SNIA | Cloud Storage Technologies

How Agentic AI Transforms the Role of Storage

Himabindu Tummala
Dell Technologies

John Cardente
Dell Technologies

Fatemeh Azmandian
Dell Technologies

Michael Hoard
SNIA CST Chair

Live Webinar
February 12, 2026
10:00 am PT / 1:00 pm ET

Today's Presenters

Michael Hoard
Moderator
SNIA
CST Chair

Fatemeh Azmandian
Distinguished Engineer
Dell Technologies

John Cardente
Senior Distinguished Engineer
Dell Technologies

Himabindu Tummala
Distinguished Engineer
Dell Technologies

 ©SNIA. All Rights Reserved.

SNIA | CST A SNIA Community

SNIA Cloud Object Storage Plugfests sponsored by CST

SNIA | Cloud Storage Technologies

https://snia.org/groups/cst

Committed to the adoption, growth and standardization of **intelligent data storage usage** in cloud infrastructures

- New simplified SNIA Membership Model
- Pay one fee for SNIA Membership – Participate in ANY group!
- Join SNIA groups at https://www.snia.org/member_com

 ©SNIA. All Rights Reserved.

The SNIA Community

200
industry leading organizations

2,000
active contributing members

50,000
IT end users & storage pros worldwide

| ©SNIA. All Rights Reserved.

SNIA | CST A SNIA Community

SNIA Legal Notice

- The material contained in this presentation is copyrighted by SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

©SNIA. All Rights Reserved.

Agenda

- Agentic AI Overview
- Multi-Agent Systems & Interoperability
- Storage in the Context of Agentic AI
- Summary

 ©SNIA. All Rights Reserved.

SNIA | CST
A SNIA Community

Agentic AI Overview

Agentic is the next wave in the evolution of AI Technologies

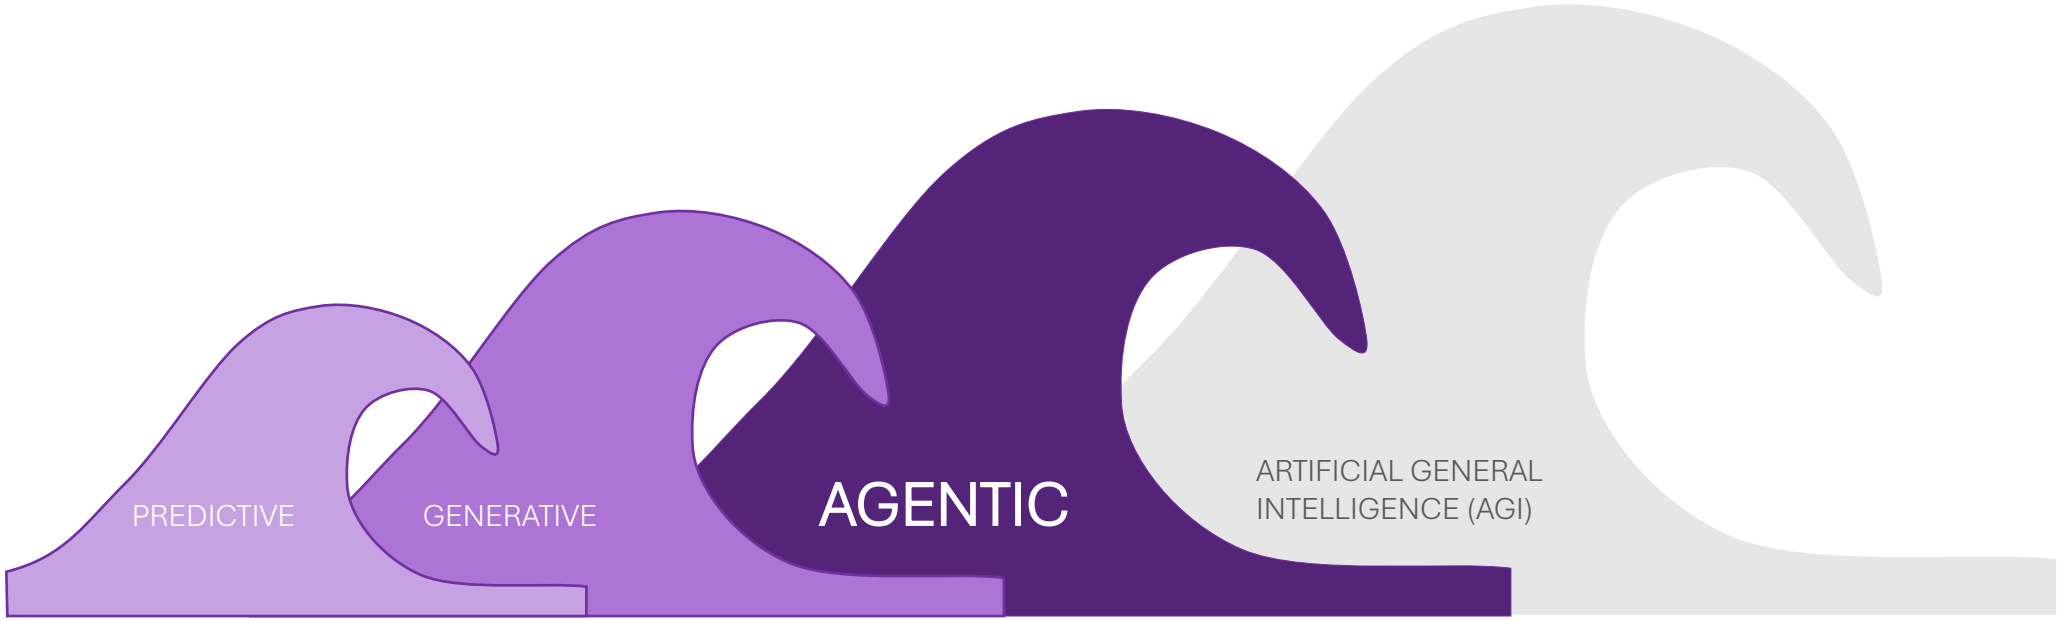

 ©SNIA. All Rights Reserved.

SNIA | CST
A SNIA Community

The ability to autonomously reason and act differentiate agents from GenAI chatbots, assistants and workflows

Directly answers questions from the model's learned knowledge

Answers questions on information outside of training dataset using Retrieval Augmented Generation (RAG) to build appropriate context

Accomplishes a task by executing a pre-defined workflow composed of a chain of LLM prompts and pre-determined actions

Accomplishes an objective by dynamically directing its own reasoning and actions based on the observed environment

 ©SNIA. All Rights Reserved.

Agents are generally built from four main components – a model, prompt, tools, and memory – integrated by orchestration logic

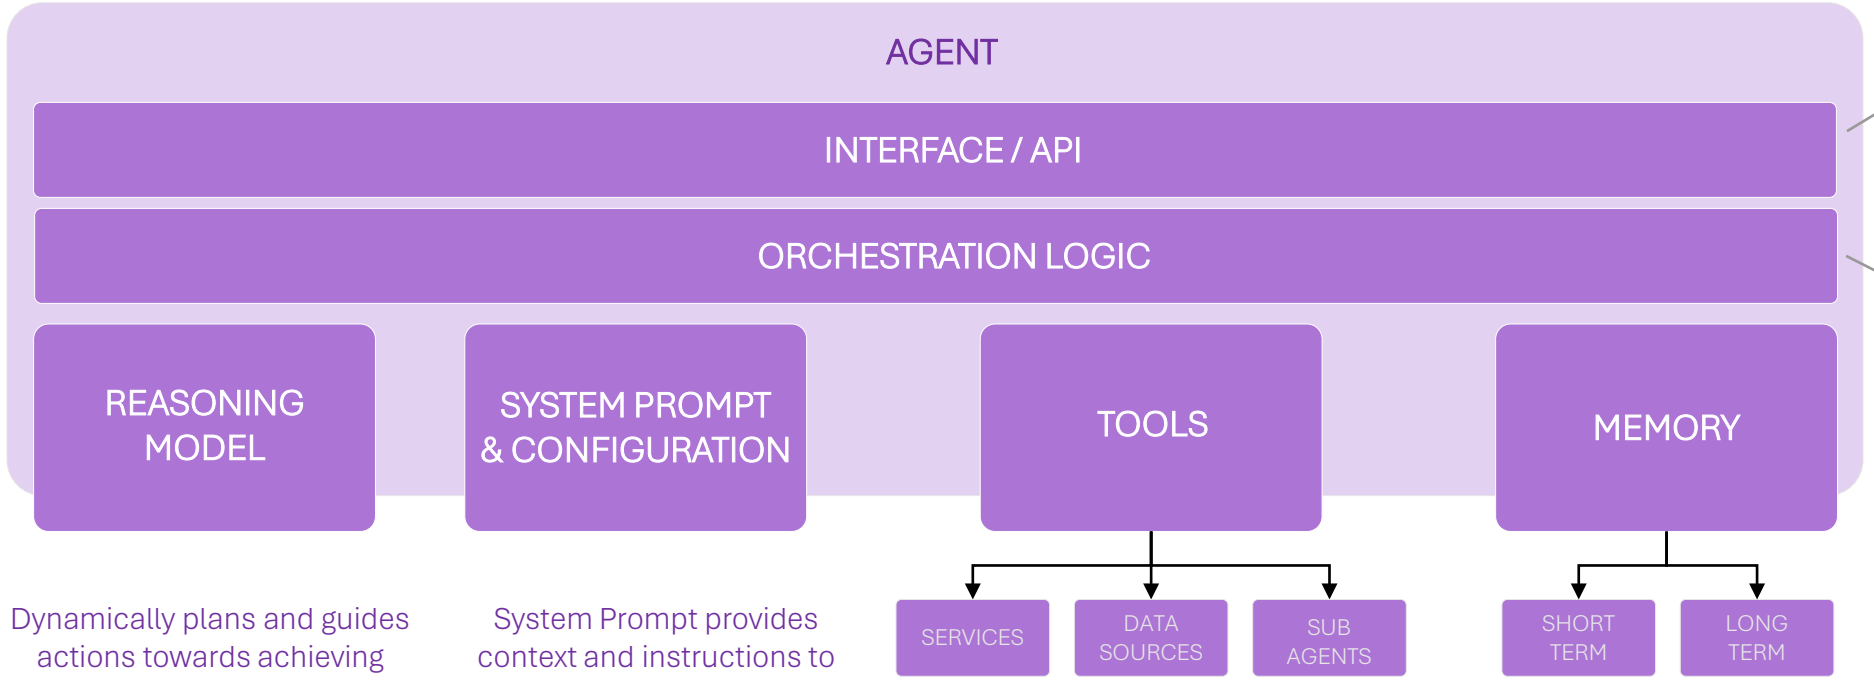

Provides a way for the agent to receive requests to achieve an objective

Implements the logic needed to integrate the components into a functional agent

Dynamically plans and guides actions towards achieving objective using an iterative reason, act, & observe process (ReAct)

System Prompt provides context and instructions to guide the LLM's reasoning.

Configuration adjusts the model's behavior (e.g. temperature, top-k, etc.)

Collects data or performs actions needed to progress towards an objective

Stores information during and between agent executions to inform planning, execution, and responding

 ©SNIA. All Rights Reserved.

An agent is a single logical construct but likely not a monolithic implementation; may utilize shared services like LLMs, databases, etc.

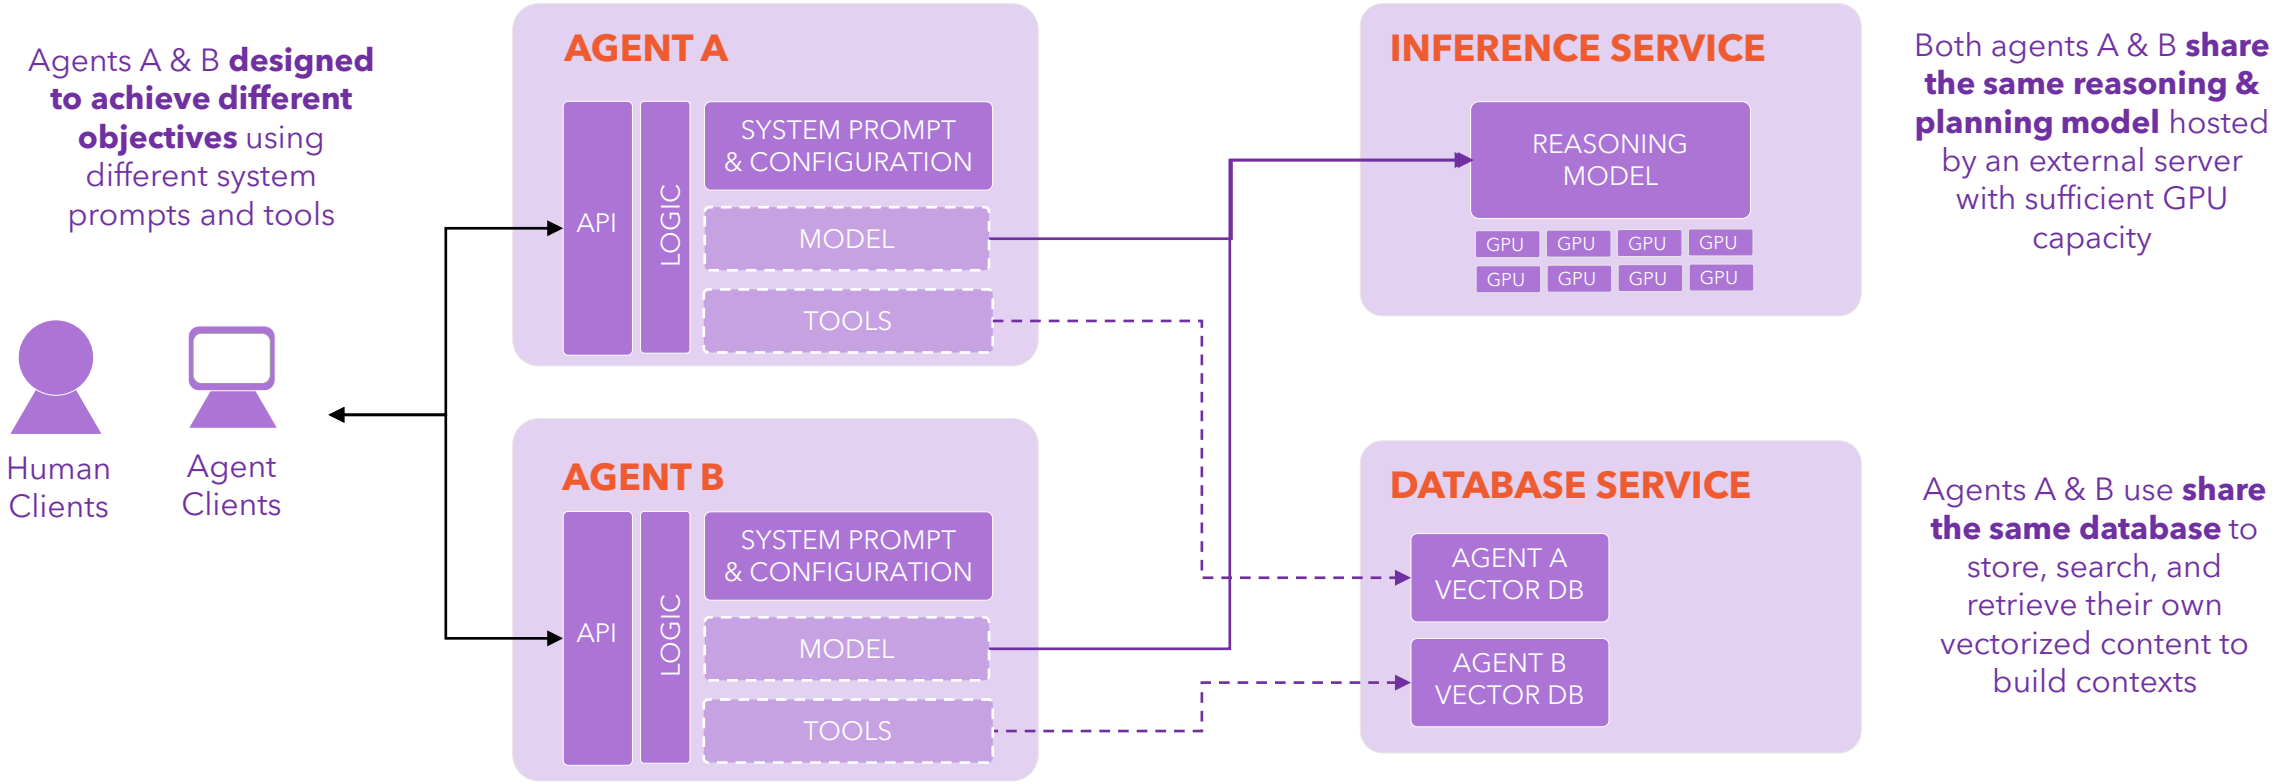

 ©SNIA. All Rights Reserved.

Reasoning models use an observe-reason-act loop to gather context and take actions needed to achieve an objective

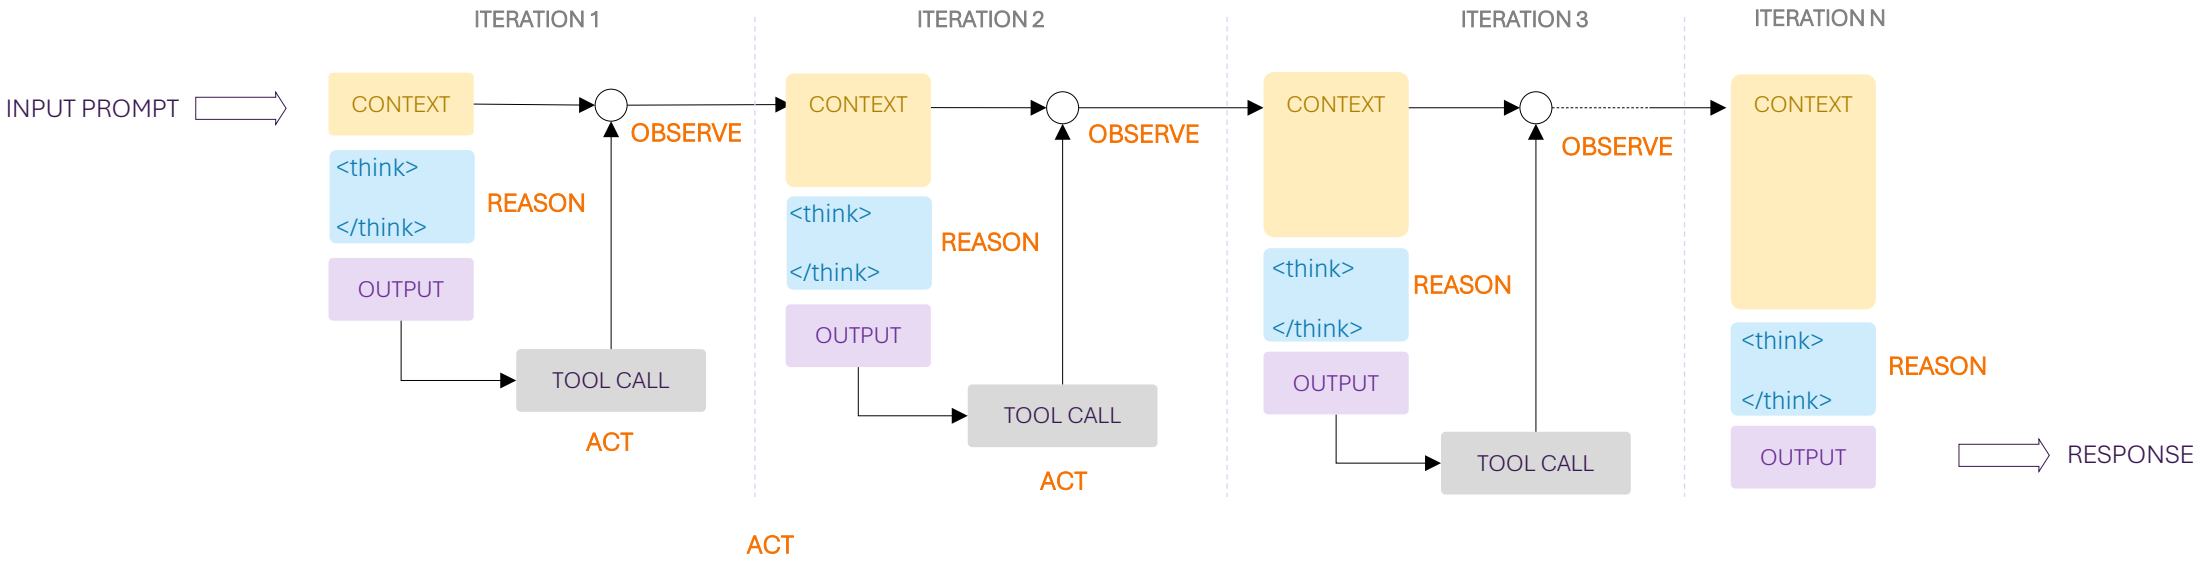

Reasoning Example

<think>
The user asked, "What's the weather like in Boston?"
To answer this accurately, I need real-time weather data.
I have access to a tool called get_weather that can retrieve this information.
I'll call this tool with the city set to "Boston" and use Celsius as the temperature unit.
</think>

Reasoning can increase the number of context tokens by 10x to 100x versus non-reasoning LLMs. This significantly increases the computational infrastructure requirements for agents at enterprise scale

 ©SNIA. All Rights Reserved.

SNIA | CST
A SNIA Community

Multi-Agent Systems & Interoperability

Automating complex tasks will require diverse and distributed teams of agents – especially for B2B use-cases

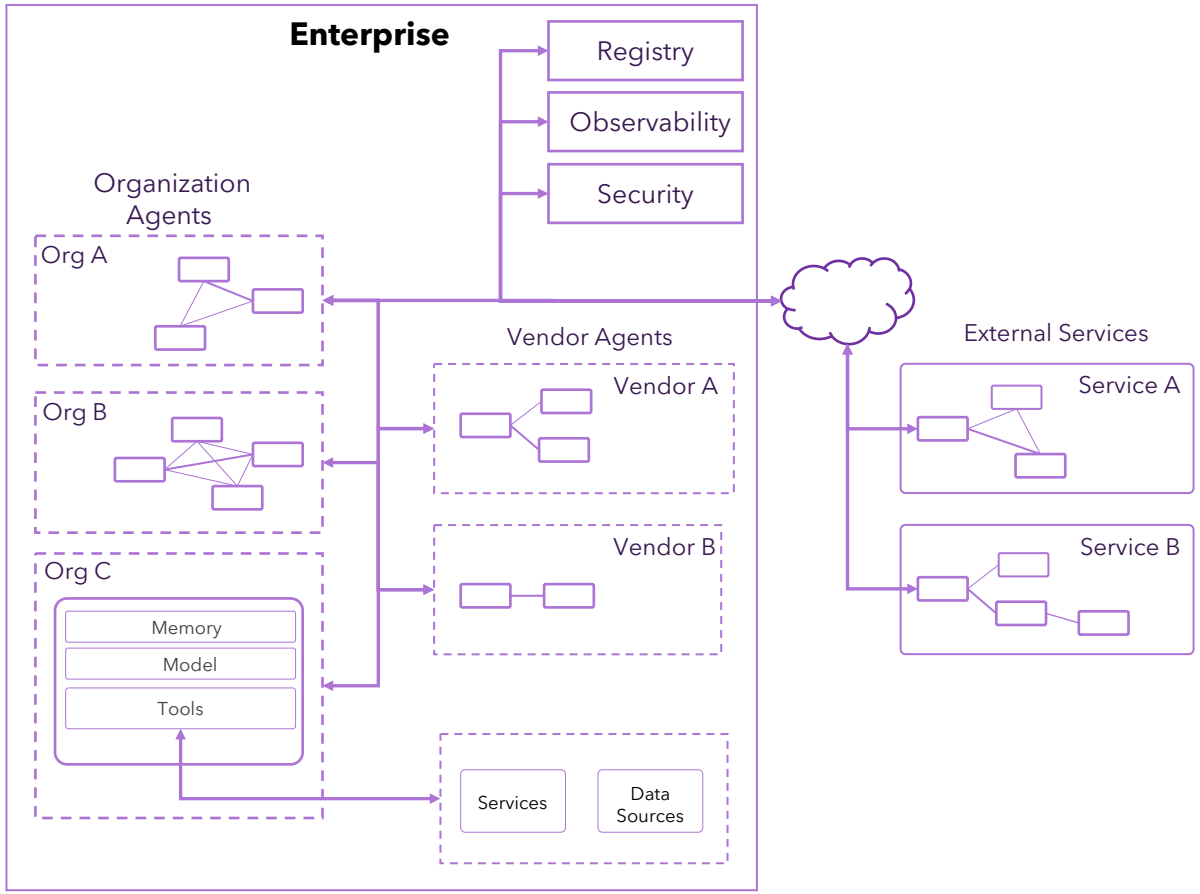

Pre-defined & Autonomous Collaboration

Agents execute pre-defined workflows or autonomously decide how to collaboratively accomplish complex business objectives

Multi-vendor & Distributed

Agents from multiple vendors deployed across different locations will work together to accomplish complex objectives

Mixed Identities & Security Domains

Agents use their identities and "on-behalf of" credentials to enable enforcement of security policies across deployment environments, organizations, & enterprises

 ©SNIA. All Rights Reserved.

Agentic Interoperability standards needed to address four key concerns

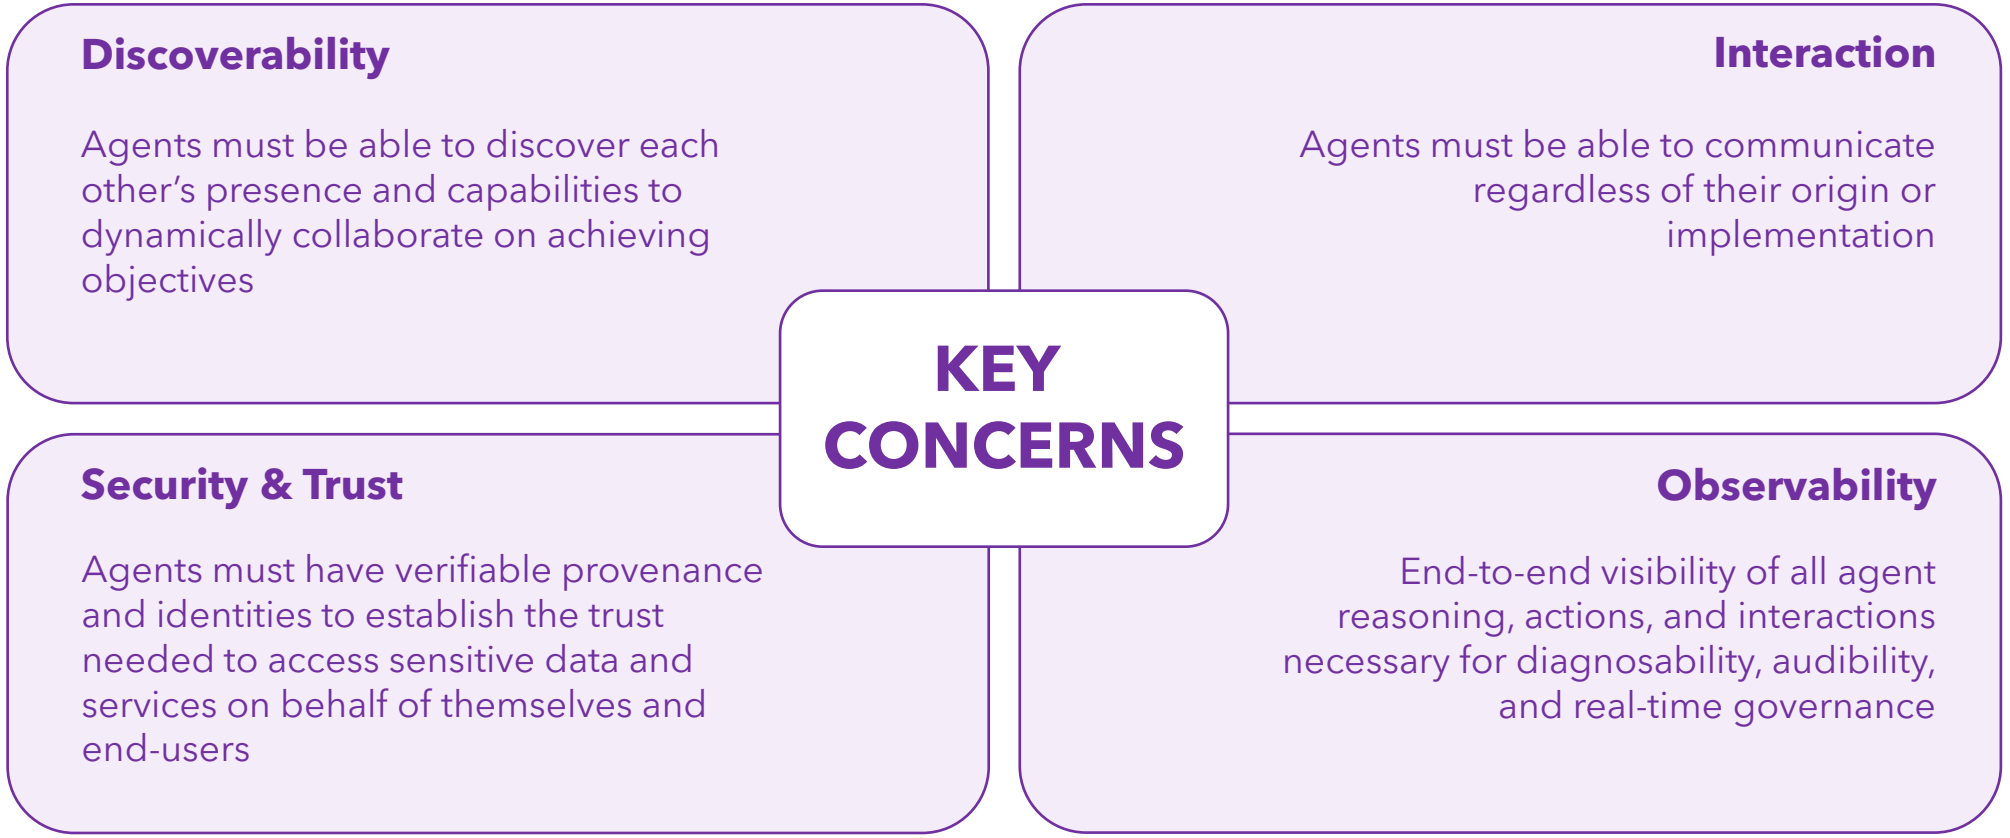

 ©SNIA. All Rights Reserved.

MCP & A2A protocols standardizing agent-to-tool & agent-to-agent interoperability

MODEL CONTEXT PROTOCOL (MCP)

ANTHROP\C

Now a Linux Foundation Agentic AI Foundation (AAIF) project

An open standard designed to securely connect AI systems with data, tools, and systems such as content repositories, databases, and service APIs

Core Design Centers:
- Retrieving data to build model context
- Synchronous function calling
- Desktop application AI assistants

AGENT TO AGENT PROTOCOL (A2A)

Google

+50 Partner Companies

Now a Linux Foundation project

An open protocol designed to facilitate communication and interoperability between AI agents, enabling them to collaborate across different frameworks and vendors without requiring custom translation layers

Core Design Centers:
- Inter-Agent communication
- Asynchronous requests for long running tasks
- Distributed multi-agent systems

 ©SNIA. All Rights Reserved.

MCP provides a standard way to connect agents with data to build context and services to perform actions

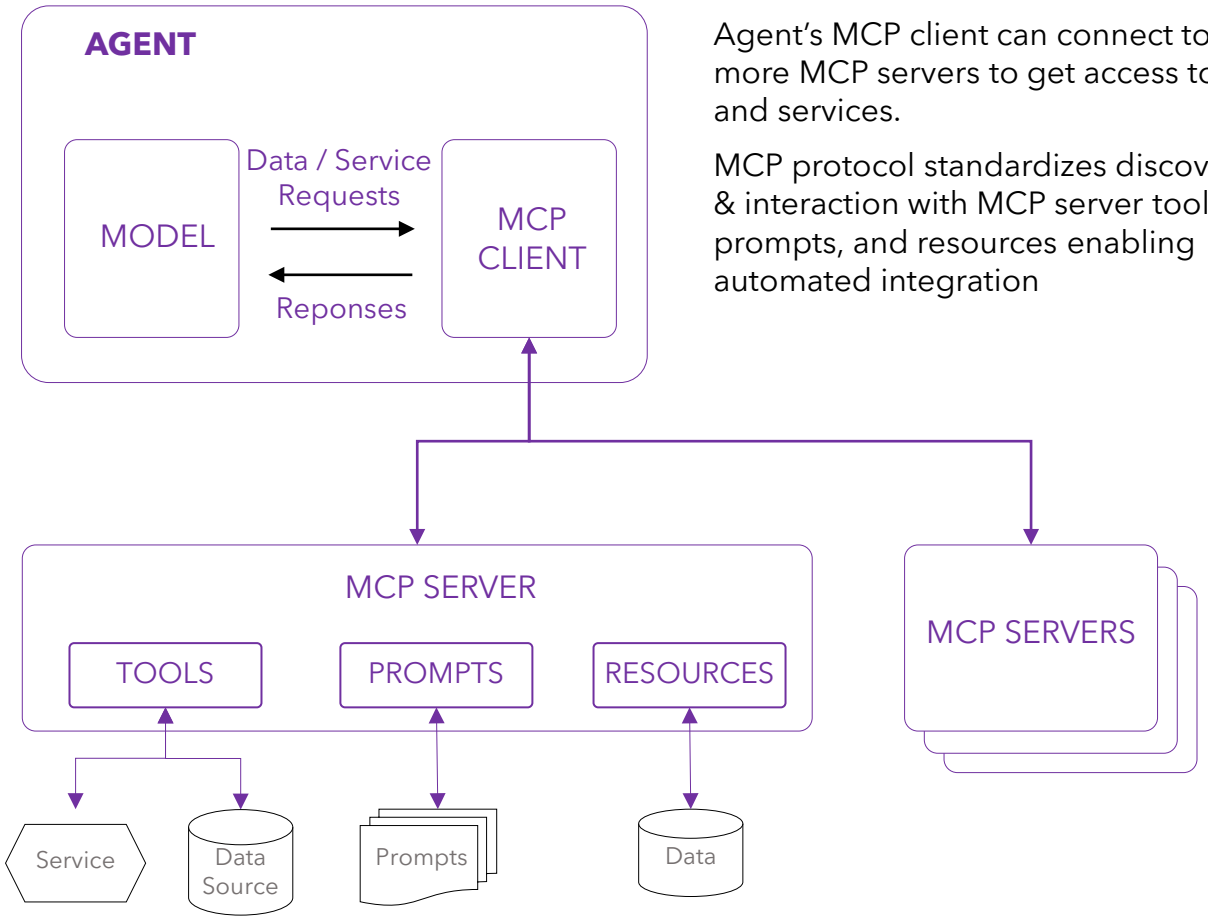

Agent's MCP client can connect to one or more MCP servers to get access to data and services.

MCP protocol standardizes discovery of & interaction with MCP server tools, prompts, and resources enabling automated integration

Seamless Integration

MCP APIs and protocols enable automatic use of services and data sources by simply connecting to new MCP servers

Re-usable Servers

MCP server implementations can be easily shared & reused to minimize development efforts. A rich ecosystem of MCP servers for common data sources and services already emerging

Shared Data & Services

Multiple AI models and applications can share a remote MCP Server using the HTTPS transport option

 ©SNIA. All Rights Reserved.

A2A enables agents to communicate and collaboratively perform complex tasks

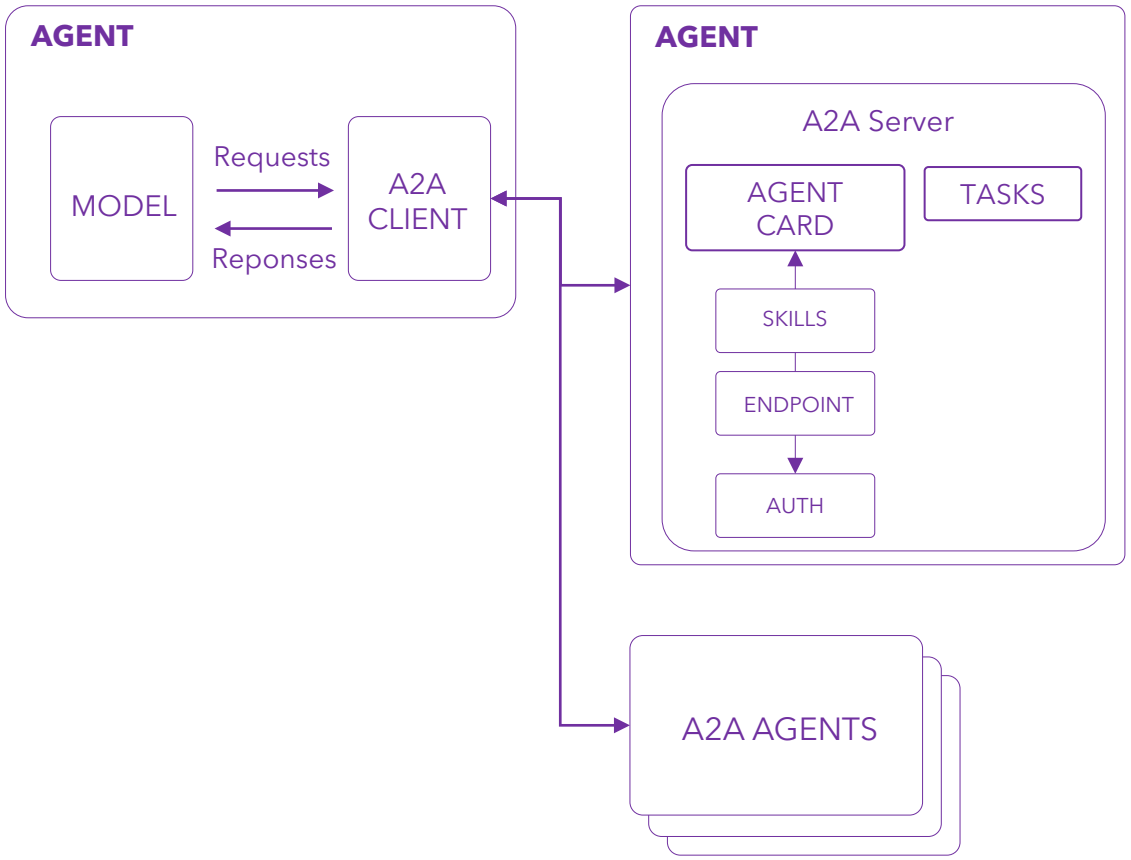

Seamless Integration

A2A APIs and protocols enable agents to discover each other's abilities and request tasks enabling a wide range of agentic architectural patterns

Complex interactions

Supports asynchronous requests, multi-turn interactions, parallel task requests, multi-modal data, streaming responses, and human-in-the-loop validations

Secure & Observable

Uses standards like TLS, OAuth2, and OTEL to securely communicate, authenticate & authorize interactions, and log activities

 ©SNIA. All Rights Reserved.

MCP & A2A solve different aspects of agentic interoperability & together enable a wide range of architectural patterns

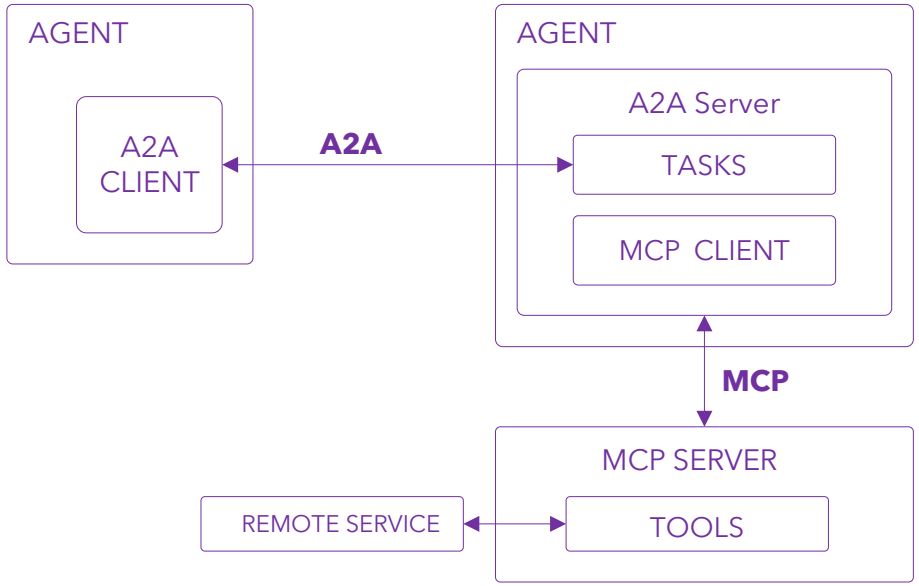

- A2A protocol handles inter-agent communication
- MCP simplifies retrieving context information from data sources and services
- Protocols compliment each other by solving different aspects of interoperability

Sequential Workflow

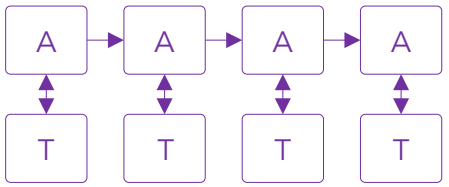

Orchestrator

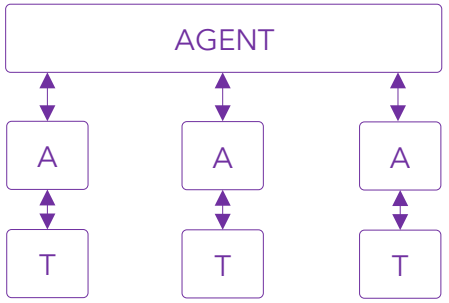

Agentic Mesh

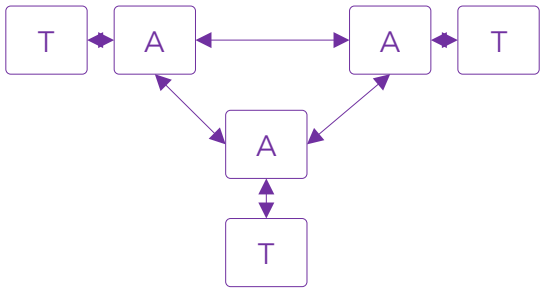

Hierarchical Decomposition

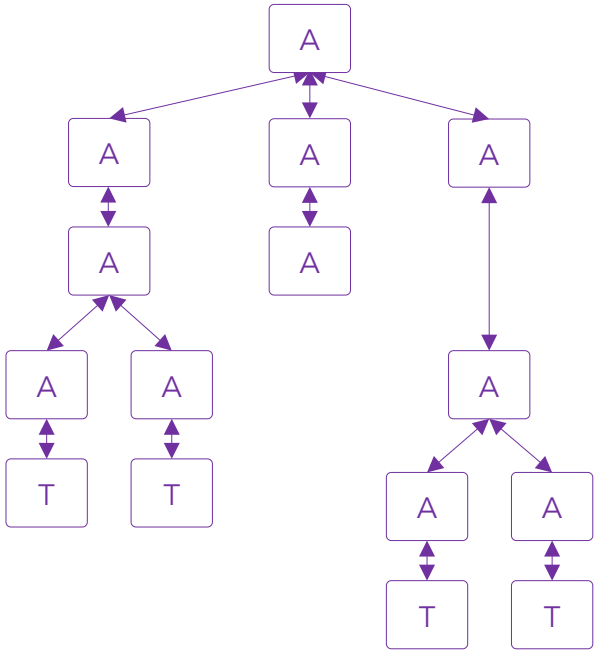

A A2A Agent

T MCP Tool Server

 ©SNIA. All Rights Reserved.

SNIA | CST A SNIA Community

The AGNTCY Linux Foundation project builds upon A2A & MCP to address discovery, identity, trust, and observability

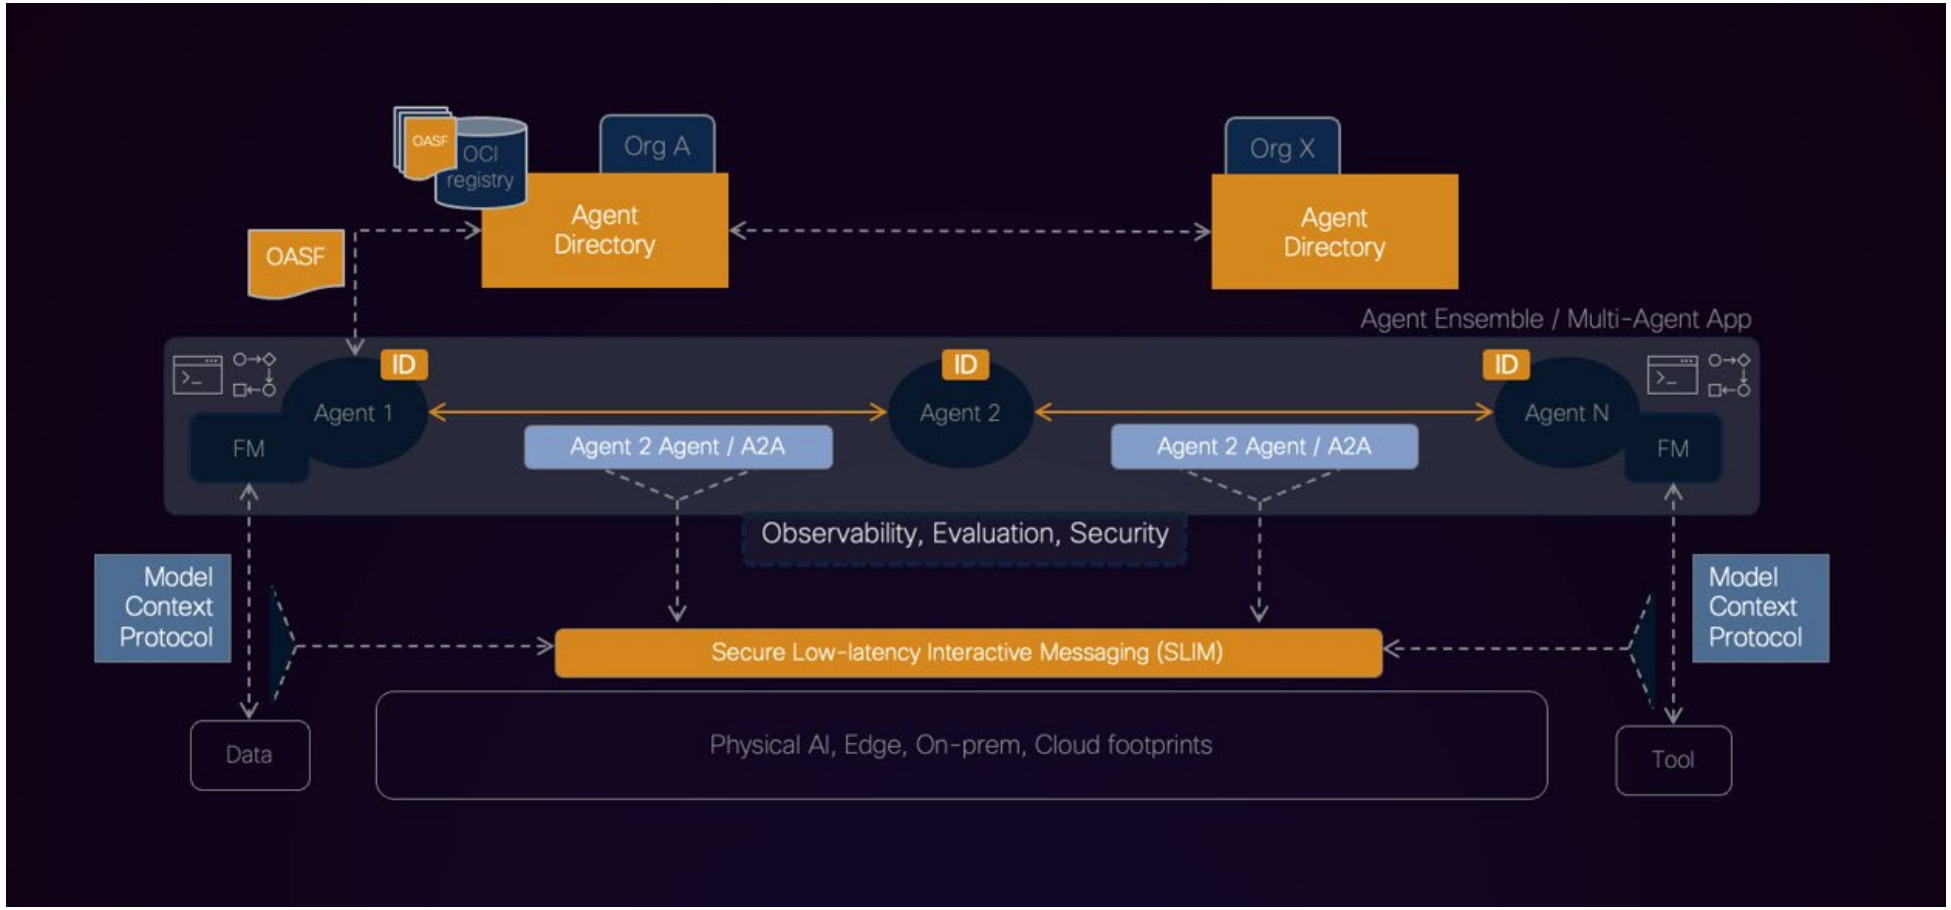

https://agntcy.org/

 ©SNIA. All Rights Reserved.

Storage in the Context of Agentic AI

Agentic AI impacts storage in two ways – new types of workloads & new ways to building context

Context Building

Agentic AI is changing the way enterprise data is discovered, retrieved, and processed to build the right context for a specific objective

EXAMPLE

AI enriched semantic search indexes and real-time data processing

Workloads

Agentic AI is creating new types of storage workloads with unique data types, access patterns, and performance requirements

EXAMPLE

KV-Cache offload to accelerate reasoning model inference

 ©SNIA. All Rights Reserved.

Customers increasingly need data platforms to not only store, but also discover / transform / search / monitor their data for AI

User Workflow	Discover Data	Curate Structured Data	Curate Unstructured Data	Search on Data	Prepare Datasets for Fine Tuning	Monitor and Optimize Data Lake
Needs	"How can I search my data sources for the data that I need?"	"How do I query structured data using LLMs?" "How do I move out of Hadoop?" "How do I combine data from disparate sources?"	"How can I transform data for GenAI?" "How do I parse a document into something usable?"	"How do I talk to my data using an LLM? (Inferencing)" "How do I keep the vector DB updated?" "How do I enforce access controls?"	"How do I feed a subset of data to GPUs to fine tune or train a model?"	"How do I optimize my data estate, correlate events and estimate future needs?"

 ©SNIA. All Rights Reserved.

Agentic AI growth drives expansion in context-building workloads.

Agents A & B **designed to achieve different objectives** using different system prompts and tools

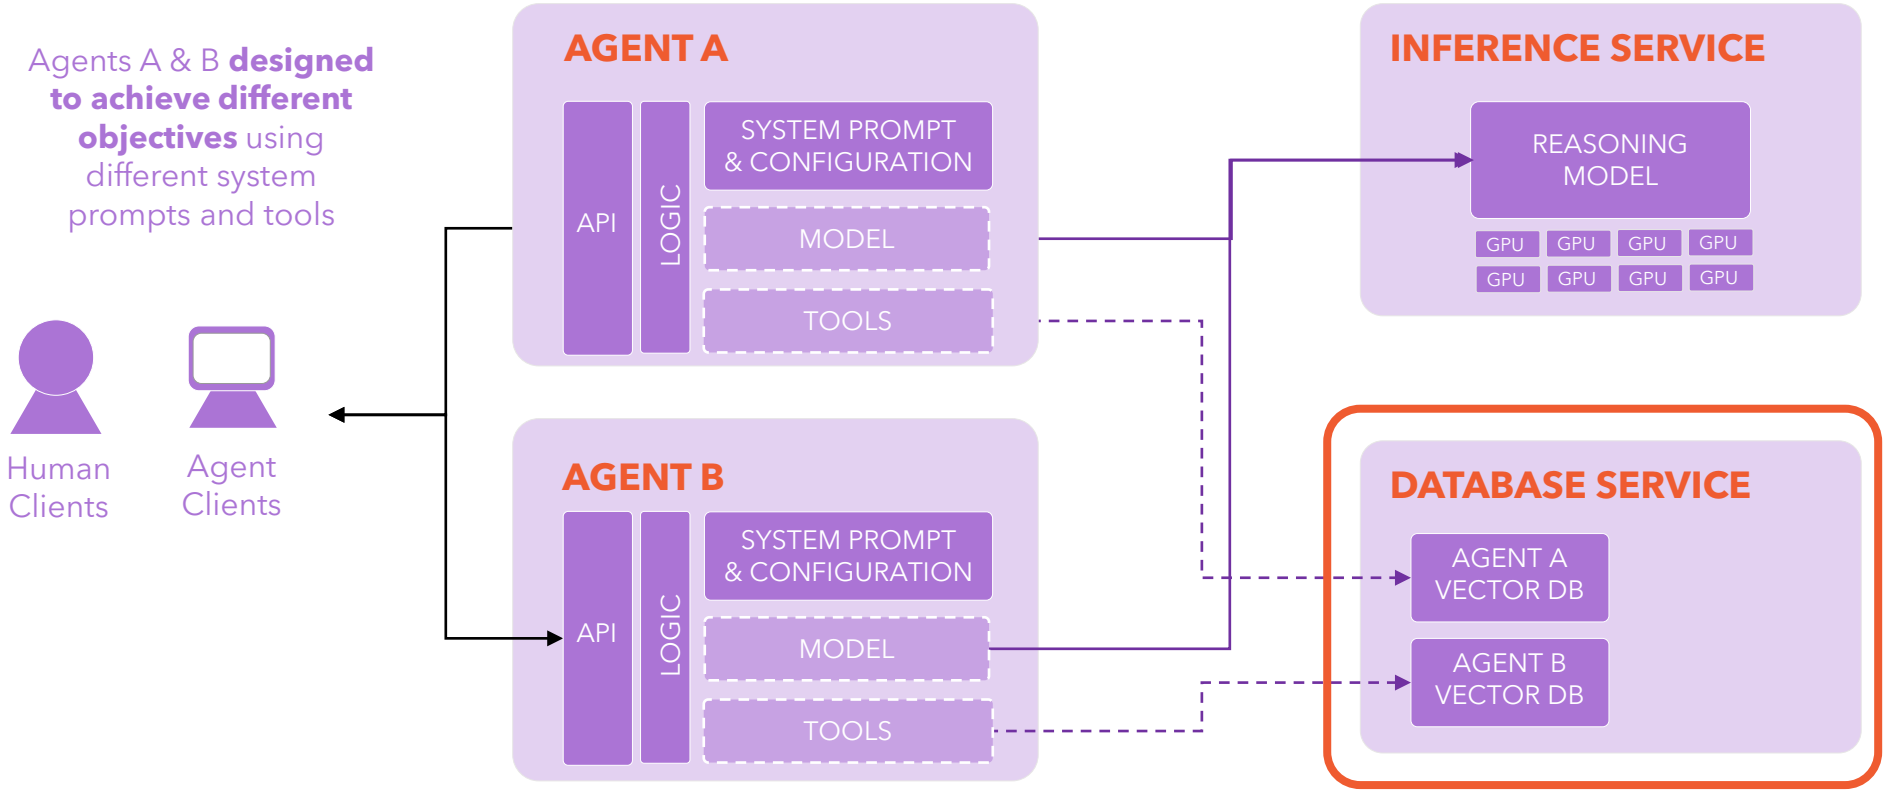

Capturing data context now requires multiple representations of **enterprise knowledge**

Deriving and maintaining context materially increases storage utilization, this "context layer" **can drive ~7×-10× more stored data**

These context + reasoning workflows introduce new I/O patterns, **mixed workload profile**

 ©SNIA. All Rights Reserved.

AI Data Preparation – Data Enrichment

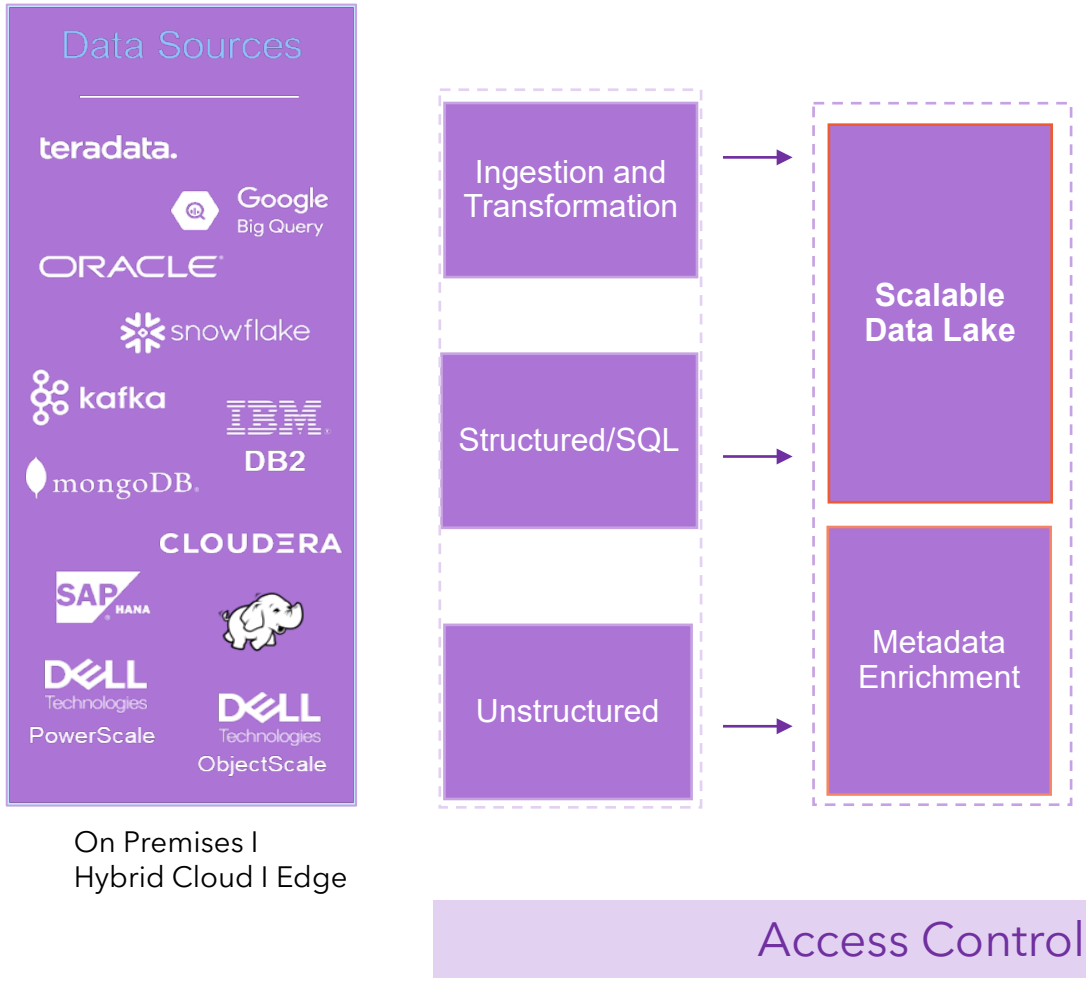

Metadata Enrichment

Adds semantic labels, classifications, timestamps, relationships, geolocation, entity tags, and more.

Tailored Actions

When data is enriched with attributes like user intent, behavioral patterns, semantic meaning, or environmental cues, agentic AI can tailor actions more effectively.

 ©SNIA. All Rights Reserved.

AI Data Preparation – Curate and Transform

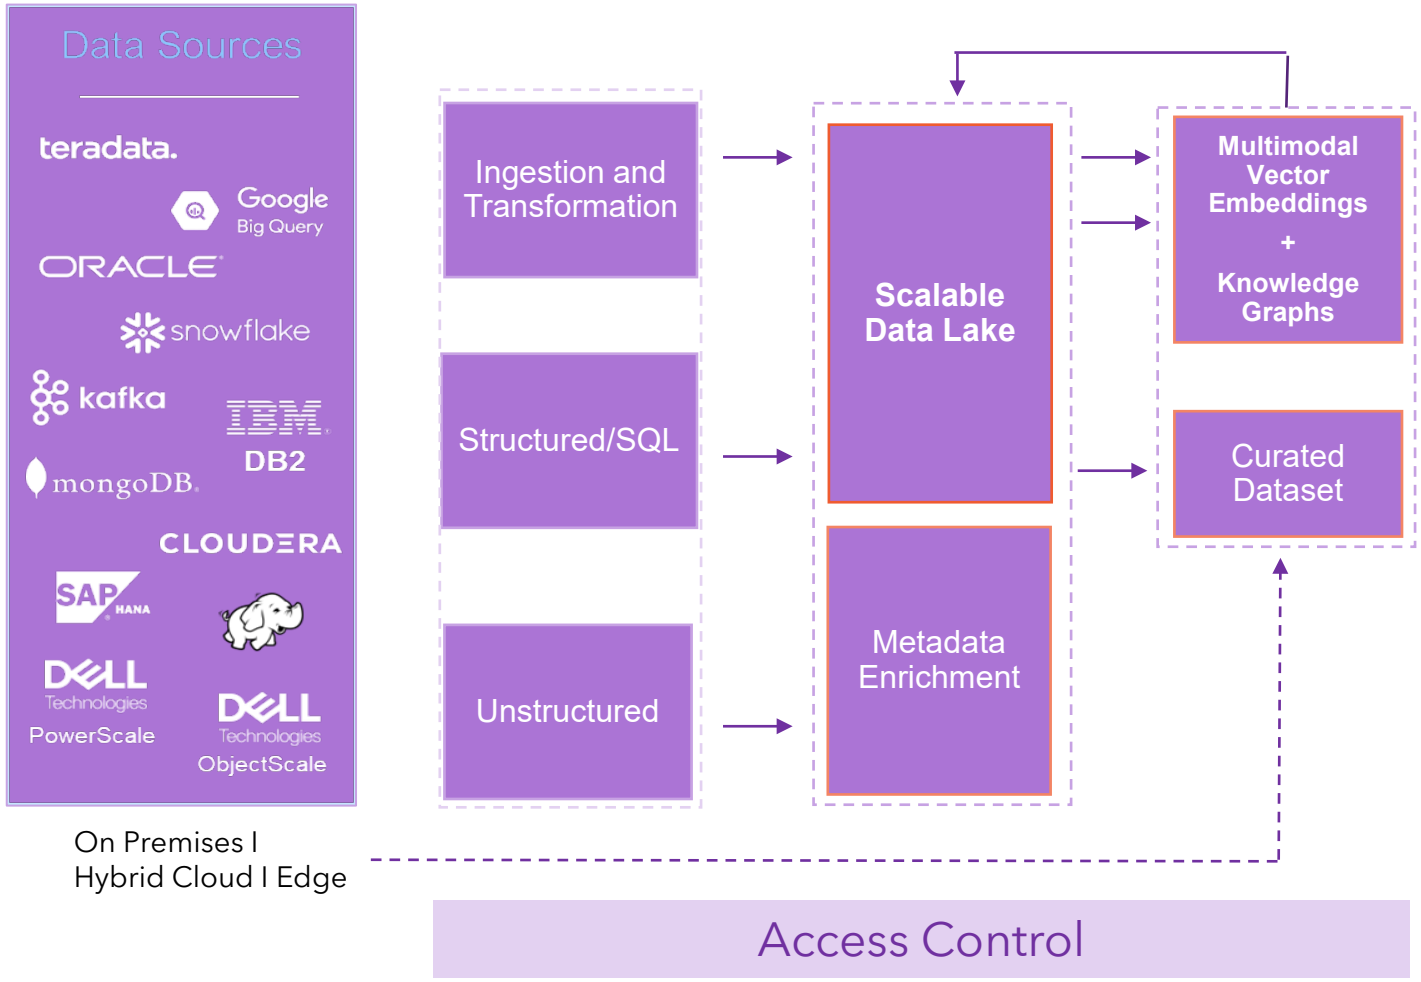

Curated Data

Data Quality is most important aspect to get accurate outcomes. Curate relevant data leveraging enriched data.

Contextual data

Transform raw data into embeddings that captures semantic context of the data.

Data Explosion

Generated data to be ~10x larger than the raw data

 ©SNIA. All Rights Reserved.

AI Data Access – Standardize way - MCP

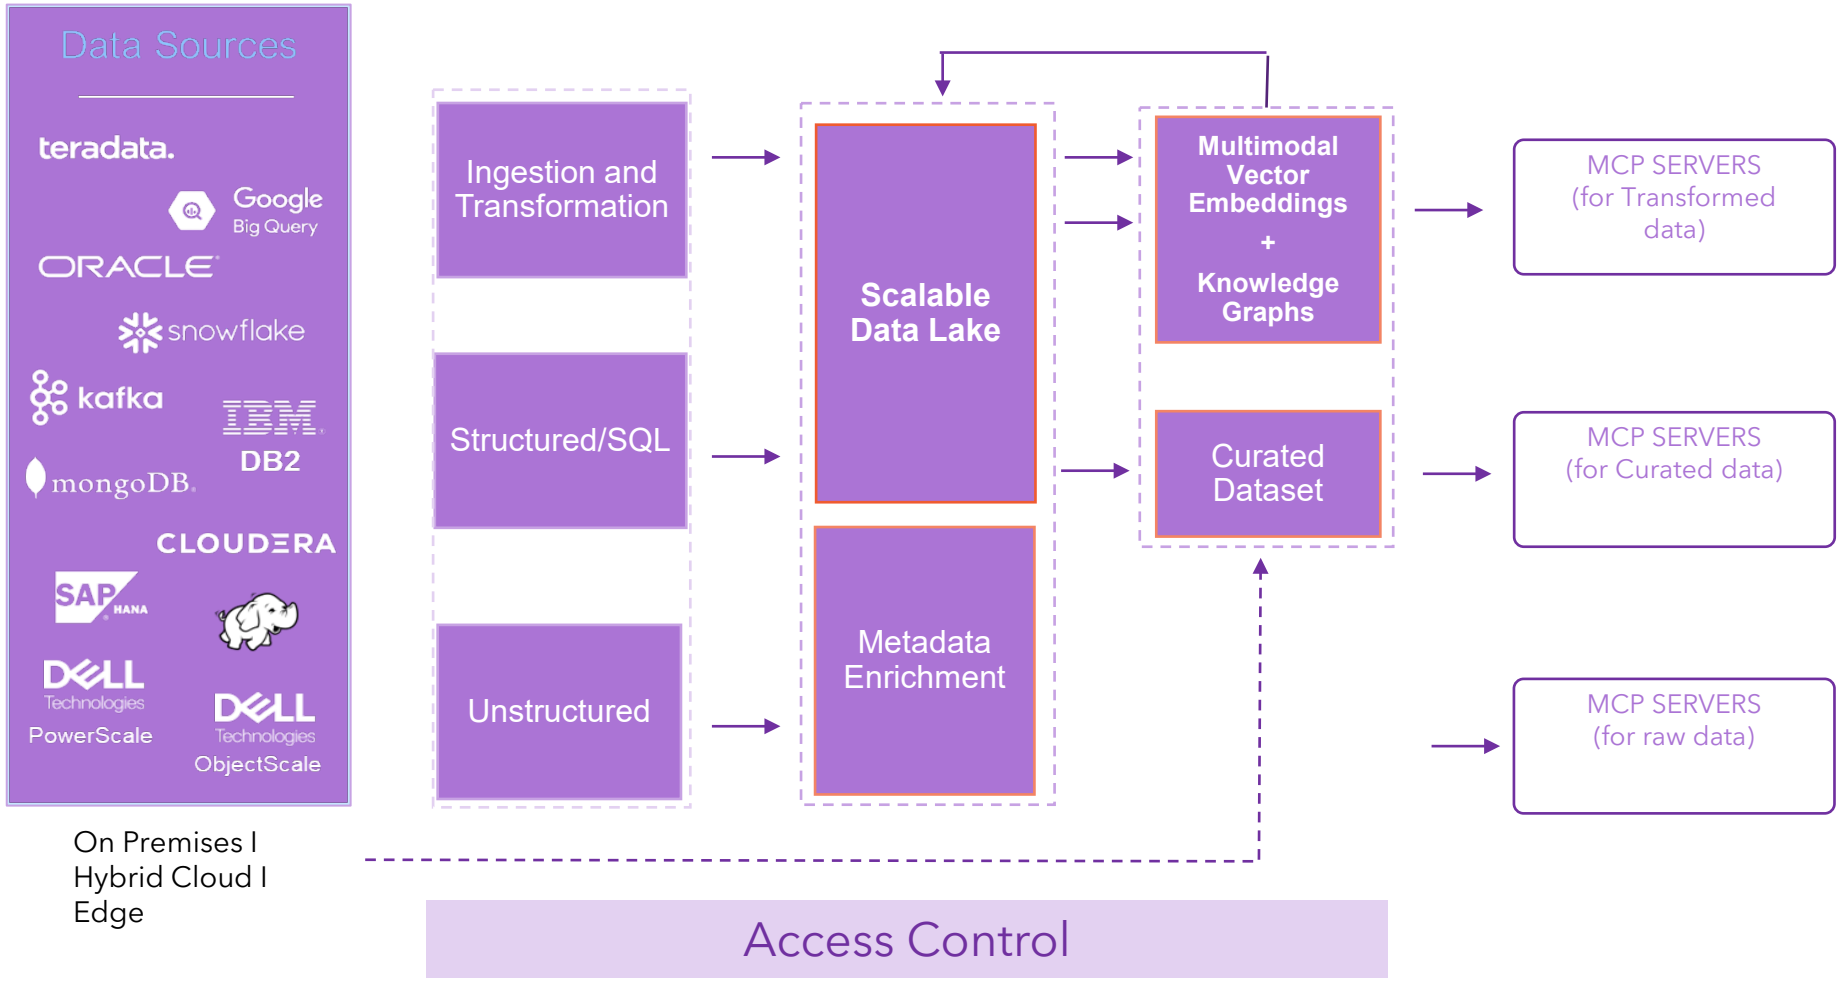

 ©SNIA. All Rights Reserved.

SNIA CST A SNIA Community

Agentic AI adoption increases, reasoning model inference workloads scale dramatically

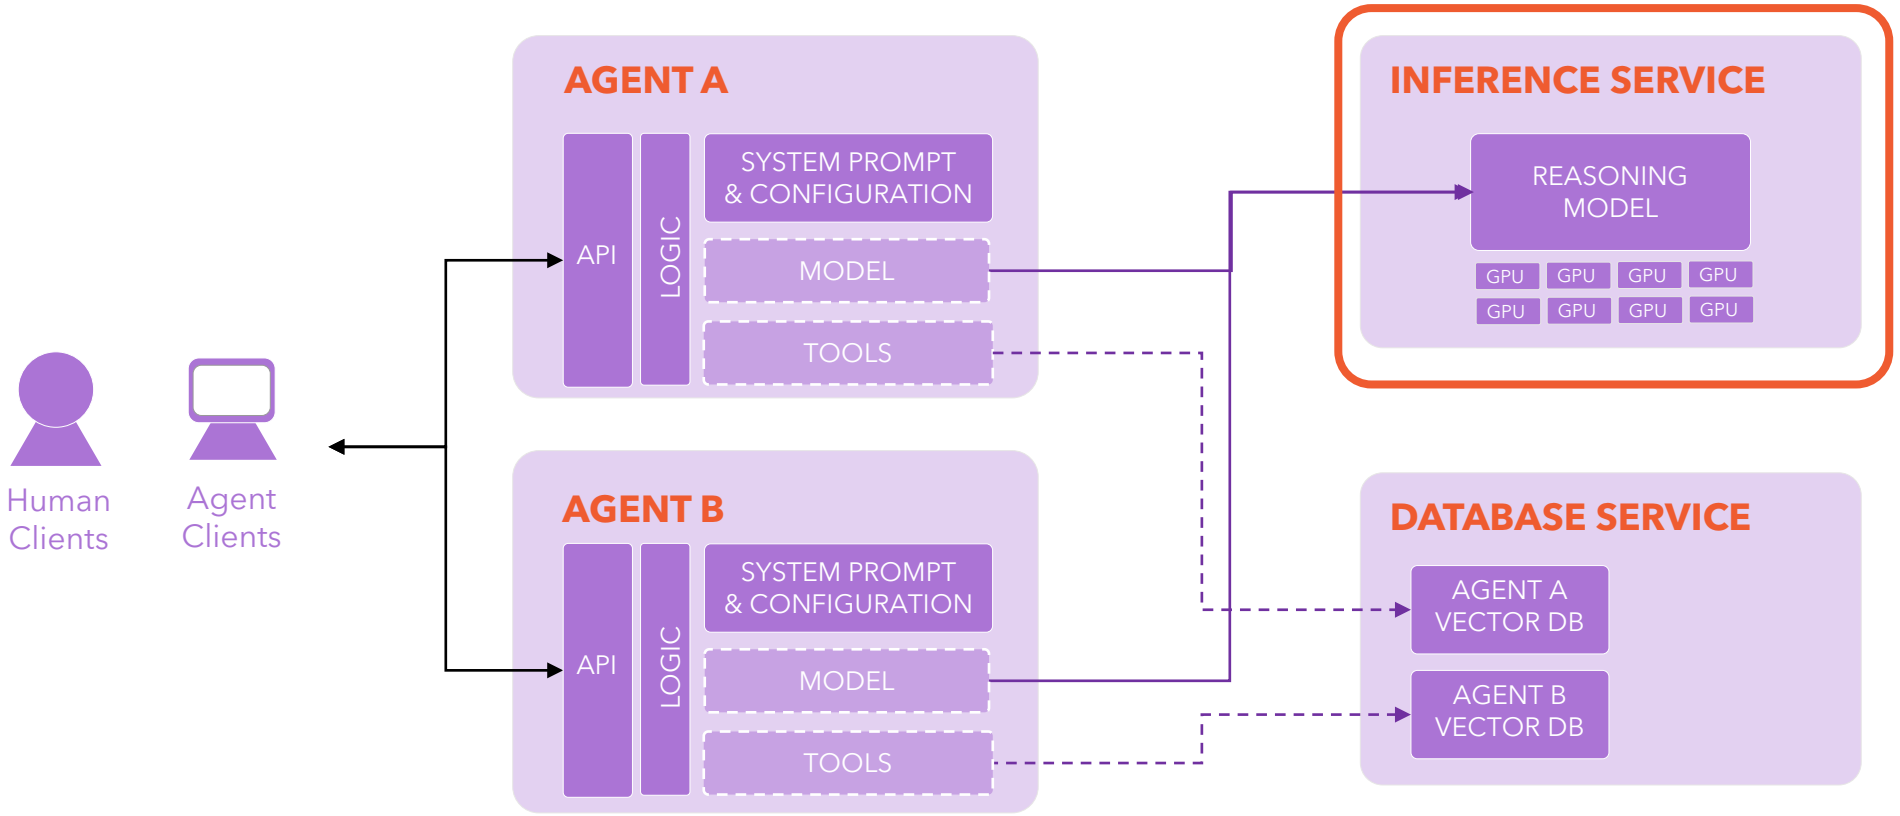

- Agents likely to use common inference services and reasoning models
- Reasoning models generate **10x to 100x more tokens** than traditional LLMs
- Aggregate reasoning workload can require substantial compute and memory resources
- Challenges scale dramatically as agentic adoption increases

 ©SNIA. All Rights Reserved.

Storing KV-Cache on shared storage provides many benefits

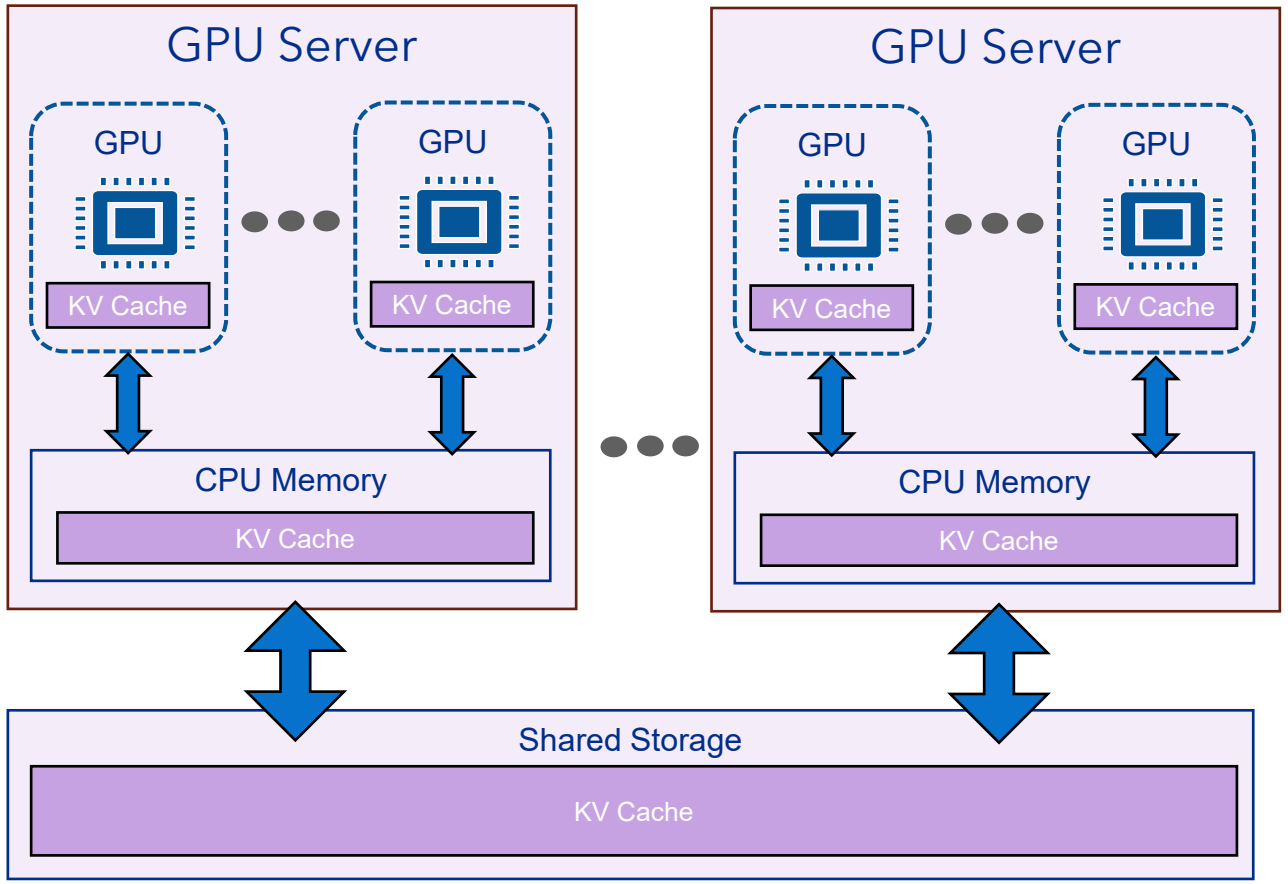

Reference

[1] Day 6: One More Thing, DeepSeek-V3/R1 Inference System Overview

Inference Sessions Scalability

Allows session isolation and scaling beyond memory limits. Providing long-term session persistence

Distributed Access

Enables cross-node access in distributed GPU architectures. Eliminates redundant data duplication across nodes

Optimized Access

Supports specialized Prefill and Decode GPU servers

 ©SNIA. All Rights Reserved.

Agentic AI: Storage providing raw data to transformed data to semantic data

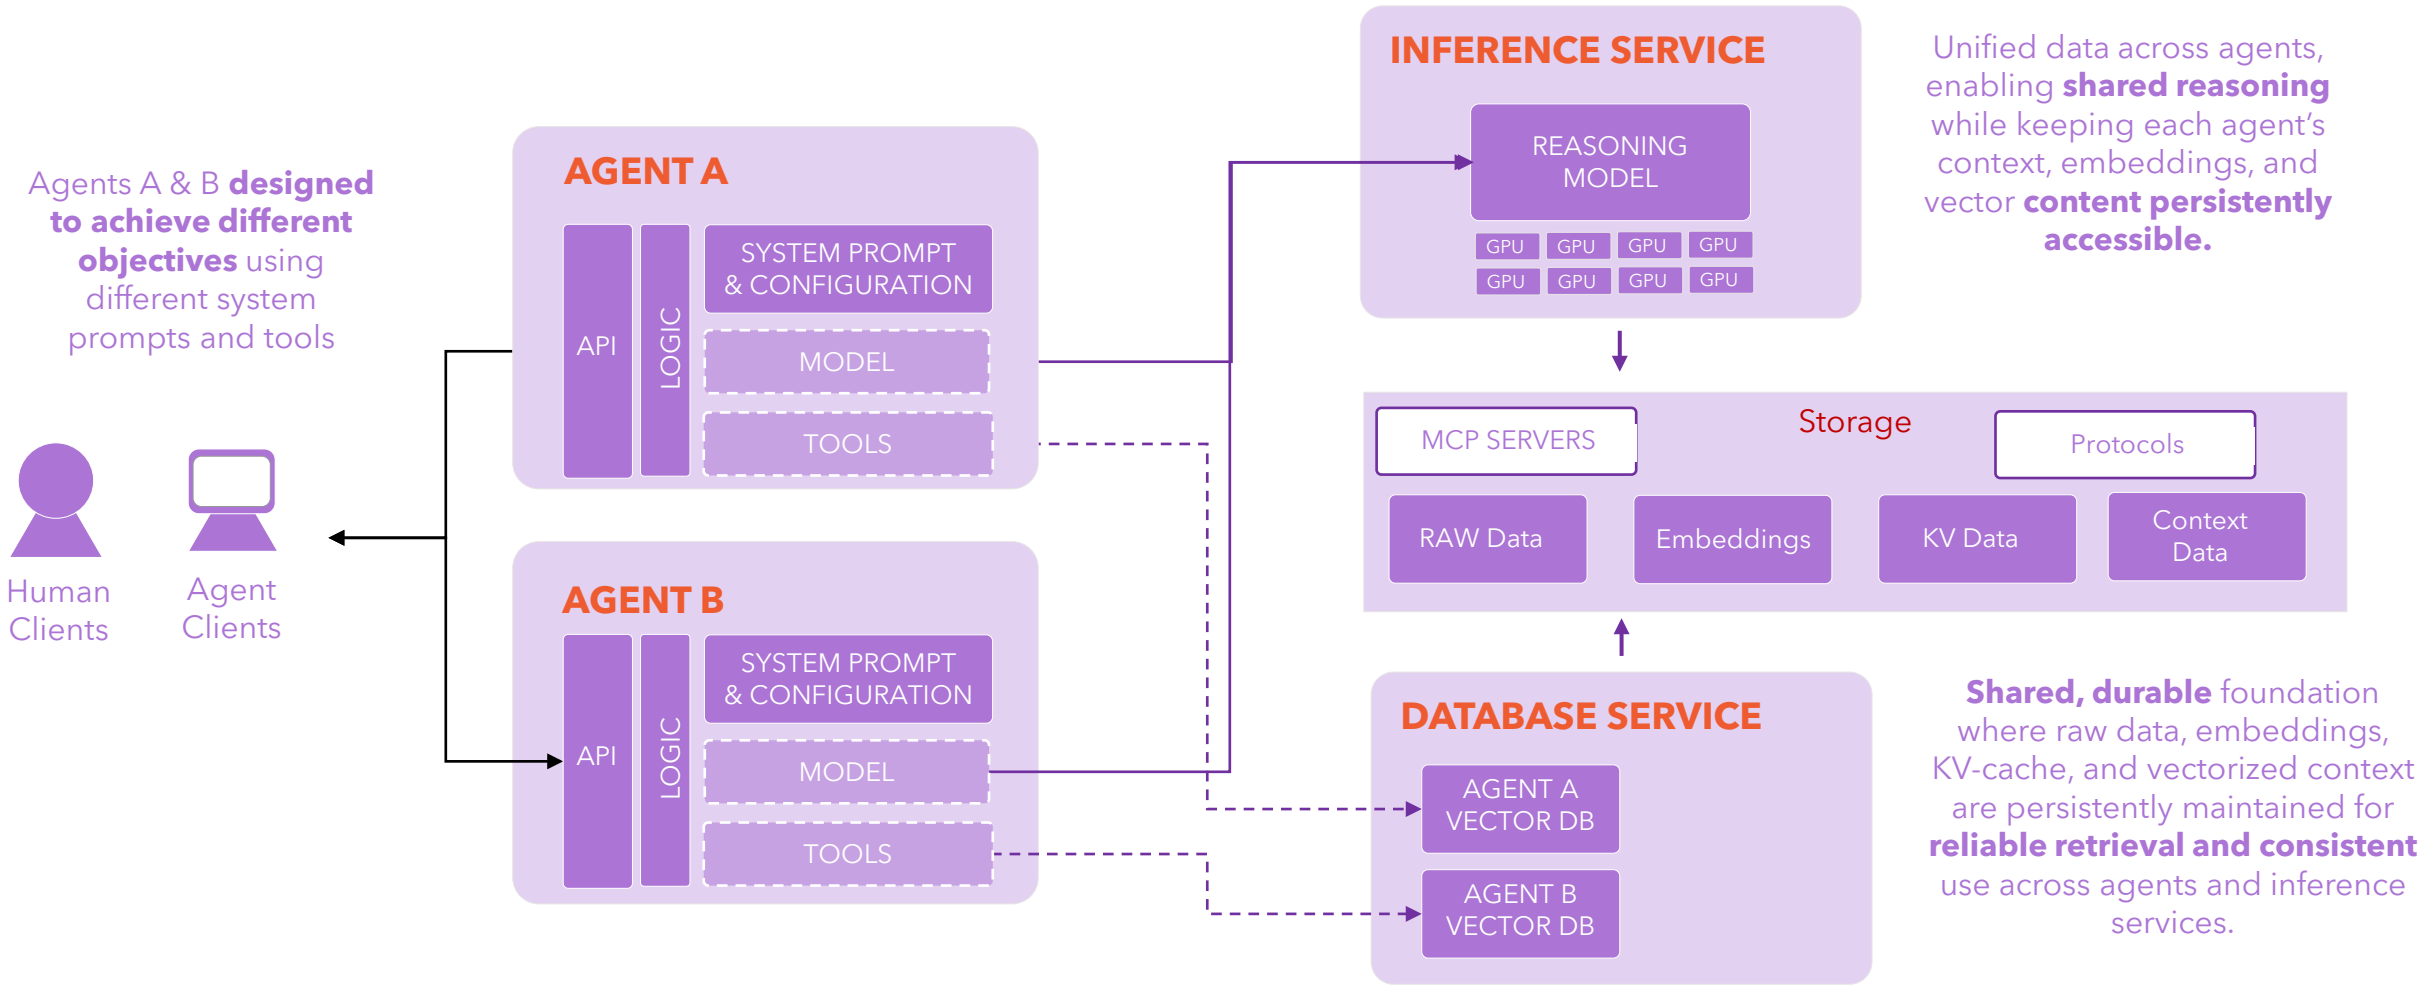

 ©SNIA. All Rights Reserved.

Summary

Agentic AI is a rapidly advancing technology that stands to impact enterprise storage across multiple dimensions

- Agentic AI is the wave of AI technology evolution
- Agents' ability to autonomously reason and act enable them to automate an ever-expanding range of complex tasks
- Diverse and distributed multi-agent systems enabled by interoperability protocols like MCP and A2A will be needed to solve complex problems
- Enabling agents to discover, process, and retrieve data from storage in real-time based on task-specific needs will require semantic data awareness and new interfaces
- Storage will play a key role as Agentic AI continues to evolve

 ©SNIA. All Rights Reserved.

Q&A

Thanks for Viewing this Webinar

- Please rate this webinar and provide us with feedback
- This webinar and a copy of the slides are available at the SNIA Educational Library
- A Q&A blog from this webinar will be posted at snia.org/blog
- Follow us on LinkedIn and X for future webinar dates

THANK YOU!

 ©SNIA. All Rights Reserved.