SingleStore ™

Proof of Value

Build

This document outlines the purpose, scope and implementation of SingleStore supported by an initial Proof of Value (POV), to be delivered by [customer name] with support from SingleStore.

Prepared by

First name Last name (Solutions Engineer)
First name Last name (Account Executive)

Table of contents

Stakeholder Confirmation

By listing my name below, I agree the POV specifications in this document are the requirements as defined by [prospect/customer company name] and SingleStore. This is not a contract, but rather a document to ensure a shared understanding of the scope and objectives of the POV.

Stakeholder Name	Role
Nick Ambrose (Customer Company Name)	Engineering Manager (email)
George Zhu (Partner Company Name)	Director Of Engineering (email)
	…

POV start date	POV end date
##/##/####	##/##/####

High-level business outcomes

Example: [Customer] is building an internal generative AI platform to process large volumes of PDF documents, loan agreements and more, giving users the ability to ask questions about document contents, summarize, extract and explain specific entities.

SingleStore will be used as the on-premises database, powering the loan document analysis application based on LangChain framework.

High-level technical outcomes and success criteria

Short hypothetical list of technical outcomes (tabular or list format), along with post-POV proven thesis.

Criteria	Desired performance	Current performance	SingleStore performance	% to desired performance (if known)
				%
				%
				%
				%

High-level execution plan (optional)

Prerequisites and dependencies; list of things we (and the customer) need to do to be successful:

Dependency	Owner
SingleStore	
Size workspace	SE
Create trial cluster and provide customer with relevant security credentials	SE
Slack channel	AE
…	…
Customer	
Private Link	Customer point-of-contact
NDA	Customer point-of-contact
…	…
SingleStore and customer jointly	
Monitor load	SE and customer POV
…	…

The hypothesis

- The detailed list of what we'll be doing along with times, if appropriate; how we'll be doing it with SingleStore.

Proven theses

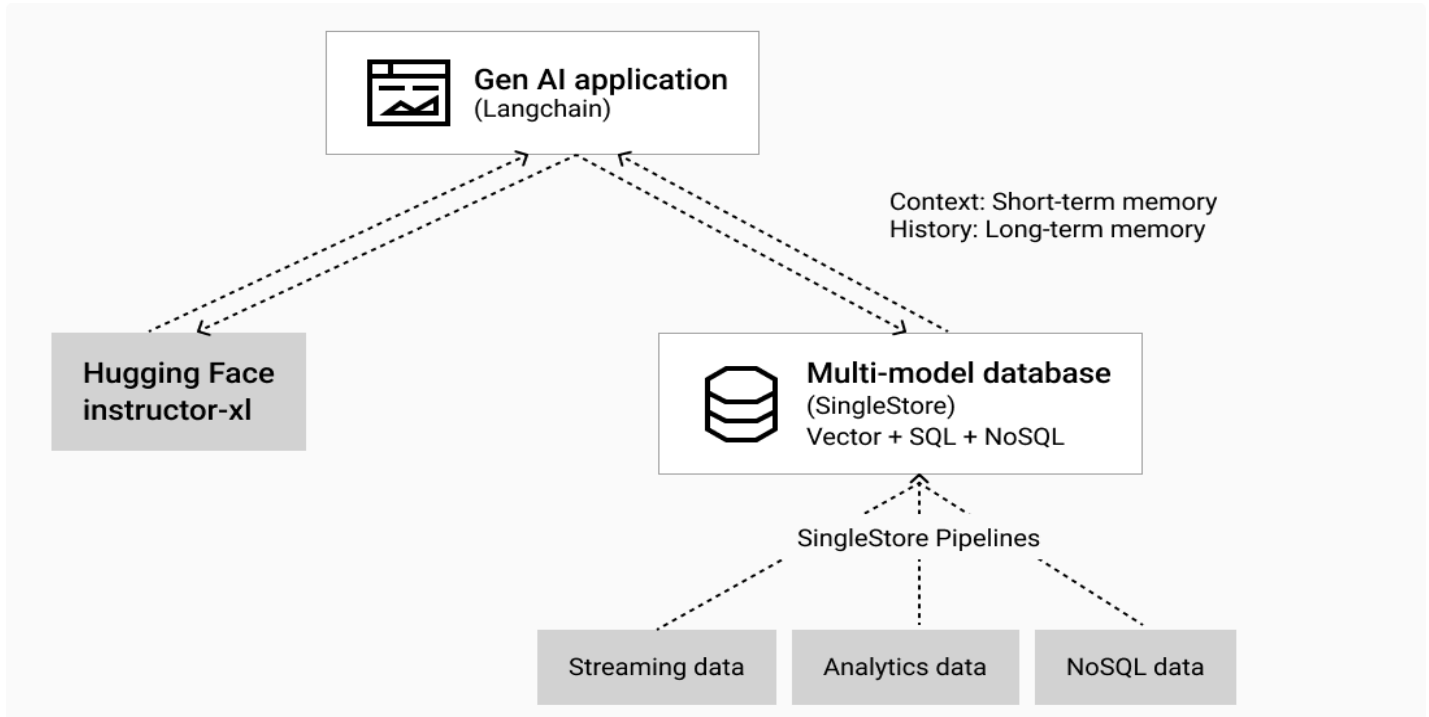- List of reasons why the hypothesis was either proved or disproved.

POV components

- List of resources required for the POV (Dropbox links, S3 buckets, workspace sizing, SingleStore org ID, cloud region for deployment, table DDLs, queries to be run, Github links, etc.)

- Producers → ingestion / tuning / scaling / integration → consumers

- Dependencies (create a list):

POV architecture

This architecture diagram represents a high-level look at what will be built for the POV:

Workspace sizing	SingleStore Cloud Size S-8+
Chunking and Embedding Strategy	What is the client's chunking and embedding strategy for this specific use-case? As SingleStore we should have an opinion or guidelines to provide them with.

POV collateral

A curated collection of links pertinent to the ongoing POV — including documents and resources encountered during the implementation phase — tailored to the specific context of the POV:

General setup

SingleStore Cloud on AWS Marketplace

Upgrade to developer support

Add users and assign RBAC

Create workspaces

Basic configurations

Workspaces

Create your first table in SingleStore

SingleStore private connectivity

SingleStore CLI

Load data/vectorize

Load data from popular sources with SingleStore Pipelines

Connect to workspace with your favorite tool

Using vector functions

Vector indexing

Vector type support in SingleStore