Siebel SR Vector Search using Oracle Database 23ai

February, 2025, Version [1.0]

Copyright © 2025, Oracle and/or its affiliates

Public

Purpose statement

This document has been created by the Siebel Center of Excellence team for informational purposes only. It aims to provide valuable insights and guidance on the subject matter while serving as a reference for stakeholders. This document does not constitute a binding agreement or official policy.

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

Summary

Oracle AI Vector Search, a feature of Oracle Database 23ai, enables data queries based on semantics rather than mere keywords. This technology can seamlessly integrate into Siebel CRM to activate AI-driven use cases like semantic searches and context-based data retrieval. In this white paper, we explore how vectorization can be used to process and retrieve related Service Requests (SRs) in Siebel CRM, leading to improved data accuracy, retrieval efficiency, and enhanced customer experience.

Overview

Traditional database searches are typically built on keyword-based search algorithms. These methods operate by matching the literal words in a query to those stored in the database. While effective for basic searches, this approach has significant limitations:

- **Limited Contextual Understanding:** Keyword searches are based on exact word matches or simple pattern recognition. They don't understand the relationships between words or their meanings in different contexts. For example, if a search query includes the word "car," a keyword search may miss results with synonyms like "automobile" or relevant phrases that imply the concept of a car without using the exact word. This can lead to incomplete or imprecise results.
- **Inability to Interpret Synonyms and Variations:** Keyword searches may struggle with variations in phrasing or synonyms. If a user searches for "purchase order," results containing "buying order" or "sales order" might be overlooked, even if they are contextually relevant.
- **Literal Matching Over Semantics:** Traditional searches do not consider the semantic meaning of the data. For instance, a keyword-based system would treat the phrases "high temperatures" and "hot weather" as unrelated, even though they convey similar meanings. This lack of semantic understanding can be a barrier in applications where nuanced interpretations of language are crucial.

Oracle Database 23ai introduces the VECTOR data type, enabling the storage and processing of vector embeddings within the database. This integration can support a range of AI use cases, including natural language processing (NLP), recommendation systems, and semantic searches. By embedding this technology into Siebel CRM, organizations can unlock powerful capabilities for finding related service requests, enabling agents and managers to respond more effectively to customer inquiries and issues.

Vector embeddings overcome these challenges by transforming data—whether words, sentences, or documents—into numerical representations that encode semantic meaning. Here's how they make searches more intuitive and contextually accurate:

- **Capturing Semantic Relationships:** Vector embeddings map data points (e.g., words or sentences) into a high-dimensional space where their proximity reflects their semantic similarity. For example, the words "king" and "queen" would be close in this space, as would "car" and "vehicle." This spatial representation allows searches to recognize related content even when exact keywords are not present.
- **Contextual Awareness:** Unlike keyword-based searches, vector embeddings can understand the context in which words appear. For instance, the word "bank" in "river bank" would be represented differently than "bank" in "financial bank" because embeddings take the surrounding context into account. This helps retrieve results that are relevant based on the query's intended meaning, not just its literal composition.
- **Handling Synonyms and Variations:** By representing data in a way that emphasizes semantic meaning, embeddings make it possible to identify matches across different expressions of the same idea. This means a search for "automobile" would still find results for "car," "vehicle," or even "SUV," improving the search's comprehensiveness.
- **Supporting Complex Queries:** Vector embeddings allow for more complex queries that go beyond simple keyword matching. For example, a search for "best ways to handle customer complaints" can return results

Copyright © 2025, Oracle and/or its affiliates / Public

related to "effective complaint resolution strategies" because the embedding system understands the underlying meaning rather than just word matches.

Solution Approach

For finding related Service Requests, we are leveraging Oracle Database 23c AI's vector embedding capabilities to implement semantic search functionality. By integrating the VECTOR data type and vector embeddings into Siebel CRM, it becomes possible to retrieve Service Requests (SRs) that are semantically related, based on their contextual meaning rather than exact keyword matches.

Initial Setup and Data Preparation

The integration of vectorization into Siebel CRM begins with configuring the VECTOR data type in Oracle Database 23ai. The process includes:

- **Defining VECTOR Data Type:** The VECTOR type allows for the storage of vector embeddings, aligning them with traditional business data for richer querying capabilities.
- **Selecting or Importing Embedding Models:** Various models such as all-MiniLM-L12-v2, OpenAI's text-embedding models, or custom-built ones can be used to create vector embeddings. These models can be imported into Oracle Database using the ONNX standard, enabling direct vector generation within SQL.

Building the Vectorization Process

- **PL/SQL Procedure for Vector Embedding Creation:** A custom PL/SQL procedure reads the SR data from the Siebel database, processes it, and generates vector embeddings. This is stored in a newly created SR_VECTOR table. The embedding process can involve pre-processing with NLP techniques to ensure data consistency and meaningful embedding generation.
- **Triggers for Automatic Updates:** Any update in the SR data triggers the PL/SQL procedure to regenerate and update the vector in the table. This keeps the SR_VECTOR embeddings up-to-date with the latest SR data, maintaining search accuracy and relevance.

Implementing Semantic Search

Semantic search in Siebel CRM leverages the generated vectors to retrieve SRs that are contextually related to the current record being viewed:

- **UI Component Integration:** A customized UI component displays related SRs on the service request detail screen. When an agent accesses an SR, this component invokes a REST API to query the related SR search microservice
- **REST API and Microservice:** The microservice, hosted in a scalable container environment, receives requests containing the parent SR details. It performs a similarity search using the VECTOR data stored in the Oracle database, returning SRs that are semantically similar to the query.

The results are then displayed in the Siebel UI, providing agents with immediate, contextually relevant information.

Copyright © 2025, Oracle and/or its affiliates / Public

Architecture and Sequence Flow

The overall architecture involves the following key components:

- **Siebel UI:** Configured to include a related SR section within the Service Request Applet.
- **Related SR Search Microservice:** Acts as an intermediary between the UI and the database, handling REST API calls and executing vector-based similarity searches.
- **In-Database Embedding Model:** Supports vector embedding creation directly within Oracle Database.
- **SR_VECTOR Table:** Stores vector embeddings of SRs, enabling efficient search queries.
- **PL/SQL Procedure:** Handles embedding generation and updates as SR data changes.

Sequence Flow

1. The initial population of the SR_VECTOR table is completed through a bulk PL/SQL procedure during system setup or data migration.
2. Updates to the SR table trigger the embedding procedure, updating the vector entry for any modified SRs.
3. When an agent accesses the SR detail page, the UI triggers a REST API call to the microservice.
4. The microservice executes a semantic search using SQL and returns the related SRs to the UI.

Benefits

- **Enhanced Search Accuracy:** Semantic search focuses on the meaning and context of words, resulting in more relevant search outcomes compared to traditional keyword-based methods. This capability helps agents find SRs that may not share common keywords but are related by context, leading to more informed responses.
- **Improved Efficiency and Productivity:** Automating the search for related SRs saves significant time for service agents who would otherwise manually search for similar cases. This not only speeds up the response process but also improves overall productivity and service quality.
- **Scalability and Flexibility:** Using vector embeddings and Oracle Database 23ai's VECTOR data type allows organizations to scale their AI capabilities. Different embedding models can be evaluated and selected based on the specific needs of the organization, ensuring an optimal balance between performance and accuracy.
- **Reduced Human Error:** The automated process of embedding generation and similarity search minimizes the risk of inconsistencies and human error. Agents can trust that the system will provide contextually relevant information every time.
- **Enhanced Customer Experience:** With faster access to related information, agents can provide quicker and more accurate responses, improving the overall customer experience. Customers benefit from reduced waiting times and more efficient issue resolution.

Copyright © 2025, Oracle and/or its affiliates / Public

Conclusion

The integration of Oracle Database 23ai's VECTOR data type and vector embeddings into Siebel CRM paves the way for advanced AI-driven capabilities. By enabling semantic searches and automating the process of finding related SRs, this solution transforms how customer service operations are managed. The result is a more streamlined workflow, reduced agent workload, and a superior customer experience. Organizations looking to innovate their CRM capabilities can use this approach as a blueprint for deploying similar AI enhancements.

For more insights or design partnerships, reach out to us at siebel_coe_grp@oracle.com

Copyright © 2025, Oracle and/or its affiliates / Public

ORACLE

Connect with us

Call +**1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

blogs.oracle.com facebook.com/oracle twitter.com/oracle

Copyright © 2025, Oracle and/or its affiliates. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.