



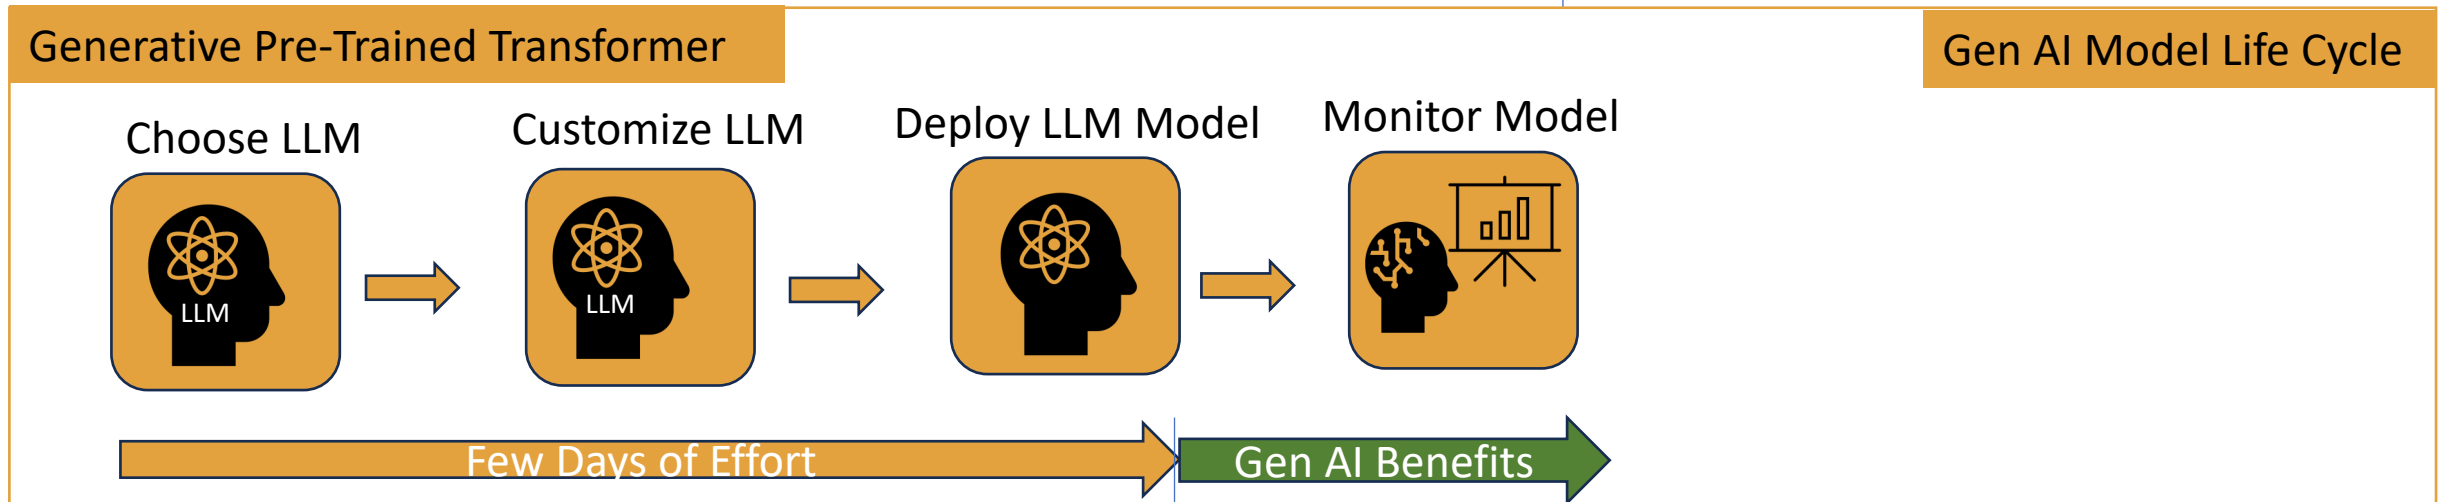
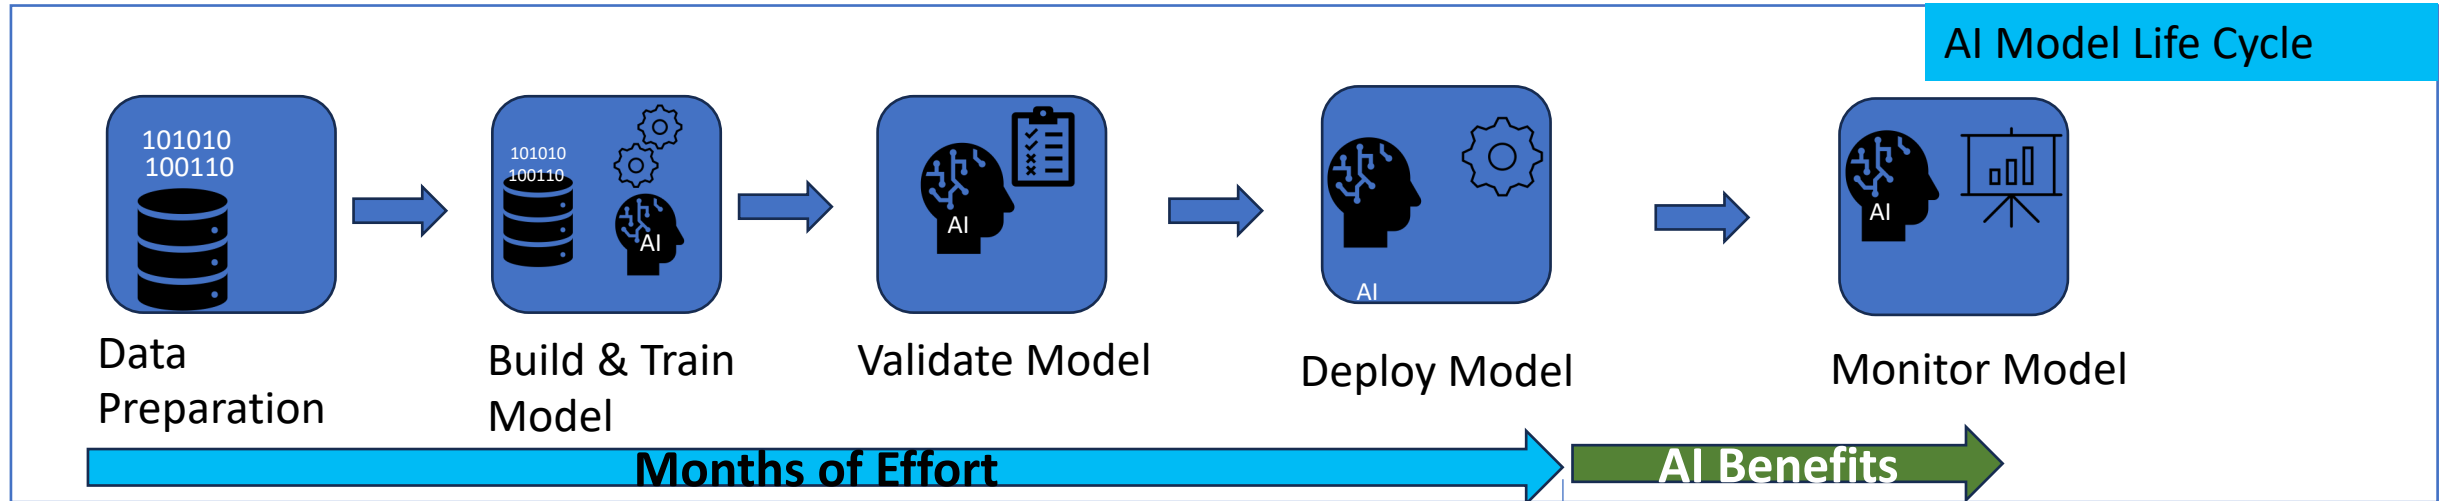
Generative AI foundation to Real World use cases

Prasad Venkatachar

Sr Director – Product & Solutions @ Pliops
IEEE Sr member & BCS Fellow



AI VS Gen AI Life Cycle



Key Pillars to Generative AI



**Enterprise Use Cases &
Business Outcomes**



**Gen AI Stack
Hardware + Software**



**Security & Regulatory
Compliance**



**Finops for Gen AI
Cost of Development &
Serving
Cost of Scaling**



**New Trends &
Upto-Date**



Gen AI Opportunity Everywhere

Industry Verticals


Banking & FSI



Retail



Content & Media



AI High Performers Response

Horizontal Functions

HR & Finance



Marketing

Customer Support

Reduce Costs in Core Business

Create New Business and/or sources of Revenue

Increase Revenue from Core Business

Increase value of Offerings by Integrating AI-Based Features or Insights

■ Respondents at AI High Performers

19%

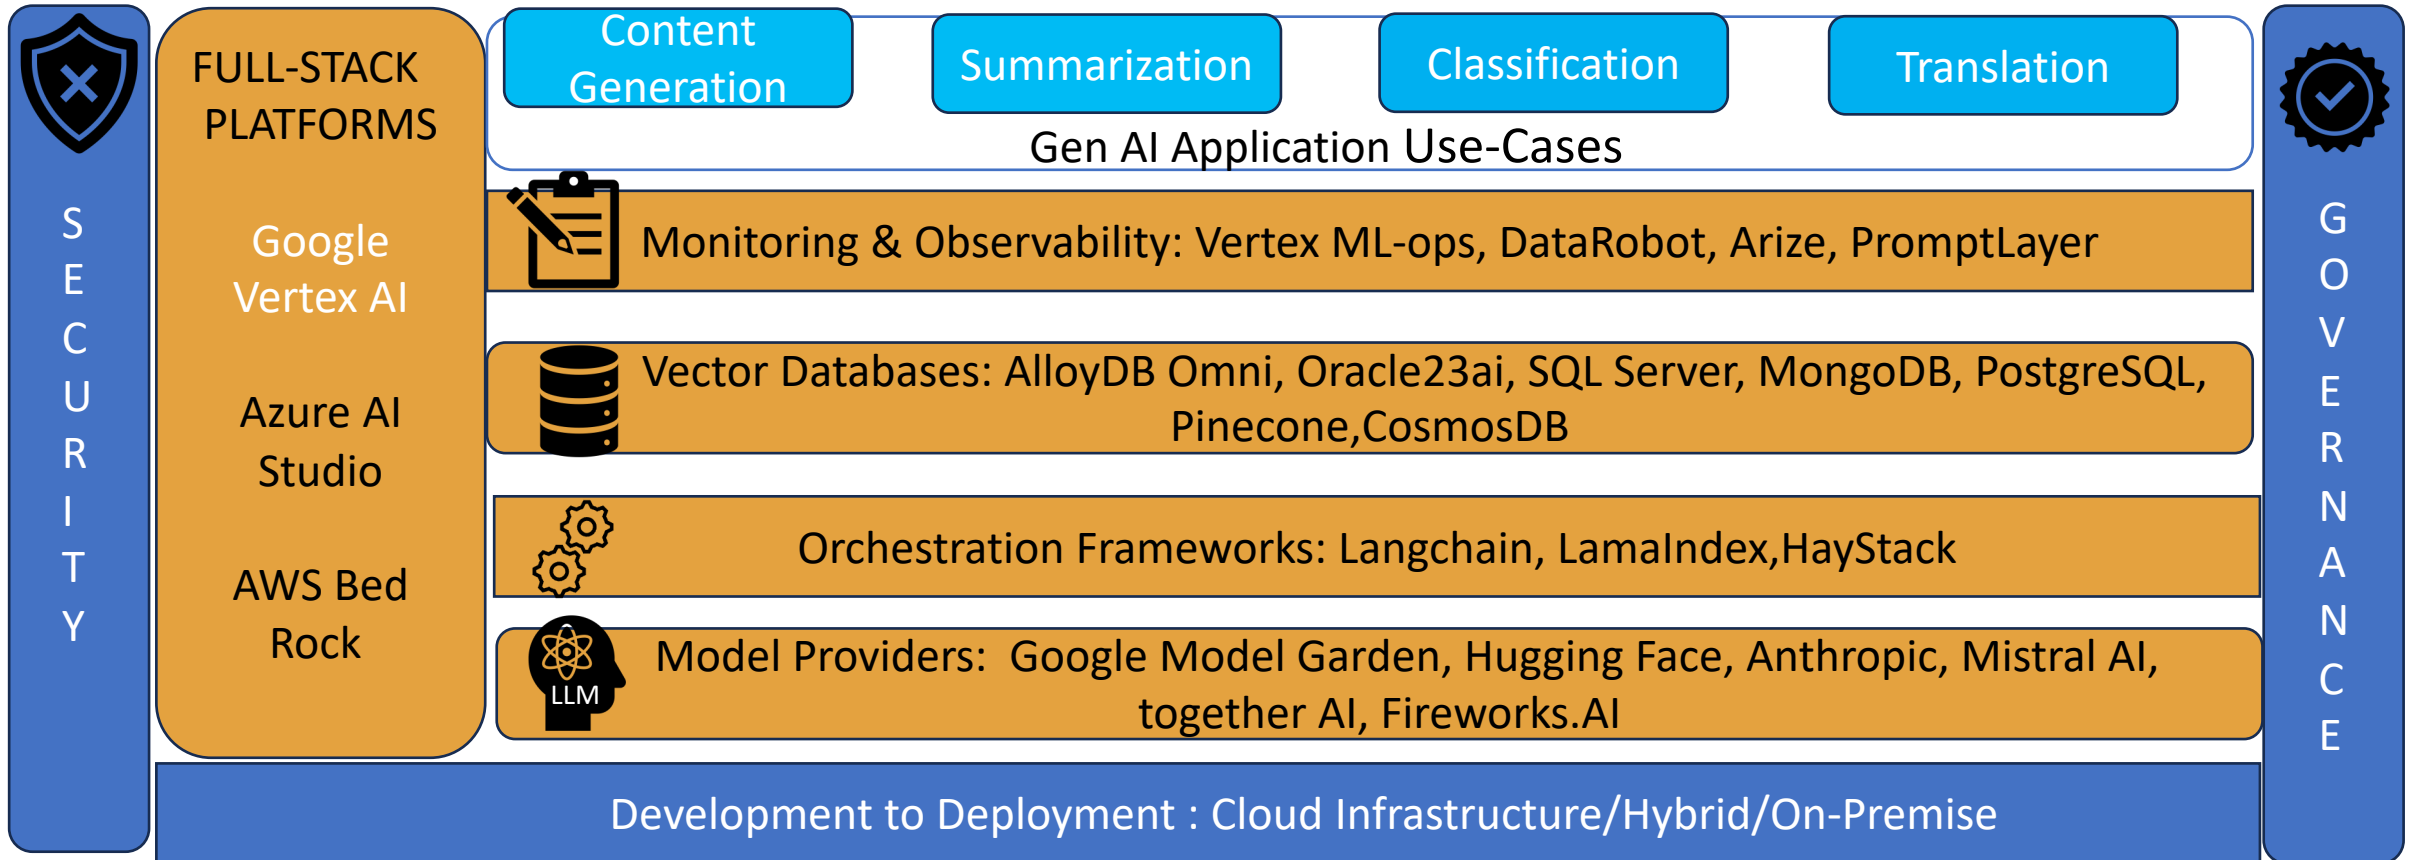
23%

27%

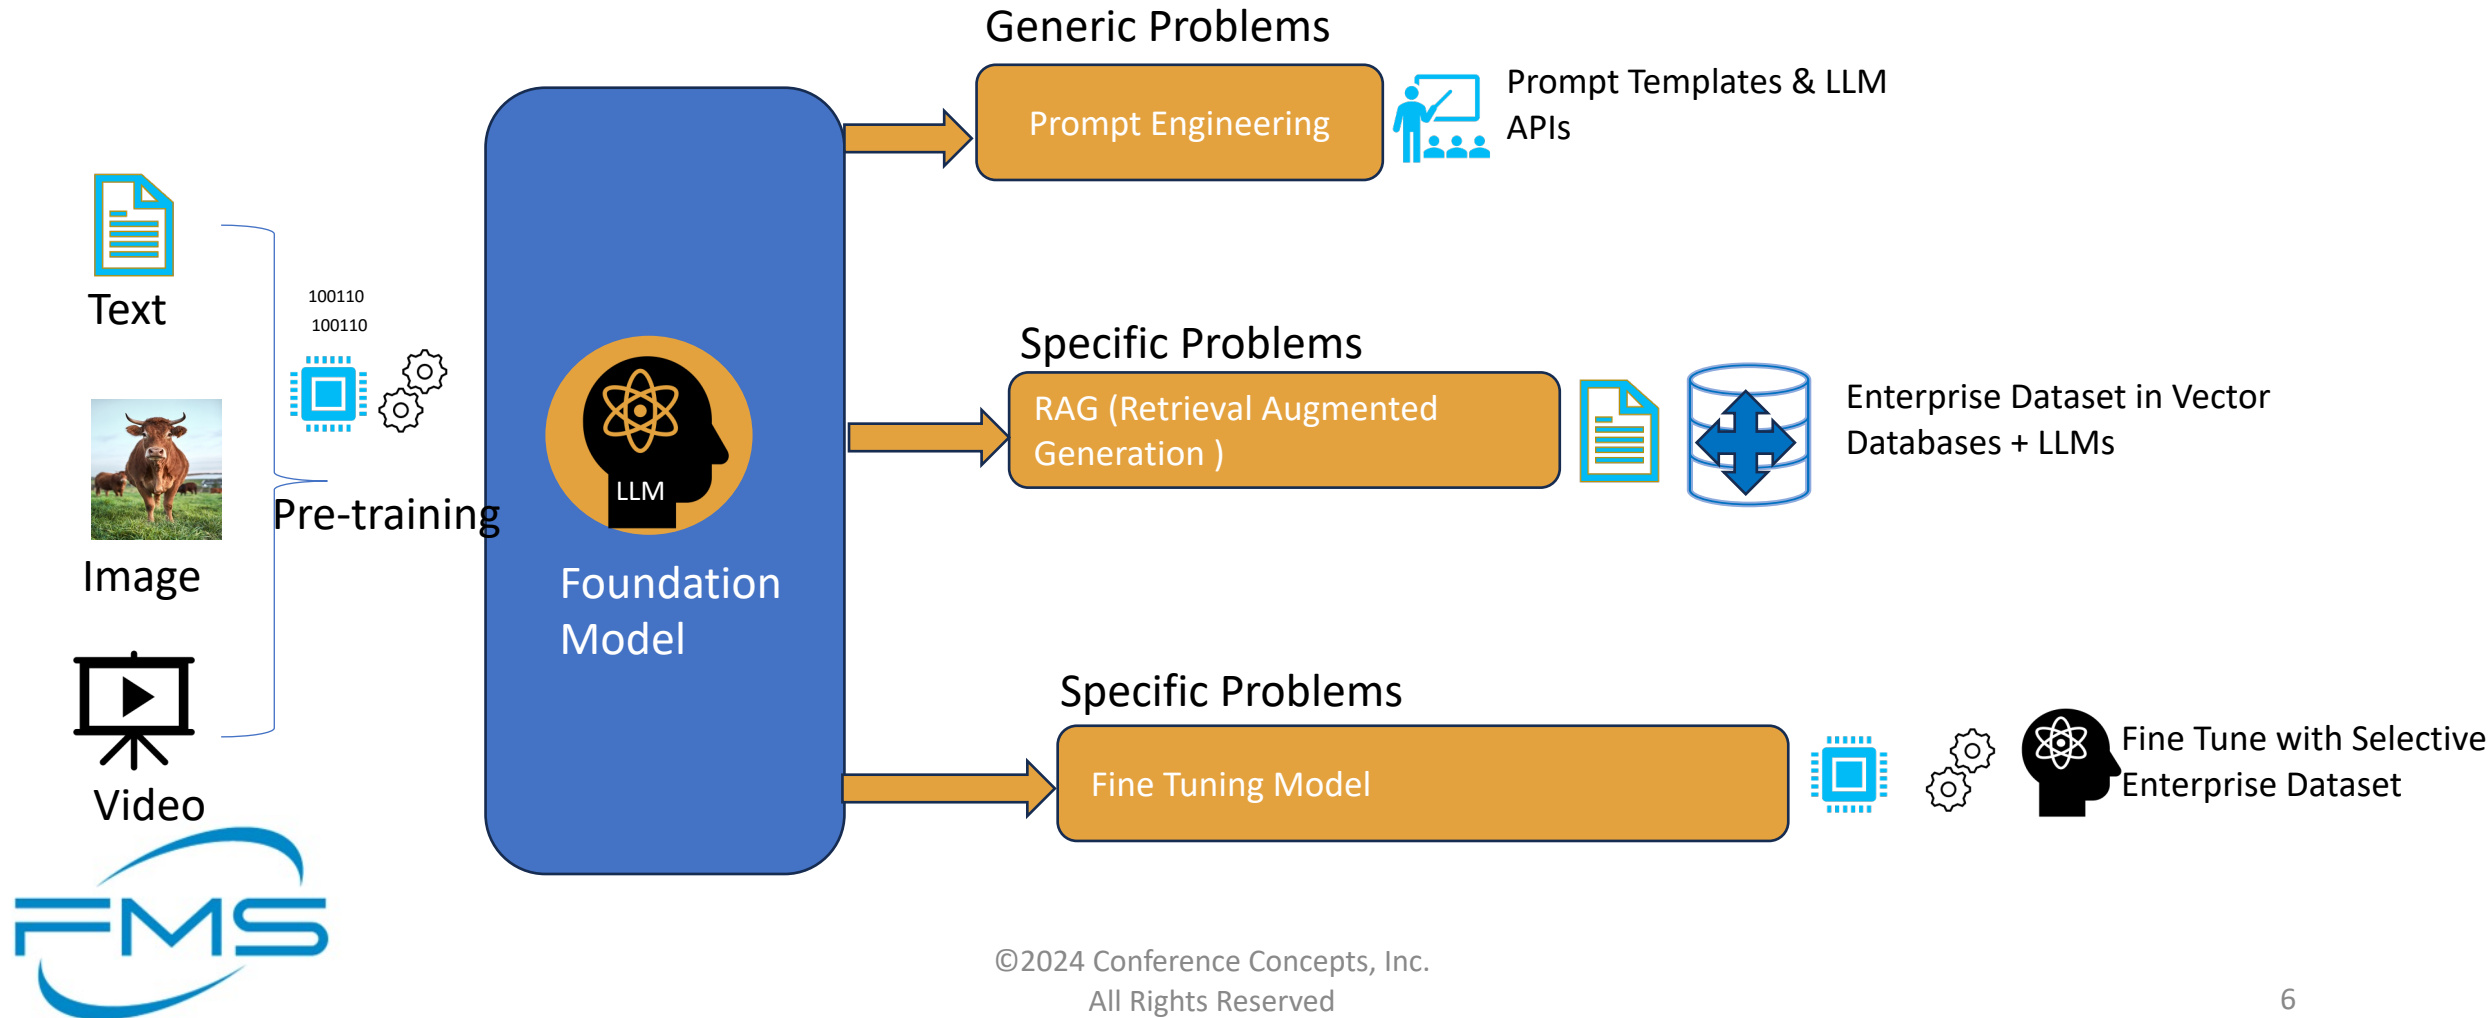
30%



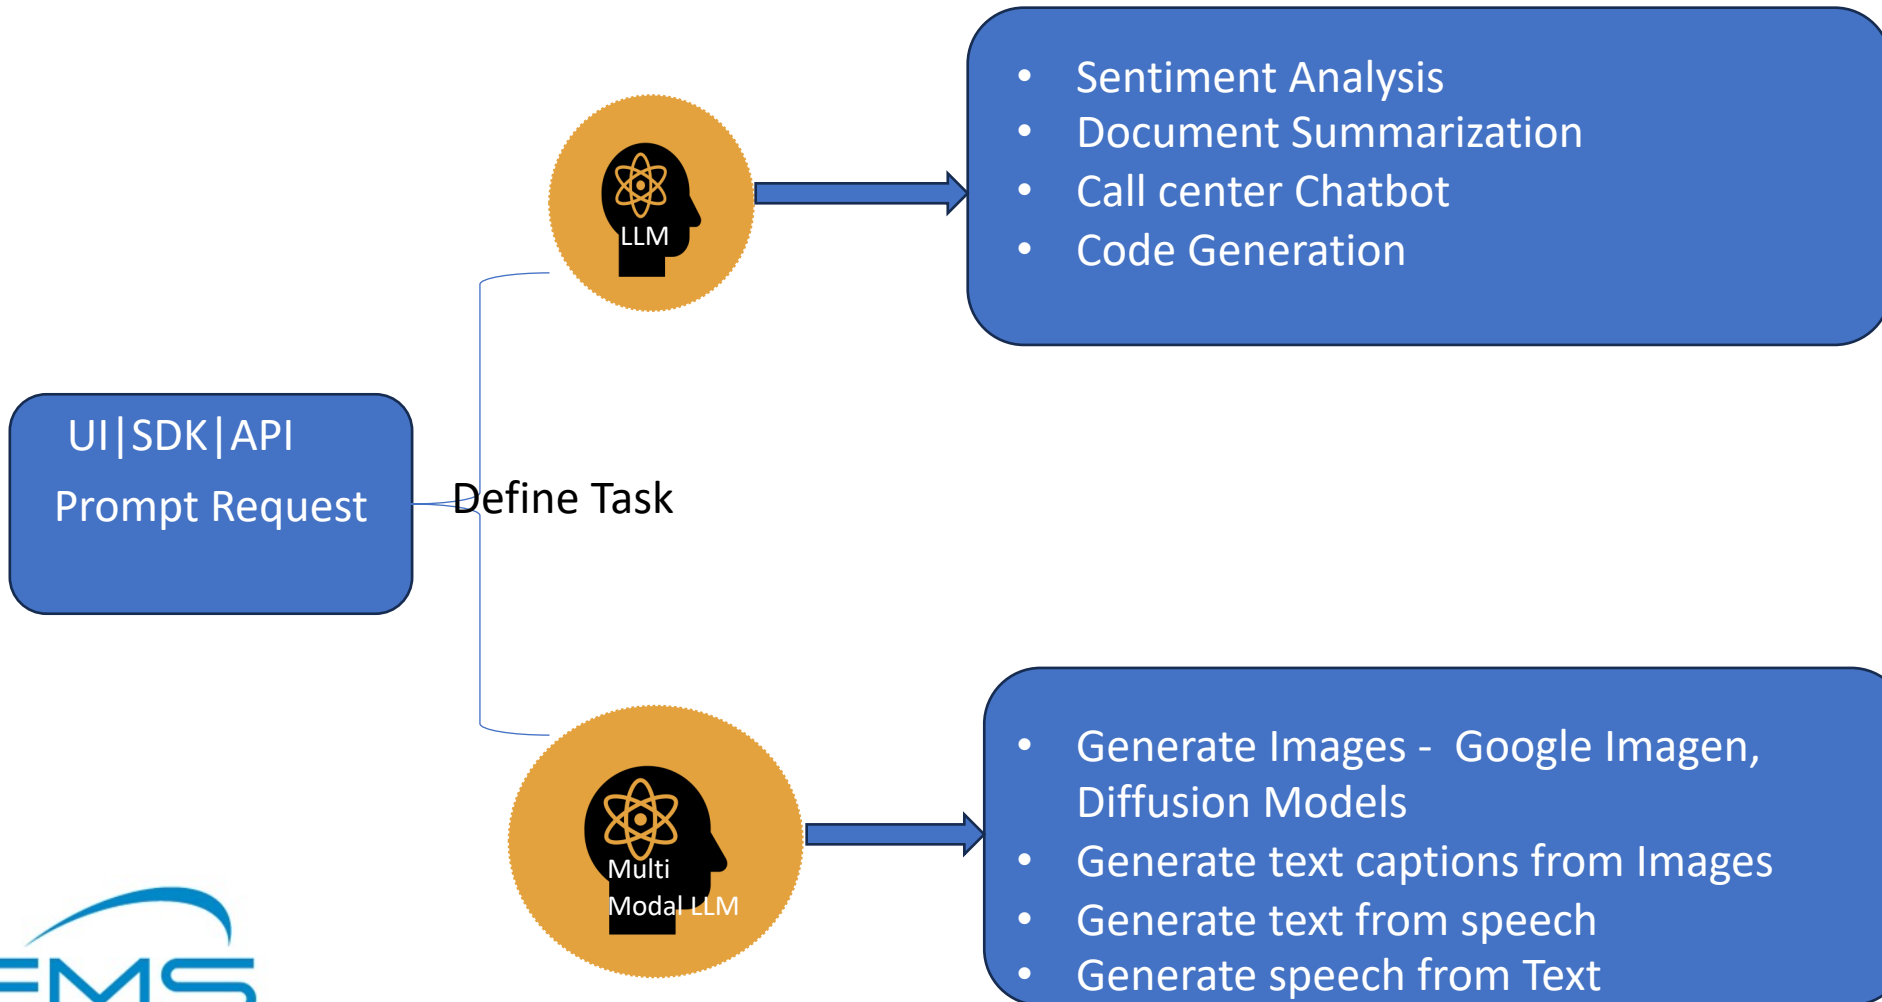
Gen AI Stack



Gen AI : Adoption Maturity



Prompt Engineering – Use cases



Electric Bike Chatbot Example

Best Seller



```
ElectricBikes = [  
{  
  "brand": "VELECTREC",  
  "type": "Mountain",  
  "features": {  
    "motor": "1000w",  
    "Max Range(Miles)": "40-50",  
    "Max Speed(MPH)": "25",  
    "Number of speeds": "7",  
    "brakes": "Front & Rear Disc Brakes",  
    "Charging Time(HRS)": "5-6",  
    "frame Type": "Aluminum",  
    "suspension": "Dual",  
    "tires": "27.5-inch all-terrain",  
  },  
  "price": 699,  
  "Sold by": "velectric inc"  
},  
]
```

E-Bike Conversation Assistant:

List Bikes: Provide list of latest models bikes matching customer query: Price, Range, Charging Time etc

addToOrder: If the Bike on the list do addToOrder, then summarizeOrder, then confirmOrder.

changeItem: Replace bike with another type: Range, Price

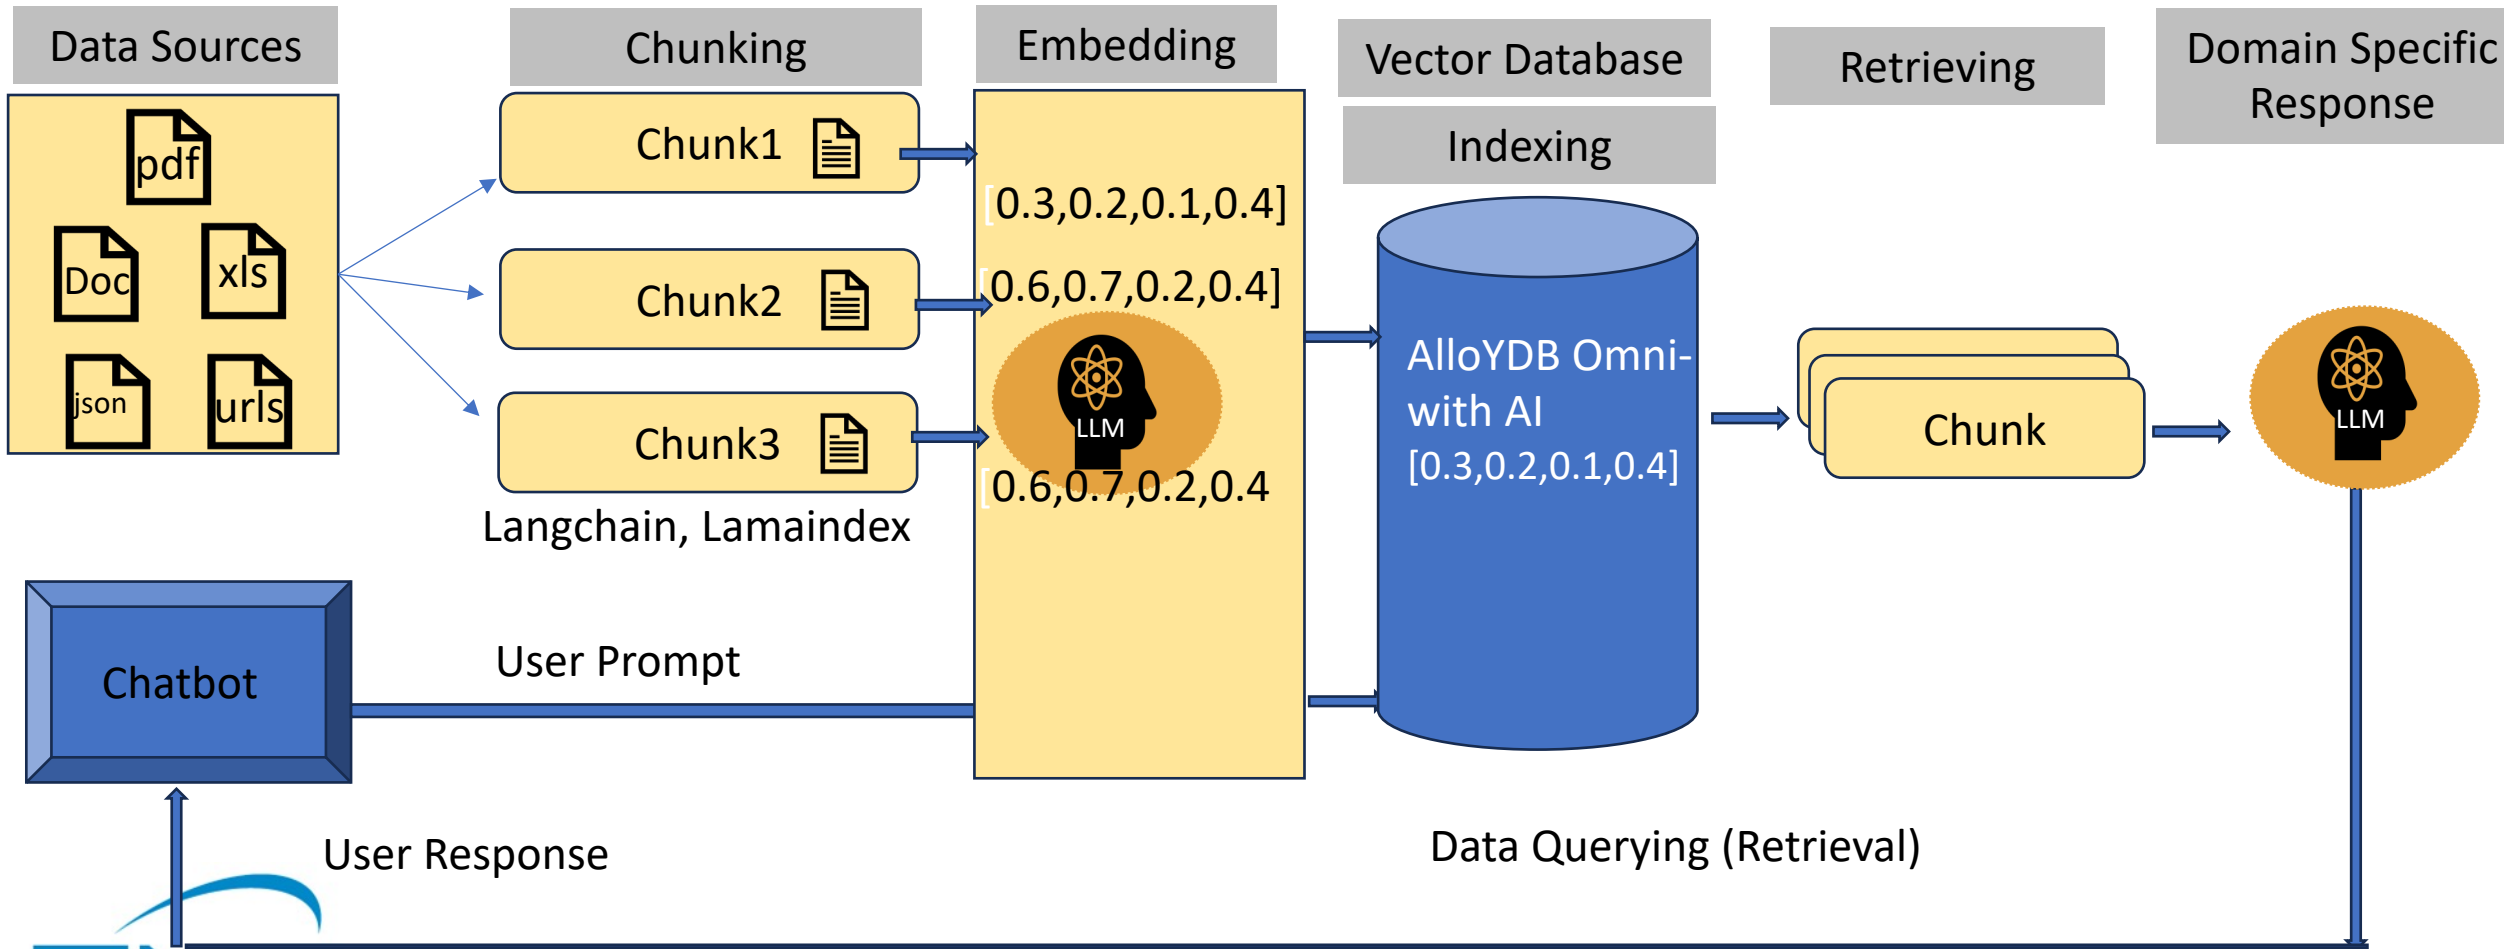
summarizeOrder: If the customer has added to the order, list each menu item and modifier added to the order. If there has been nothing ordered, redirect.

confirmOrder: Ask the customer to confirm the order details are correct.

finishOrder: Tell the user the order has been sent to the Delivery



AlloyDB AI – Vector Database

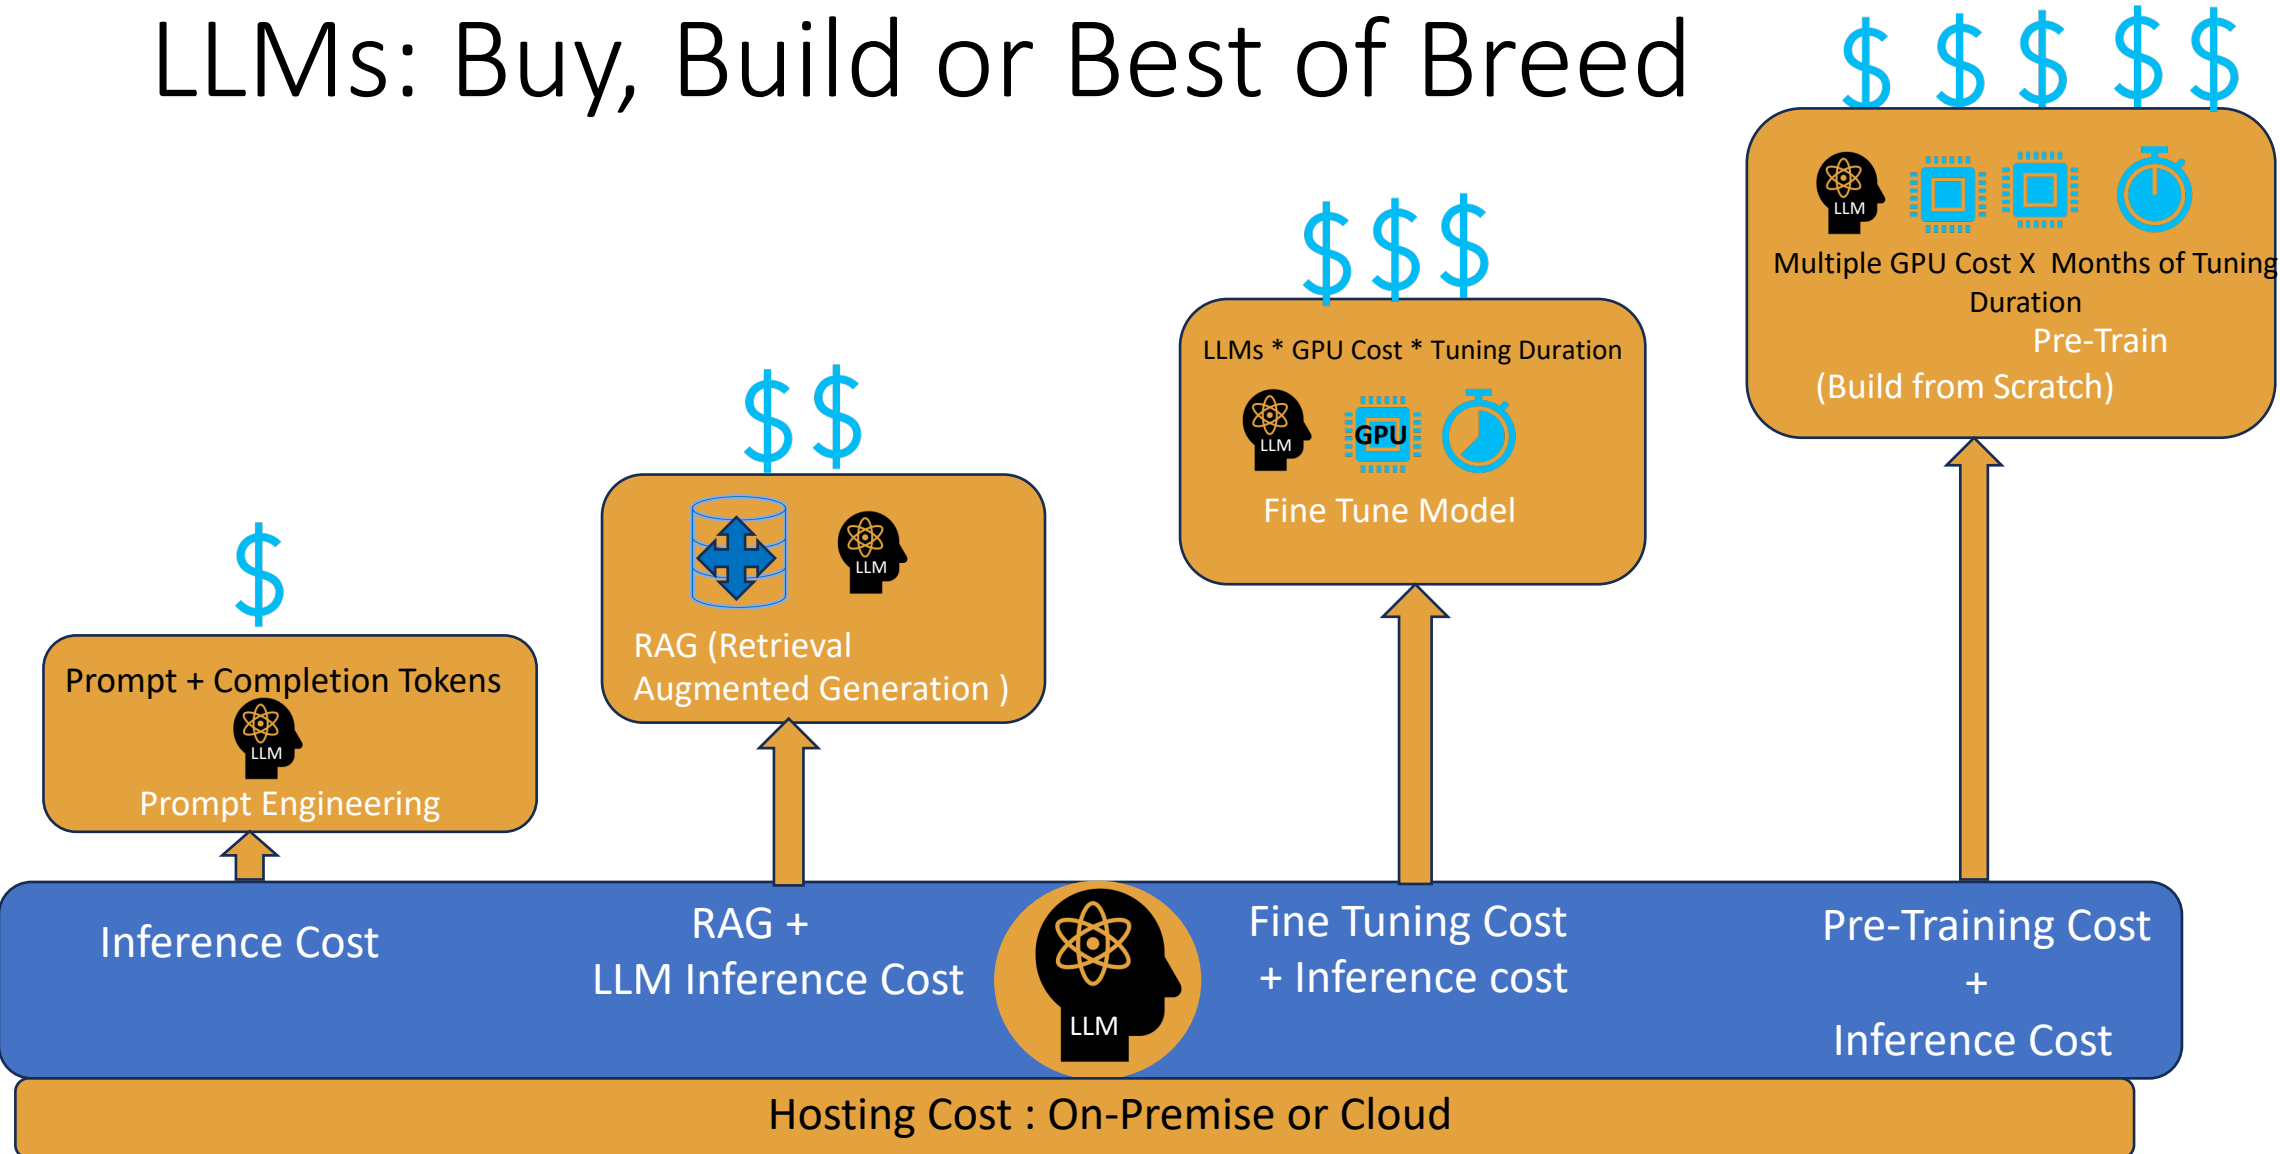


RAG Architecture

©2024 Conference Concepts, Inc.
All Rights Reserved



LLMs: Buy, Build or Best of Breed



One Platform: OLTP/Analytics/Gen AI

E-Commerce Applications



Real Time Biz Analytics



Gen AI Chatbots



 AlloyDB Omni

 AlloyDB



Lenovo

XDP Data Accelerator + SSDs



Lenovo Edge Server



XDP Data Accelerator + SSDs



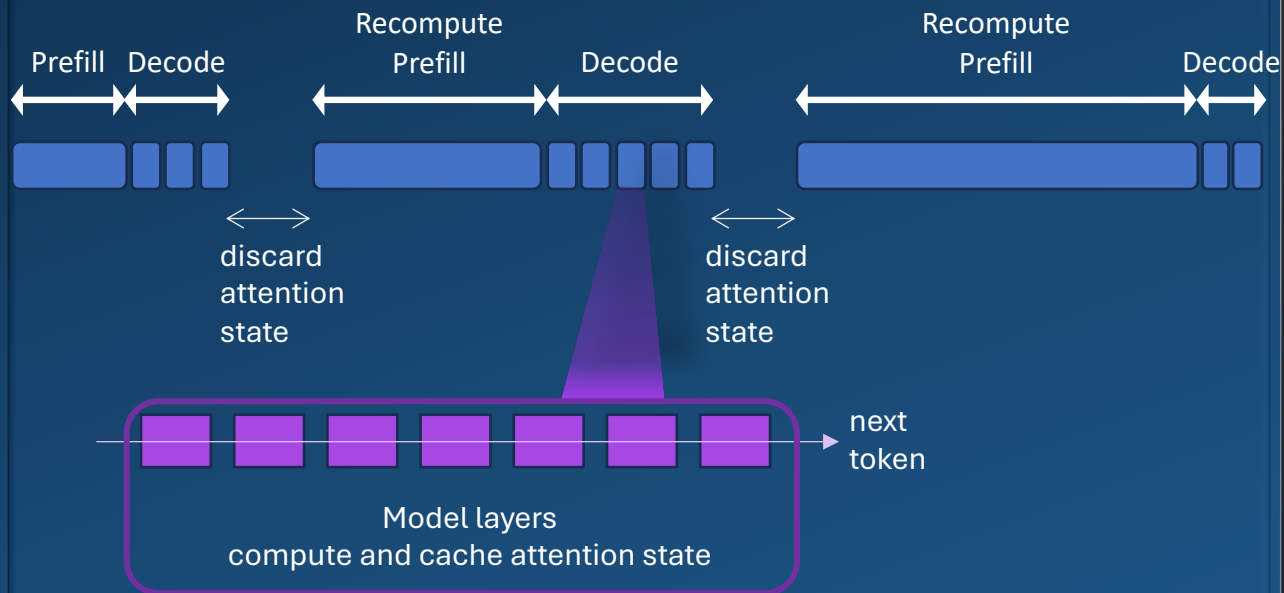
Lenovo Datacenter Server



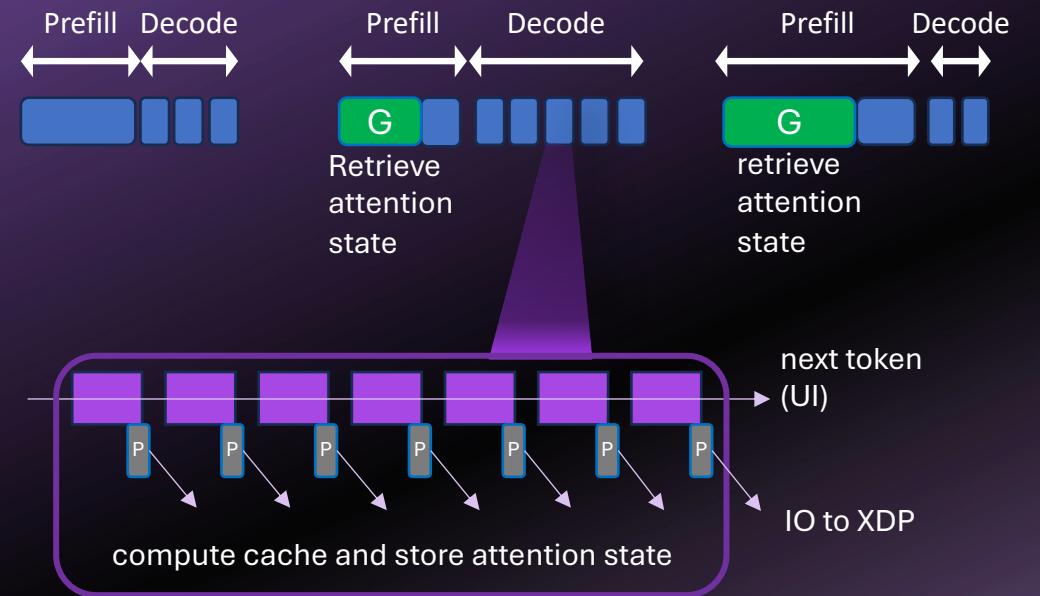
Long-Context Attention

conversation or multi-shot task agent prefills attention score based on expanding history

Today: Compute-Based Prefill



Pliops LightningAI-Based Prefill



Summary & Q/A

- Gen AI is rapidly adopted in Enterprises
- Identifying the use case and business outcomes is the key
- Building skillset for Gen AI adoption maturity is critical
- Enterprises has to focus on Fin-ops for Gen AI
 - Building the Gen AI LLM solution
 - Serving the Gen AI LLM Solution

