LumenVox®
by capacity

Regulatory **Best Practice Recommendations for GDPR / PCI / HIPAA** in On-Premise Installations

Executive Summary:

This paper describes best practices for securing sensitive data when working with LumenVox on-premise software and outlines the extensive functionality that our LumenVox on-prem software makes available for protecting sensitive data.

Audience:

This white paper is a technical document that is written for developers and DevSecOps personnel. It assumes some experience with speech & voice biometric technologies, including familiarity with basic speech recognition grammars and the Speech Semantic Markup Language (SSML) for TTS as well as voiceprints and other information used by Voice Biometrics applications.

Caveats:

This white paper pertains to LumenVox's on-premise licensed software; because of the nature of cloud computing, different considerations are applicable to our SaaS-licensed software. This information is not legal advice. While we do our best to provide useful information to use as a starting point, LumenVox advises all customers to obtain professional legal advice to ensure that their businesses operate in full compliance with all applicable laws. It is the responsibility of customers licensing, deploying and using LumenVox software to perform their own due diligence and use the information contained in this paper as a resource only.

Introduction

Whenever developing any application, not just applications using LumenVox technologies, it is important to be aware of potentially sensitive data that may be processed by the application, such as credit card information, personally identifiable information ("PII"), protected health information, personal data or sensitive personal data. For purposes of this white paper, we will refer to "sensitive data", but it is important to understand that the various regulatory schemes have different definitions for categories of data and different regulatory approaches.

LumenVox on-premises software is used in a large number of diverse applications, from simple demonstrations to integrated banking applications, and as such, there is the possibility that sensitive data may pass through LumenVox software. It is essentially impossible for the software itself to know whether any data should be considered sensitive or not, so the customer application developer and DevSecOps personnel have a responsibility to understand what the risks are and determine whether data being processed could be considered sensitive, and more importantly, take measures to limit exposure of that data.

Developers can utilize some of the built in features of LumenVox software to help customers comply with applicable law and regulatory schemes. LumenVox, working in conjunction with some of our major clients and partners have developed a number of features that assist in securing or suppressing sensitive data.

How can LumenVox Help

This white paper was created to provide our users that are concerned with PCI, HIPAA or GDPR compliance with information that may be useful when performing a security evaluation relating to the use of LumenVox on-prem software and to assist in protecting potentially sensitive data.

This paper provides guidance in the following ways:

1. **Identify areas of potential risk**

   Throughout the remainder of this paper, various areas of potential risk are identified allowing users, application designers and IT managers to better understand where potentially sensitive data may be exposed when using LumenVox software.

2. **Provide details of any risk mitigation measures that can be taken**

   Along with each potential risk that is identified, measures embedded in our software that can be taken to limit or prevent exposure of potentially sensitive data are also described.

In general, if the application developer knows that data could be potentially sensitive for an ASR, TTS or Voice Biometric interaction, then the LumenVox software can be configured to enable various security features to protect that sensitive data. Configuration settings can be specified at startup that enforce sensitive data protection functionality by default, and also disable the logging or recording of events as well as suppressing potentially sensitive data. This level of protection may be too limiting in certain situations, since it may prevent diagnostic reporting, but being aware of this option is important if your primary goal is data security.

Data Discovery

Determining where potentially sensitive data may be located is crucial when analyzing security risks associated with any software. Specifically, a number of areas have been identified where potentially sensitive data may be found, which this guide will attempt to describe in as clear a manner as possible in the following sections.

Note that in most respects, the Call Progress Analysis (CPA) functionality may be considered an automatic speech recognition (ASR) resource for the purposes of this document, so references to ASR resources also cover the CPA resource type as well.

RTP Stream Data (ASR Resources)

As with any other network communication traffic, when receiving DTMF data over an RTP stream, the RTP-Events themselves may contain the numbers of a credit card or other sensitive information when viewed or recorded with a network protocol analyzer utility.

In addition to DTMF events, voice audio is generally used with ASR resources which may be contained within a similar or the same RTP stream and may similarly contain the numbers of credit cards or other sensitive information when accessed with a network protocol analyzer utility.

Risk Mitigation:

LumenVox software cannot prevent exposure of any data that is being sent in an unencrypted manner across the network. LumenVox software supports and recommends the use of optional encryption to secure traffic, including audio/DTMF (RTP) streams as well as session control streams (RTSP / SIP / MRCP). Please review the various secure connectivity options when designing your system architecture with security in mind.

RTP Stream Data (TTS Resources)

TTS audio being sent to an RTP stream from LumenVox software may contain the numbers of credit cards or other sensitive information when accessed with a network protocol analyzer utility.

There are no DTMF-Events generated by the LumenVox software that are sent to an RTP Stream.

Risk Mitigation:

LumenVox software cannot prevent exposure of any data that is being sent in an unencrypted manner across the network. LumenVox software supports and recommends the use of optional encryption to secure traffic, including audio/DTMF (RTP) streams as well as session control streams (RTSP / SIP / MRCP). Please review the various secure connectivity options when designing your system architecture with security in mind.

RTP Stream Data (Voice biometrics)

As with any other network communication traffic, when receiving audio data over an RTP stream, the audio may contain the numbers of a credit card or other sensitive information when recorded with a network protocol analyzer utility.

Risk Mitigation:

LumenVox software cannot prevent exposure of any data that is being sent in an unencrypted manner across the network. LumenVox software supports and recommends the use of optional encryption to secure traffic, including audio/DTMF (RTP) streams as well as session control streams (RTSP / SIP / MRCP). Please review the various secure connectivity options when designing your system architecture with security in mind.

Grammar Data

Grammars are often being fetched across network resources and may also contain potentially sensitive information.

Risk Mitigation:

LumenVox gRPC and MRCP interfaces are extremely flexible in the URI references that may be used to specify grammar locations, for example remote HTTP/HTTPS URIs. Storing sensitive grammars on a HTTPS server will offer some protection as it is traversing the network. We do not recommend the use of unencrypted HTTP URI references for grammars. By using TLS secured HTTP URI references, this offers the best protection for the potentially sensitive information that may be within grammars being used.

SSML Data

Text to be synthesized for a TTS resource may be either plain text or a markup language (SSML) used to express more clearly the way in which things should be pronounced.

Many of the LumenVox gRPC and MRCP functions allow users to specify URIs to use when fetching SSML data being used for speech synthesis using a TTS resource. This SSML data may contain potentially sensitive data, and since this is outside of the LumenVox code, it may be viewed using network analysis tools.

Risk Mitigation:

Storing this SSML data (containing the desired SSML content) on web servers utilizing a TLS secured protocol such as HTTPS may be used to protect the data as it traverses the network.

Log Data (API Users)

Logging information from containerized services are recorded centrally within the Kubernetes platform, in JSON format.

Logging within a Kubernetes environment can be configured for various levels of verbosity. The default verbosity level is Warning, which records minimal information relating to any errors or warnings encountered. Setting higher levels of verbosity will log more information, including diagnostic information at higher verbosity settings.

Since these logs are recording activity within various parts of LumenVox software, it is important to understand that sensitive information may be contained within these logs, such as the result of ASR recognition, or the raw text to be synthesized by a TTS resource.

> **Risk Mitigation:**
>
> There are a number of things that can be done to prevent potentially sensitive data from being recorded via logging. Since Kubernetes by design requires RBAC, this is one of the main ways to protect access to log information from the software, and should be configured appropriately by your security team. Another option may be to limit the verbosity setting to a minimal setting (e.g. Warning), which will only record errors and warnings, which do not contain any sensitive data.

In addition to these options, there are specific API functions that are designed to force a specific ASR or TTS resource into a “secure_context” mode, where potentially sensitive information will be suppressed from the logs, replacing potentially sensitive data simply with the word “_SUPPRESSED”.

Note that this “secure_context” functionality can be specified at an interaction level, allowing fine-grained control over each interaction.

When dealing with an ASR resource, it can be tasked to process both speech and DTMF requests, and when operating in secure_context mode, these DTMF events will be suppressed in the same fashion.

Global configuration settings are also available to allow this secure mode to be enabled by default if this is desired.

Log Data (MRCP-API Users)

The LumenVox MRCP-API can be considered a wrapper around the gRPC API for ASR, TTS and other resources. In addition to the API functionality, this service provides additional connectivity to support MRCPv1 and MRCPv2 clients and their respective media streams and control sessions. This service therefore uses many of the same API logs described above, and uses the same “secure_context” settings in addition to some others specific to MRCP.

The MRCP-API has its own set of additional log files that are created when operating. Verbosity of these logs can be controlled by the similar configuration settings as used for the API.

In addition to API logging information, the MRCP-API can log inbound and outbound SIP / RTSP / MRCP / RTP-Event traffic, depending on its logging verbosity setting. If configured for minimal verbosity, only errors and warnings will be recorded, however with more verbose settings, the information passed over the network within the SIP / RTSP / MRCP / RTP-Event traffic may also be logged, which could contain potentially sensitive data.

Some specific areas where sensitive data may be logged include the logging of RTP-Events corresponding to DTMF keys, which could be used to expose potentially sensitive information. Similarly, the logging of recognition results from either speech or DTMF activity in the logs may also expose potentially sensitive data.

Risk Mitigation:

To avoid logging sensitive data, select a minimal verbosity setting to ensure that only errors and warning messages are logged, which should not pose a security risk.

At any logging verbosity setting, the "secure_context" configuration may be used to suppress logging of potentially sensitive data, in the same way as is described above for non-MRCP-API users. Since MRCP-API users do not have direct access to the API functions allowing the modification of "secure_context" settings on a per-interaction basis, when running the MRCP-API, there is additional functionality provided to allow these settings to be controlled using **Vendor-Specific-Parameters.** See the Vendor Specific Parameters knowledge base article for more details.

MRCP Specific Audio Recording

MRCP save-waveform functionality may be used to store (RTP) audio sent into the ASR resource. The MRCP-API can be configured to store these waveforms to files on disk, potentially exposing sensitive information via the files on disk.

Risk Mitigation:

The save-waveform functionality in MRCP needs to be explicitly configured and also activated in order to be used. This means that for potentially sensitive data to be exposed, the MRCP-API configuration needs to specify a target location for these waveforms. By default this is not set, which effectively disables the functionality. In addition, the MRCP application developer needs to activate the save-waveform session parameter for these files to be generated.

Also, once these files are generated, the RECOGNITION-COMPLETE message may report the Waveform-URI that can be used to access these files, which poses an additional potential data risk. In such cases, implementing access restrictions to these exposed URI's may also be desirable.

Disabling the save-waveform functionality for potentially sensitive interactions removes this risk, since these files would not be created.

The knowledge base article on **MRCP-API** settings can be used as a reference to disable or configure the waveform saving functionality; specifically the save_waveform and waveform_url_location settings. Basically, if the waveform_url_location setting is set to a blank string, this will disable any possible save-waveform

functionality within the MRCP interfaces in a way that ASR audio will not be recorded using that feature.

Also, using TLS secured connectivity for your MRCP and RTP traffic is recommended.

MRCP Specific Inline Grammars

Inline grammars that are sent to the recognizer resource may include sensitive data, which can appear in log output (described above) as well being archived, and therefore accessible via the Analysis Portal (see section below).

Risk Mitigation:

If inline grammars are specified, and the “secure_context” option is enabled for the resource, the grammar will be suppressed from logging. The recommended option is to avoid using inline grammars in favor of grammars stored in a secured (HTTPS) location to further enhance security. Also, using TLS secured connectivity for your MRCP and RTP traffic is recommended.

MRCP Specific Inline TTS

Inline SSML or plain text may be specified in certain MRCP requests, such as the SPEAK request. These contain potentially sensitive information, which can appear in the logs.

Risk Mitigation:

To prevent this exposure, users are encouraged to specify remote (HTTPS) locations to store such sensitive SSML data in favor of inline text to offer protection as they traverse the network. If inline SSMLs are specified and the “secure_context” option for the resource is enabled, the SSML will be suppressed from logging, as well as from being archived (see below). Also, using TLS secured connectivity for your MRCP and RTP traffic is recommended.

MRCP Specific Inline Interpret-Text

The MRCP INTERPRET-TEXT request is used to request that the recognizer resource perform a text parse of the specified text. This text may contain potentially sensitive data, especially as it traverses the network.

Risk Mitigation:

Enabling the “secure_context” for the resource will prevent this text from being recorded in logging output. Also, using TLS secured connectivity for your MRCP and RTP traffic is recommended.

Voice Biometric Data

Enrollment and verification audio, as well as related metadata for voiceprints and identities, may be optionally stored within the system when using Voice Biometrics.

When using Passive Voice Biometrics, specifically, it is important to note that recordings could contain potentially sensitive information like credit card details, if spoken by the user during the enrollment or verification process.

Applications have the option of using customer identifiable information as the enrollment tag which may expose potentially sensitive information, such as a mobile phone number, or other sensitive data, causing this enrollment tag to potentially be recorded in logging or the database.

LumenVox recommends that any voiceprint creation, amendment or deletion is closely monitored for any potential exposure of sensitive data.

Risk Mitigation:

LumenVox recommends that verification audio is stored for audit and troubleshooting purposes. Enrollment audio should also be stored for troubleshooting purposes as well as to allow re-creation of voiceprints should the models be upgraded or upon major software version upgrade; provided however that the security risk associated with this storage should be considered.

To avoid exposing sensitive data, such as PII, via the logging tags associated with Voice Biometrics, LumenVox recommends that only non-PII references are used and that any correlation with that information happen outside of the LumenVox system, for example by the use of anonymous tokens.

Audit tables in the Postgres database and changes to configurations should be monitored to mitigate exposure of PII.

Data Archiving

LumenVox software has the ability to archive session and interaction-specific audio and related metadata for future analysis. This functionality can be useful from a diagnostic perspective, as well as a system accuracy monitoring perspective.

Archiving may be explicitly enabled at a global or session level, and by default this is not enabled.

Archived data may contain detailed information relating to all sessions and interactions, which may also contain potentially sensitive data. This data is stored in a combination of MongoDB (for binary data) and Postgres (corresponding metadata) Databases. This data may be accessible by authorized users via the Analysis Portal, when used. See sections below relating to MongoDB and Postgres data security.

Risk Mitigation:

Disabling the archive_session setting within the session configuration options prevents data from being archived, so this is the primary mechanism to mitigate archived data risk.

Analysis Portal

The LumenVox Analysis Portal is designed to analyze the audio and other information from ASR, transcription and CPA interactions in order to diagnose any performance issues. This, by definition, means that it may need access to potentially sensitive data because those interactions also need to be analyzed and diagnosed at times.

The Analysis Portal utilizes interaction and session data that have been archived, including grammars, phrase lists, audio, results and other metadata which may contain potentially sensitive data. Anyone using, viewing or downloading data from the Analysis Portal may intentionally or unintentionally gain access to this potentially sensitive data.

Risk Mitigation:

We recommend that sensitive session or interaction data should not be archived, and that archiving only be enabled for troubleshooting purposes for such sensitive interactions. Access to the Analysis Portal should also be strictly restricted using Kubernetes RBAC and other policies. Access to archived data can also be accessed via the reporting API, which should have similar security restrictions in place to mitigate risk.

Admin Portal

The LumenVox Admin Portal is an administrative web interface running within the Kubernetes stack alongside other LumenVox services. This is designed to provide high-level administrative access to the system, including tasks such as adding or modifying deployments (tenants) on the system, as well as configuring connectivity to external resources, such as RabbitMQ, Redis, MongoDB and Postgres instances. These connectivity settings define credentials allowing connectivity to those resources so they are by definition sensitive in nature.

Risk Mitigation:

Access to the Admin Portals should be strictly controlled using Kubernetes RBAC and other policies your IT team may require. User access to this portal should be limited to IT and Security personnel only. Specifically, only those who require access to configuring deployments and/or how connectivity is defined for external resources should be allowed access.

TLS secured connectivity (HTTPS) is recommended when this portal is accessed.

Access to this portal can be controlled by the corresponding Kubernetes Ingress configuration, allowing system administrators to determine who may have access.

Deployment Portal

The LumenVox Deployment Portal allows browser-based control over settings and also provides access status information for a specific deployment (tenant). In addition, diagnostic tests are available to validate the functionality of the deployment, if needed. The Deployment Portal can also provide access to the Analysis Portal for a specific deployment.

Risk Mitigation:

Access to the Deployment Portals should be strictly controlled using Kubernetes RBAC and other policies your IT team may require. Access to this portal is not essential for system operation, so providing access will typically be determined by the system administrator.

TLS secured connectivity (HTTPS) is recommended when this portal is accessed.

Access to this portal can be controlled by the corresponding Kubernetes Ingress configuration, allowing system administrators to determine who may have access.

External Resources - Redis

Redis is used as a requirement for the LumenVox services to communicate efficiently when processing interactions. Redis holds transient information about interactions as they are being processed. This also means that this data, which includes audio as well as request and result metadata, may include potentially sensitive data.

Risk Mitigation:

Interconnectivity between LumenVox services and external resources such as Redis are designed to use TLS encryption to protect data in transit. Also, access to the external resource itself should be secured with unique and secure credentials. There is no user-serviceable information held within Redis, and the intent is that only LumenVox software be allowed access to it within the confines of the configured and secured Kubernetes environment. We recommend adopting this approach to securing this resource to mitigate risk of data exposure.

External Resources - RabbitMQ

RabbitMQ is used as a requirement for the LumenVox services to communicate efficiently when processing interactions. RabbitMQ holds transient information about interactions as they are being processed. Utilization of RabbitMQ also means that this data, which includes request and result notifications, may include potentially sensitive data.

Risk Mitigation:

Interconnectivity between LumenVox services and external resources such as RabbitMQ are designed to use TLS encryption to protect data in transit. Also, access to the external resource itself should be secured with unique and secure credentials. There is no user-serviceable information held within RabbitMQ, and the intent is that only LumenVox software be allowed access to it within the confines of the configured and secured Kubernetes environment. We recommend adopting this approach to securing this resource to mitigate risk of data exposure.

External Resources - Postgres Database

Postgres Database is used as a requirement by the LumenVox services to operate efficiently when processing interactions. Postgres holds persistent metadata about interactions that are processed, including license usage information metrics (counters). In addition, when archiving is enabled, Postgres includes session and interaction metadata for later analysis. This also means that this data, which includes request, result and configuration metadata, may include potentially sensitive data.

Risk Mitigation:

Interconnectivity between LumenVox services and external resources such as Postgres are designed to use TLS encryption to protect data in transit. Also, access to the external resource itself should be secured with unique and secure credentials. There is no user-serviceable information held within Postgres, and the intent is that only LumenVox software be allowed access to it within the confines of the configured and secured Kubernetes environment. We recommend adopting this approach to securing this resource to mitigate risk of data exposure.

External Resources - MongoDB Database

MongoDB Database is used as a requirement by the LumenVox services to operate efficiently, when processing interactions. MongoDB holds persistent binary data relating to sessions and interactions only when archiving is enabled. This binary data, in conjunction with related metadata (within Postgres),can be used for later analysis. This also means that this data, which includes audio and other binary data for interactions, may include potentially sensitive data.

Risk Mitigation:

Interconnectivity between LumenVox services and external resources such as this are designed to use TLS encryption to protect data in transit. Also, access to the external resource itself should be secured with unique and secure credentials. There is no user-serviceable information held within this resource, and the intent is that only LumenVox software be allowed access to it, within the confines of the configured and secured Kubernetes environment. We recommend adopting this approach to securing this resource to mitigate risk of data exposure.

In addition, all binary data stored within MongoDB for a LumenVox application is encrypted within the MongoDB database, using a deployment (tenant)-specific encryption key, which can be periodically rotated when needed.

LumenVox® by capacity

LumenVox is a speech & voice biometric technology company offering a range of products including the ASR, transcription, Text-to-Speech, Call Progress Analysis and Voice Biometrics. Based on industry standards, LumenVox's core Speech & Voice Biometric Software is recognized as one of the most accurate, and reliable solutions in the industry.

For more information, visit **www.lumenvox.com**