check out these links:
    - ISO Global certification
    - Security processes
    - SOC compliance

- Customer case studies

We use many AWS provided solutions for our compute & storage, such as - EC2, SQS, DynamoDB, S3, Cloudfront, API Gateway, Lambda. We use AWS's identity management for providing access to services based on least privileged access. We have our ML models running in the virtual private cloud. Download Murf's Network Diagram.

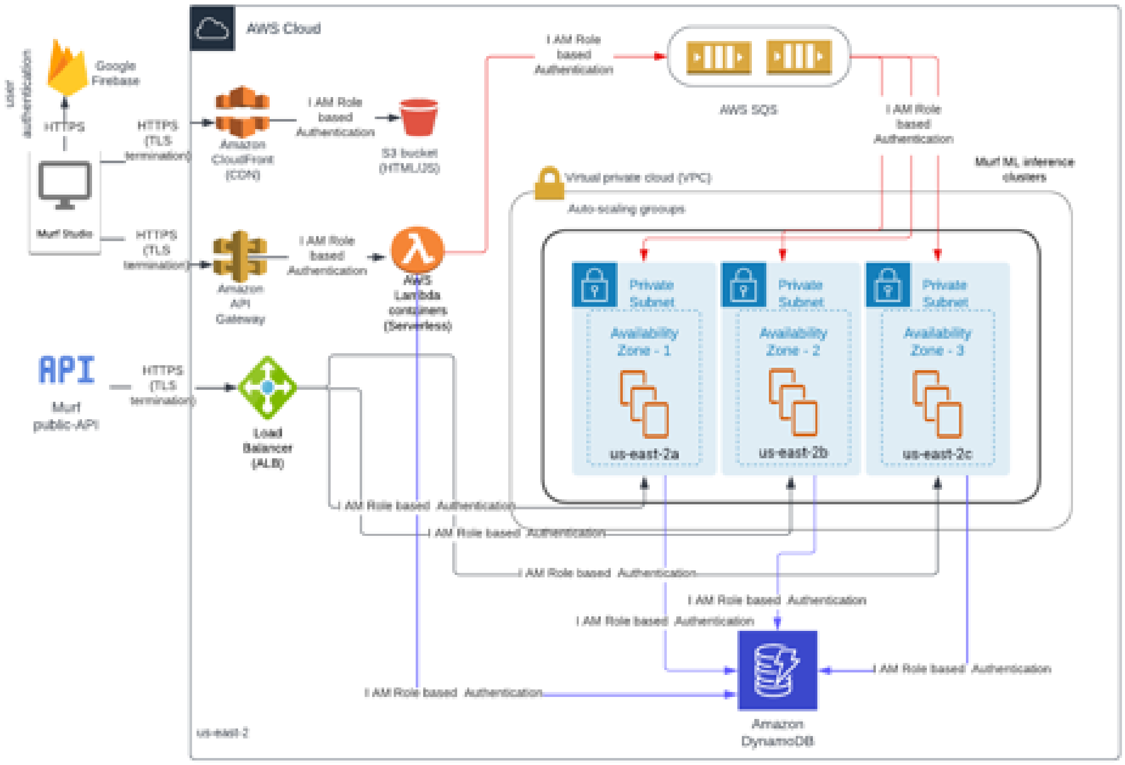
AWS Cloud
user authentication
Google Firebase
HTTPS
HTTPS (TLS termination)
Amazon CloudFront (CDN)
I AM Role based Authentication
S3 bucket (HTML/JS)
Murf Studio
HTTPS (TLS termination)
Amazon API Gateway
I AM Role based Authentication
AWS Lambda containers (Serverless)
I AM Role based Authentication
AWS SQS
I AM Role based Authentication
Murf ML inference clusters
Virtual private cloud (VPC)
Auto-scaling groups
Private Subnet
Availability Zone - 1
us-east-2a
Private Subnet
Availability Zone - 2
us-east-2b
Private Subnet
Availability Zone - 3
us-east-2c
API
Murf public-API
HTTPS (TLS termination)
Load Balancer (ALB)
I AM Role based Authentication
I AM Role based Authentication
I AM Role based Authentication
I AM Role based Authentication
I AM Role based Authentication
I AM Role based Authentication
I AM Role based Authentication
Amazon DynamoDB
us-east-2

Architecture components

1. Murf UI

This is the Murf website (https://murf.ai/) via which the customers interact with Murf. For user authentication we leverage Google's Firebase (https://firebase.google.com/docs/auth) and Google cloud identity for Single sign-on. All the communication to Murf backend is made via HTTPs (with TLS 1.2) and signed via AWS Sigv4. The UI assets are stored in AWS S3 (private endpoint) and served via AWS Cloudfront (CDN) to serve content globally.

2. Murf Services & Middleware

These are a group of services exposing APIs to process & orchestrate client requests, authentication and audio/video/image processing. We use AWS API gateway and AWS Application load balancer as a reverse proxy to our private clusters. Inter-machine communication is made via HTTPs (with TLS 1.2) and signed via AWS Sigv4. It uses AWS IAM for authorization. We leverage AWS SQS (simple queue service) for asynchronous communication and also to have good resiliency and differentiated scaling. We host various types of auto-scaled clusters of EC2 instances for various types of workloads - varying in terms of compute & storage. These clusters have dynamic auto-scaling enabled based on the current & predicted load to achieve maximum availability and performance. All of these are part of our virtual private cloud (VPC) and are replicated 3 times across 3 availability zones in us-east-2. Replication makes our systems resilient to zone failures. We leverage Stripe for payments (https://stripe.com/)

3. Murf ML inference

This is the part responsible for running our ML models for speech synthesis. We have clusters of EC2 machines of high configurations to run ML models at scale. The compute is replicated 3 times across 3 availability zones for better resiliency and has dynamic auto-scaling enabled.

4. Murf Storage

These are a group of services exposing APIs to process & orchestrate client requests, authentication and audio/video/image processing. We use AWS API gateway and AWS Application load balancer as a reverse proxy to our private clusters. Inter-machine communication is made via HTTPs (with TLS 1.2) and signed via AWS Sigv4. It uses AWS IAM for authorization. We leverage AWS SQS (simple queue service) for asynchronous communication and also to have good resiliency and differentiated scaling. We host various types of auto-scaled clusters of EC2 instances for various types of workloads - varying in terms of compute & storage. These clusters have dynamic auto-scaling enabled based on the current & predicted load to achieve maximum availability and performance. All of these are part of our virtual private cloud (VPC) and are replicated 3 times across 3 availability zones in us-east-2. Replication makes our systems resilient to zone failures. We leverage Stripe for payments (https://stripe.com/)

Software quality assurance

We design and develop all software in-house using AWS CodeCommit as our source code repository. We have established our Software Development Life Cycle (SDLC) process as below, in chronological order.

- **Prioritise**

  Chalk out the deliverables and functional requirements

- **Design**

  Cover aspects such as scale, infrastructure, security and instrumentation.

- **Develop**

  Local development with test cases.

- **Code review**

  Code review is done by some other member in the development team who is normally familiar with the development area. If not approved we move to #3.

- **QA testing**

  Post code review approval, the code is merged into a development branch and hosted in our development environment where the QA team performs functional & non-functional testing.

- **Release**

  Release is done by dedicated release managers who are usually the tech leads. The code is merged to our prod branch and deployed via AWS codePipeline to our production environment. Post deployment, our on-call monitors the health of the system via metrics and alarms.

User authentication

While signing up with Murf, a user can use their google account or provide email id. For email id sign up we send a verification link for users to verify and login into their account. We also enable enterprise customers to do single sign on (SSO) using their identity providers.

Murf doesn’t store user passwords. We rely on a third party service, Google Firebase, as our authentication provider. We use a secure firebase SDK which sends user credentials to the firebase server and returns a JWT token. Murf servers use this JWT token to determine the user identifier. The communication between Murf UI and Firebase server is all TLS 1.2 encrypted. For enterprise customers, Murf supports SSO with SAML 2.0 integrating with major identity providers. SAML is an open standard for exchanging authentication and authorisation data between a SAML IdP and SAML service providers. Murf leverages Google Cloud identity. More information on how it works visit - https://cloud.google.com/architecture/identity/single-sign-on

Certification

SOC 2 Type 1 Certification

We have undergone the rigorous SOC 2 Type 1 examination—a procedure designed to ensure that service providers can securely manage data to protect the interests and privacy of their clients—and have demonstrated that we have the necessary controls and structures in place to meet the high-security standards expected by our customers. Our security processes have been independently inspected and validated, confirming that we adhere to the trust services criteria set by the American Institute of Certified Public Accountants.

ISO 27001 Certification

ISO/IEC 27001 is an international standard to manage information security. Murf is ISO 27001 certified which signifies a best-practice approach in managing information security by addressing people, processes, and technology

GDPR

We value your privacy and your rights as a data subject and process all personal data in accordance with the principles of GDPR. We have also appointed local partners as our privacy representative and your point of contact. If you want to contact us via our representative, Prighter or make use of your data subject rights, please visit the following website. https://prighter.com/q/14532374752

Data protection & privacy

Murf is committed to treat information of employees, customers, stakeholders and other interested parties with the utmost care and confidentiality.

- We have established a privacy policy that documents and effectively communicates to our users the amount and type of personal data collected, the obligations of the company, the rights of individuals to access, modify, or delete their personal information, and an up-to-date contact point to direct their inquiries, appeals, or concerns.
- Your project information remains confidential, with the exception of instances where you specifically ask us to look at a project for the purpose of customer service.
- When you sign up to our service and set up an account, we collect the name and email address information. When you make a purchase or attempt to make a purchase through the Site, we collect name, email address and payment or credit card information (we don't have access to your entire credit card number or CVC as this is handled by our payment gateway - Stripe).
- We only work with third-party suppliers that have strict data protection policies and meet the data privacy conditions described in our Privacy Policy. Our privacy policy can be viewed here https://murf.ai/resources/privacy_policy/

Application security

Our system employs encryption in rest and transit using 256 bit TLS. Inter machine communication happens via IAM policies. Murf developers have role based access to the system according to the principle of least privilege. Murf partners with external penetration testing vendors to conduct annual tests. We rely on AWS Guardrail for vulnerability scans. Servers are auto-patched via AWS patch manager.

Resiliency & recovery

Our compute clusters are replicated 3 times over 3 availability zones in AWS cloud. All our databases have point-in-time-recovery (PITR) enabled with periodic backups for disaster recovery.

We have dynamic auto scaling enabled with asynchronous batching to provide high availability in case of sudden traffic spikes. Our public endpoints are inside an AWS WAF (Web application firewall) to protect from DDoS attacks.

To ensure operational effectiveness, we publish relevant metrics and alarms with each release. We have a weekly on-call process with an escalation matrix to address any production issues.

Hiring & training

We have a comprehensive hiring process which involves multiple rounds of interactions. This helps us in raising the bar of our software products. We perform background checks of the candidates prior to employment. Our development team follows OWASP secure coding standards and any new joinee undergoes relevant training. All our employees go through regular security training and adhere to our Information security policies.