



# AI COMMUNITY CONFERENCE

Vancouver, BC  
October 18, 2024

## Azure AI Search and Open AI

Peter Carson

#AICommunityConference

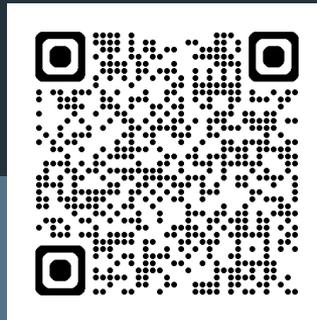


Microsoft

# Peter Carson

- President, Extranet User Manager
- 14-time Microsoft M365 MVP
- [peter@envisionit.com](mailto:peter@envisionit.com)
- [blog.petercarson.ca](http://blog.petercarson.ca)
- President, Toronto Microsoft 365 User Group

Connect with me  
on LinkedIn



AI Community Conference Vancouver 2024  
# AICommunityConference



# Thank You, Sponsors

PLATINUM

Extranet  
User Manager



BRONZE



COMMUNITY



AI Community Conference Vancouver 2024  
# AICommunityConference



# Shout Outs

Share your experience with your peers through social media. Make sure to tag one of the following to enter a raffle as a social media trailblazer

@AICommunityEvent, #AICommunityEvent



# How much time do employees spend searching for what they need for work?

**15**  
minutes

**30**  
minutes

**1**  
hour

**2**  
hours



# Agenda

- Copilot for Microsoft 365
- Build Your Own Chat
- Azure Open AI and AI Search
- Real World Demonstrations
- Closing and Questions



# Copilot



AI Community Conference Vancouver 2024  
# AICommunityConference

# Copilot

For Microsoft 365

 **Follow up strong!**

Write a follow up email to `meeting` to drive continued progress.

 **Get calendar info**

When's my next meeting with `person` ?

 **Understand quickly**

Provide a comprehensive summary of my emails about [the FY24 Sales Report]

 **Decode acronyms**

What does [LLM] stand for?

 **Help me write**

Write an email to my team about our top priorities for next quarter from `file`

 **Prep for that meeting**

Help me prepare for `meeting`

 View prompts

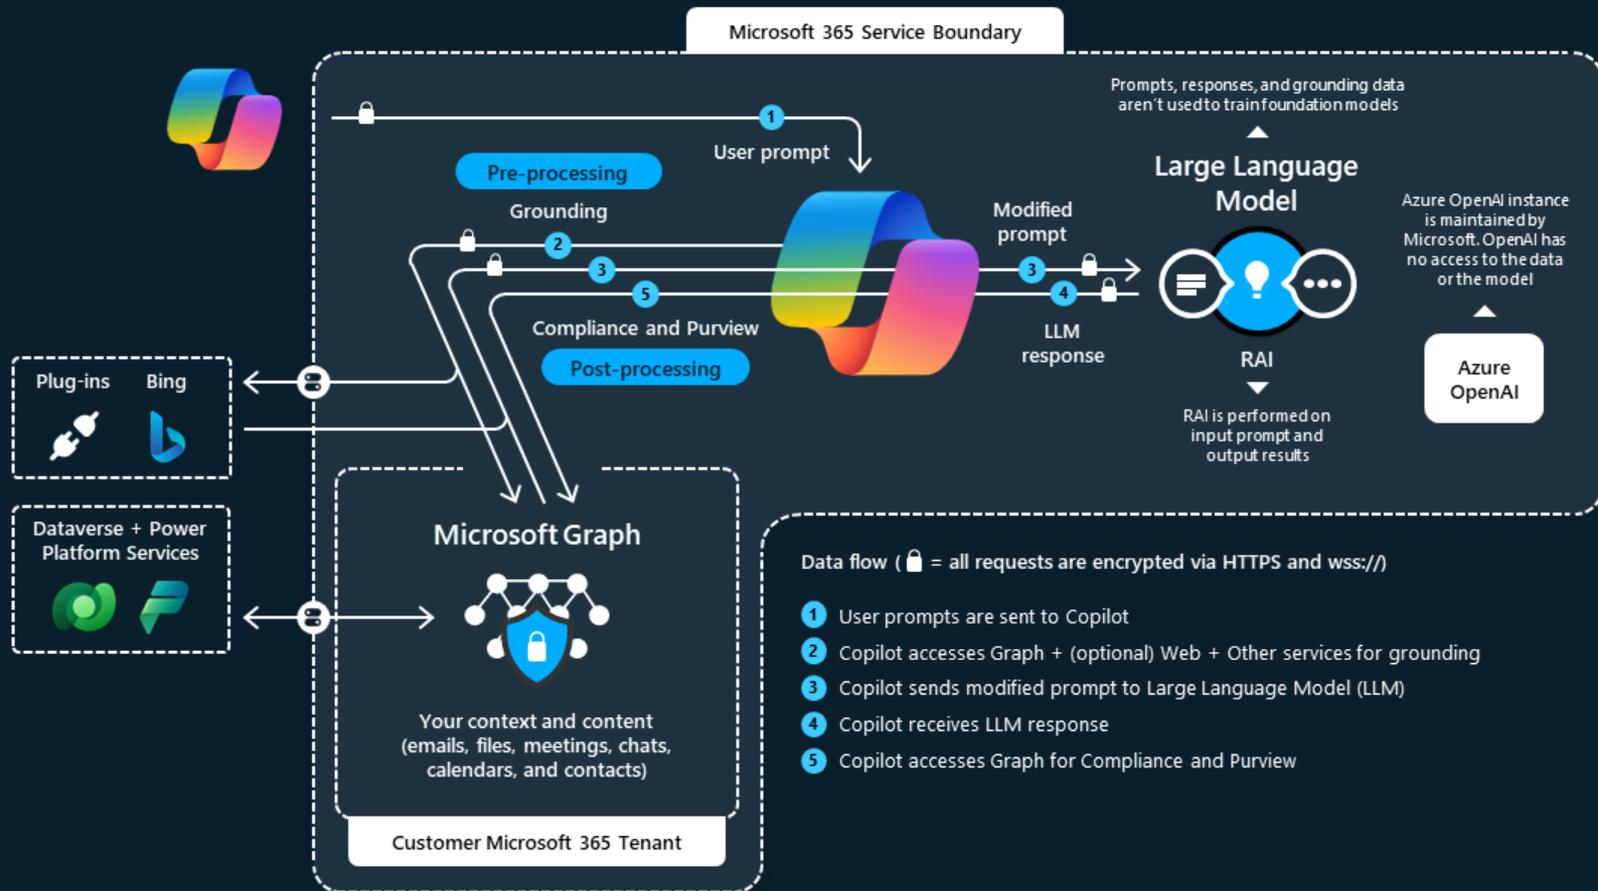
Use / to insert people, files and more

 Add people, files, and more for better results

0 / 2000

# Microsoft Copilot for Microsoft 365 architecture



# Azure OpenAI



- [Azure OpenAI Service – Large Language Models for Generative AI \(microsoft.com\)](https://microsoft.com)
- Azure OpenAI service features offers REST API to GPT-4 and GPT-35
- Both language models co-develop their APIs, guaranteeing compatibility and easy transitions
- Azure OpenAI ensures the same security standards as Microsoft Azure
- Azure OpenAI offers private networking, regional availability, and responsible AI content filtering



## Chat with our support team

How to search and book rentals?

[Ask now](#)

What is the refund policy?

[Ask now](#)

How to contact a representative?

[Ask now](#)

Type your question, eg. "How to search and book rentals?"



# RAG Retrieval-Augmented Generation



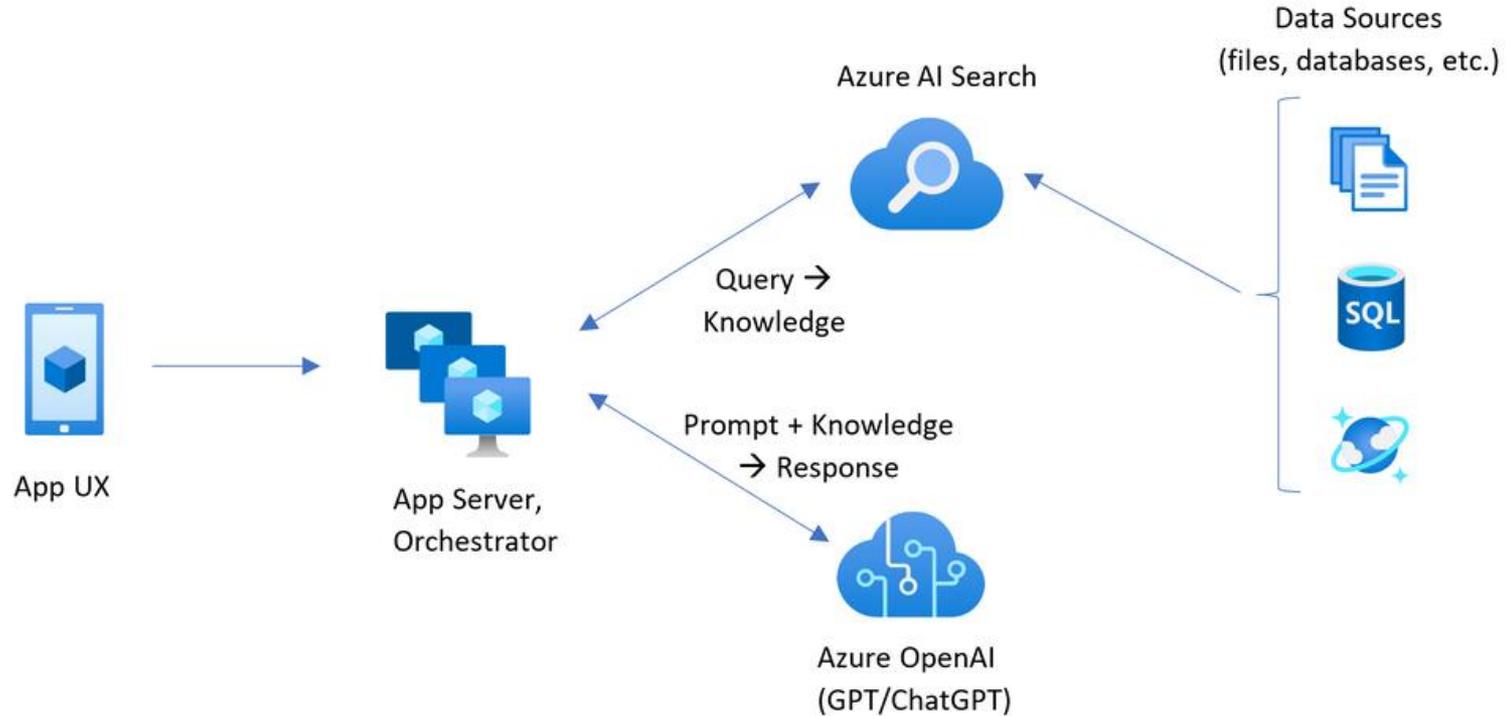
**Retriever**



**Generator (LLM)**

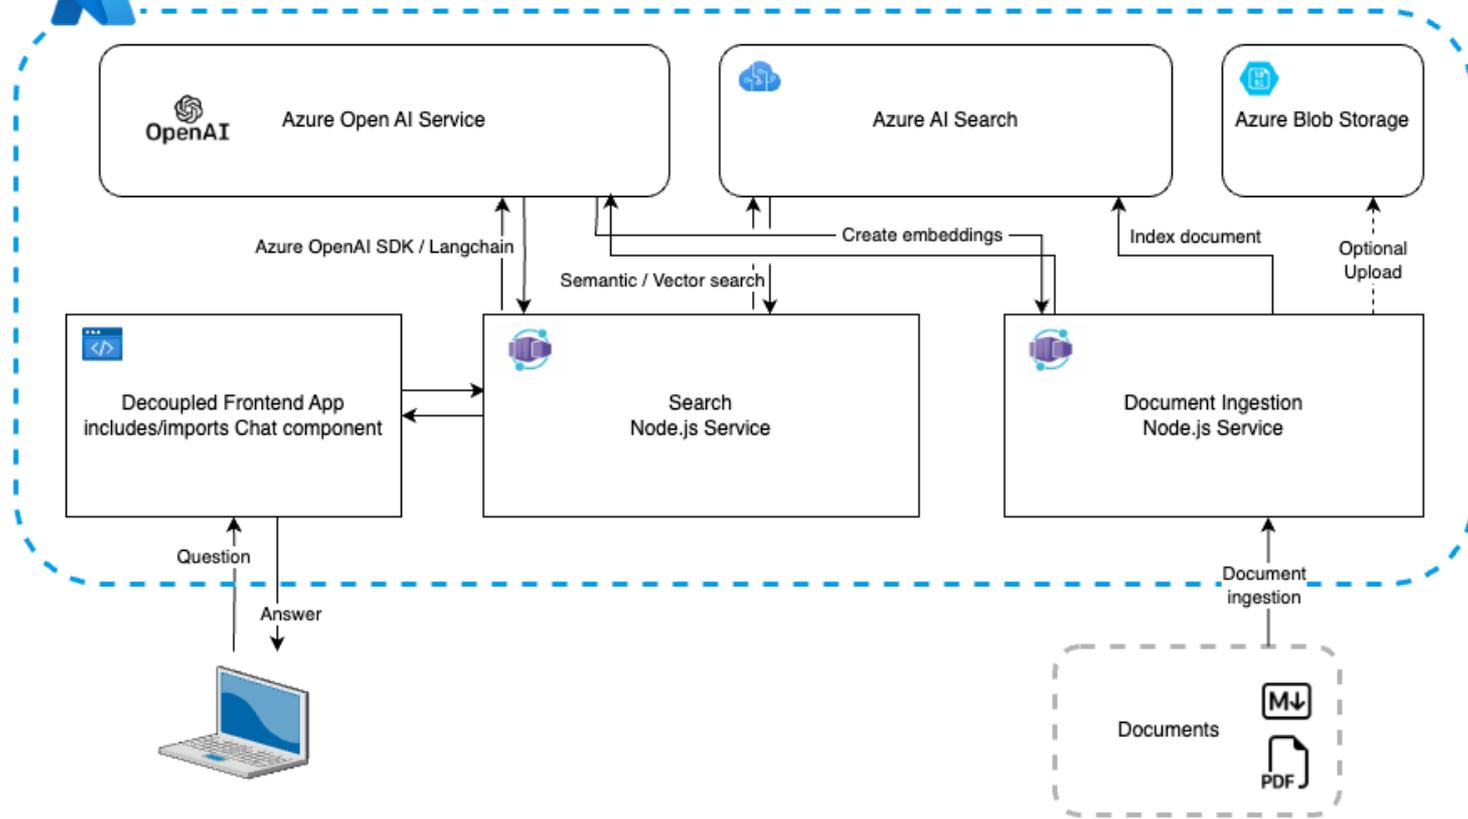


# Architectural overview





## ARCHITECTURE DIAGRAM



# Pricing

## [Azure OpenAI Service - Pricing | Microsoft Azure](#)

### Priced and scaled by tokens and context

- 1 token  $\sim$  4 chars in English
- 1 token  $\sim$   $\frac{3}{4}$  words
- 100 tokens  $\sim$  75 words

Models	Context	Input (Per 1,000 tokens)	Output (Per 1,000 tokens)
GPT-3.5-Turbo-0125	16K	\$0.0007	\$0.0021
GPT-3.5-Turbo-Instruct	4K	\$0.0021	\$0.003
GPT-4	8K	\$0.042	\$0.083
GPT-4	32K	\$0.083	\$0.165

[Calculating Azure OpenAI Service Usage Costs: A Comprehensive Guide | by Oleksandr Krakovetskyi | The DevRain Tech Blog | Jul, 2023 | Medium](#)



Search 



AI Community Conference Vancouver 2024  
# AICommunityConference

# Transitioning from Keyword to Semantic and Vector Search

**Keyword Search:** Direct matching of query terms to terms in documents or databases.

**Use Cases:** Early search engines, basic database queries.

**Advantages:** Simple, quick, easy to implement.

**Limitations:**

- Requires exact term matches
- Can yield irrelevant results if terms are out of context
- Can miss relevant documents if they don't use the exact terms





## Semantic Search

Understands the meaning and intent behind user queries, rather than just matching keywords.

**Use Cases:** modern search engines, CMS, e-commerce recommendations.

**Advantages:**

- Relevant results
- Can answer queries instead of just returning files
- User-friendly with natural language queries

**Limitations:**

- Requires more sophisticated algorithms
- Potential for ambiguity in interpretation



## Vector Search

Finds items in a dataset based on their similarity in meaning or context, rather than exact keyword matches.

**Use Cases:** advanced data retrieval systems, document clustering, and topic modelling.

**Advantages:**

- Relevant results
- Understands context, relationships between words
- Supports languages and terms with sparse data

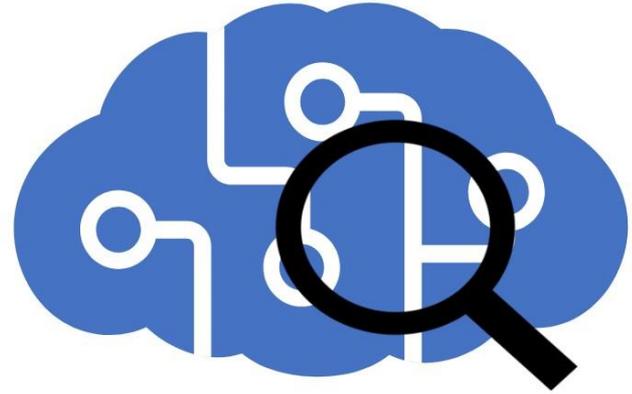
**Limitations:**

- Requires substantial computational power
- Vector models need training and refining



# Azure AI Search

- Cloud service providing a full-text search engine, offering semantic and vector search capabilities
- Intended for developers needing application/based search experiences
- Benefits from vast data sources and AI-enhanced indexing

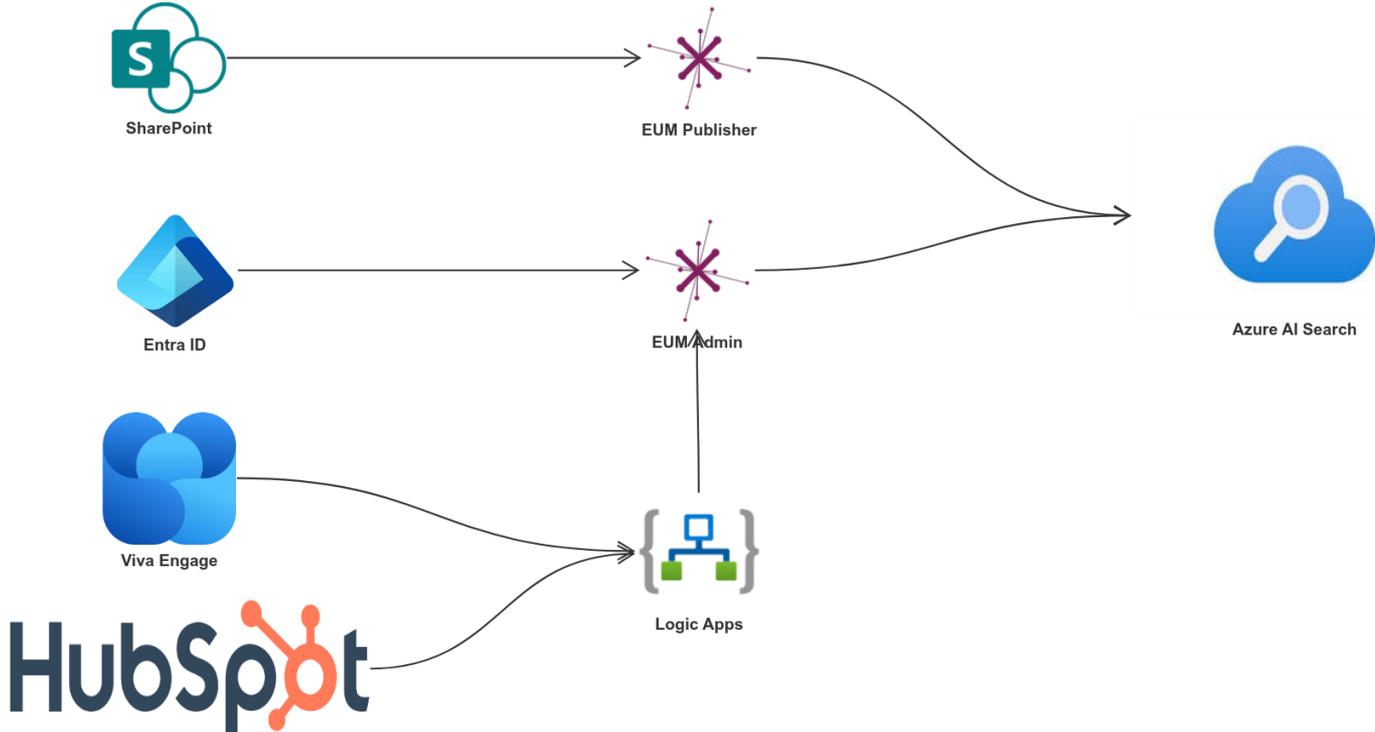


# Azure AI Search Pricing

	Free	Basic	Standard S1	Standard S2	Standard S3	Storage Optimised L1	Storage Optimised L2
Storage <sup>1</sup>	50 MB	15 GB (max 45 GB per service)	160 GB (max 1.9 TB per service)	512 GB (max 6 TB per service)	1 TB (max 12 TB per service)	2 TB (max 24 TB per service)	4 TB (max 48 TB per service)
Max indexes per service	3	15	50	200	200 or 1,000/partition in high density mode	10	10
Scale out limits	N/A	Up to 9 units per service <sup>2</sup> (max 3 partition; max 3 replicas)	Up to 36 units per service (max 12 partition; max 12 replicas)	Up to 36 units per service (max 12 partition; max 12 replicas)	Up to 36 units per service (max 12 partition; max 12 replicas) up to 3 partitions in high density <sup>3</sup> mode	Up to 36 units per service (max 12 partition; max 12 replicas)	Up to 36 units per service (max 12 partition; max 12 replicas)
Price per SU (Scale Unit)	\$0/hour	\$0.14/hour	\$0.47/hour	\$1.85/hour	\$3.69/hour	\$5.26/hour	\$10.52/hour



# EUM AI Search Architecture



# Azure AI Search Demo



AI Community Conference Vancouver 2024  
# AICommunityConference

Home > Search > Search Snippet Builder

## Search Snippet Builder

Snippet

Articles List Search Snippet

Query

Template

Search

SearchText

Filter

EUMSearchSource eq 'Publisher Pages' and ContentType eq 'EUM Article' and search.in(SiteURL,{'currentSite'})

Select

Title, RelativeURL, RollupContent, RollupImage, PublishedDate1, SiteURL

OrderBy

PublishedDate1 desc

Size

JSON Data Table Preview Fullscreen

Snippet



### [Best Practices in Secure File Sharing for Government](#)

June 27, 2024

In 2023, private data exposure incidents reached a record high, with 100 reported cases impacting a ...



### [EUM Sponsored Mercedes-Benz World Driving Experience](#)

June 26, 2024

EUM is giving away a half hour driving experience at Mercedes-Benz World in Weybridge, UK. The draw ...



### [What's Going on with Microsoft Delve?](#)

June 25, 2024

Microsoft will be retiring the Delve web app as of December 2024. Read to find out why, how this may...

### [Configuring Power Pages Entra ID B2B Authentication](#)



EUM Admin Search Explorer Search

Home Reports Users Groups Admin Email Templates

SearchText

Request Body

```
{
  "searchOptions": {
    "filter": "EUMSearchSource eq 'Publisher Pages'",
    "size": "10"
  },
  "select": [],
  "orderBy": []
}
```

JSON Data Table

```
{
  "totalCount": 899,
  "facets": {},
  "value": [
    {
      "score": 1,
      "highlights": null,
      "semanticSearch": {
        "rerankerScore": null,
        "captions": null
      },
      "document": {
        "id": "14626fa4-5c33-4a19-a3d6-5a647ca5205d-porta1dev",
        "company": null,
        "companyID": null,
        "partner": null,
        "partnerID": null,
        "status": null,
        "estimatedCloseDate": null,
        "guestCount": null,
        "installationPackage": null,
        "licenseTerm": null,
        "notes": null,
        "waitingon": null,

```

Support User Guide License valid until 2024-12-31 © 2024 Extranet User Manager - v7.1.1654.1



Home > About Us > Directory > Envision IT Staff Directory

## Envision IT Staff Directory

Enter a search term



Found 24 results in 0.58 seconds



Aaron Silver  
Developer



Alina Raza  
UI / UX Developer



Amanda Da Costa  
Executive Assistant



Anish Birju  
Developer



Artur Pinheiro  
Developer



Denesh Sohan  
Director of Engineering



Devraj Banerjee  
Customer Success Specialist



Don Harkness  
Chief Financial Officer



Eman Abdeen  
Quality Assurance and Support



Faye Yuan  
Senior Accountant



Gurinderpal Dhillon  
Senior Developer



Jeff Clement  
Senior UI/UX Developer



Joel Stickland  
Sales Development Representative



Kate Carson  
Developer



Kevin Smith  
Chief Marketing Officer



Logan Guest  
VP Sales and Customer Success



Maha Elkoshairi  
Senior Consultant



Mark Campbell  
Senior Consultant



Niki Patel  
Quality Assurance Specialist



Peter Carson  
President & CEO



## Service Areas

 Accounts Payable	 Building Permits and Inspections	 By-Law Enforcement	 Child Care
 Clerks	 Culture	 Customer Service	 Emergency Medical Services (EMS)
 Emergency Shelters	 Facilities	 Fire and Rescue Services	 Fleet
 General Government	 General Revenues	 Human Resources	 Information Technology

<https://mbncanada.ca>



Home &gt; Search

## Search

### Filters

- All
- PDF
- Webpages

parking

✕

Found 353 results in 0.33 seconds

[Parking App](#)

No coins?

[Disputing a Parking Ticket](#)Quick Link for Dispute Forms: Request for Screening Officer Review Form Request for Hearing Officer Review Form Why are **Parking** Tickets issued?[JMCC Parking Lot](#)JMCC **Parking** Lot at Malden Park Why was the **Parking** Lot Closed?[Accessible Parking Permits and Veteran Plates](#)**Parking** privileges are available to those with a valid accessible **parking** permit or a veteran plate (Ontario poppy plate or other demarcation).[Parking Officers Misconceptions](#)There are a lot of common misconceptions regarding the **Parking** Enforcement outside contractor, the Commissionaires of Ottawa - Windsor Division.[3-Day Parking Infraction and Trailers](#)What is a 3-day **parking** infraction?<https://www.citywindsor.ca>



[Home](#) > [Directory](#) > [Speakers](#)

## Speakers

Enter a search term



Found 41 results in 2.27 seconds



Amit Vasu



Brad Orluk



Brian Lalancette



Craig Lussier



Dave Healey



David Drever



Denis Molodtsov



Denzil Fernandes



Doug Davis



Doug Ferrier



<https://tm365.org>



# Key Points

- Search is a key component of a chat experience
- Think about grounding and security
- There are many ways to do this within Microsoft
  - Copilot
  - Copilot Studio
  - Custom dev and Azure
- Azure AI Search is incredibly powerful



# Summary of Key Benefits



Enhanced  
Productivity

Improved  
Decision-  
Making

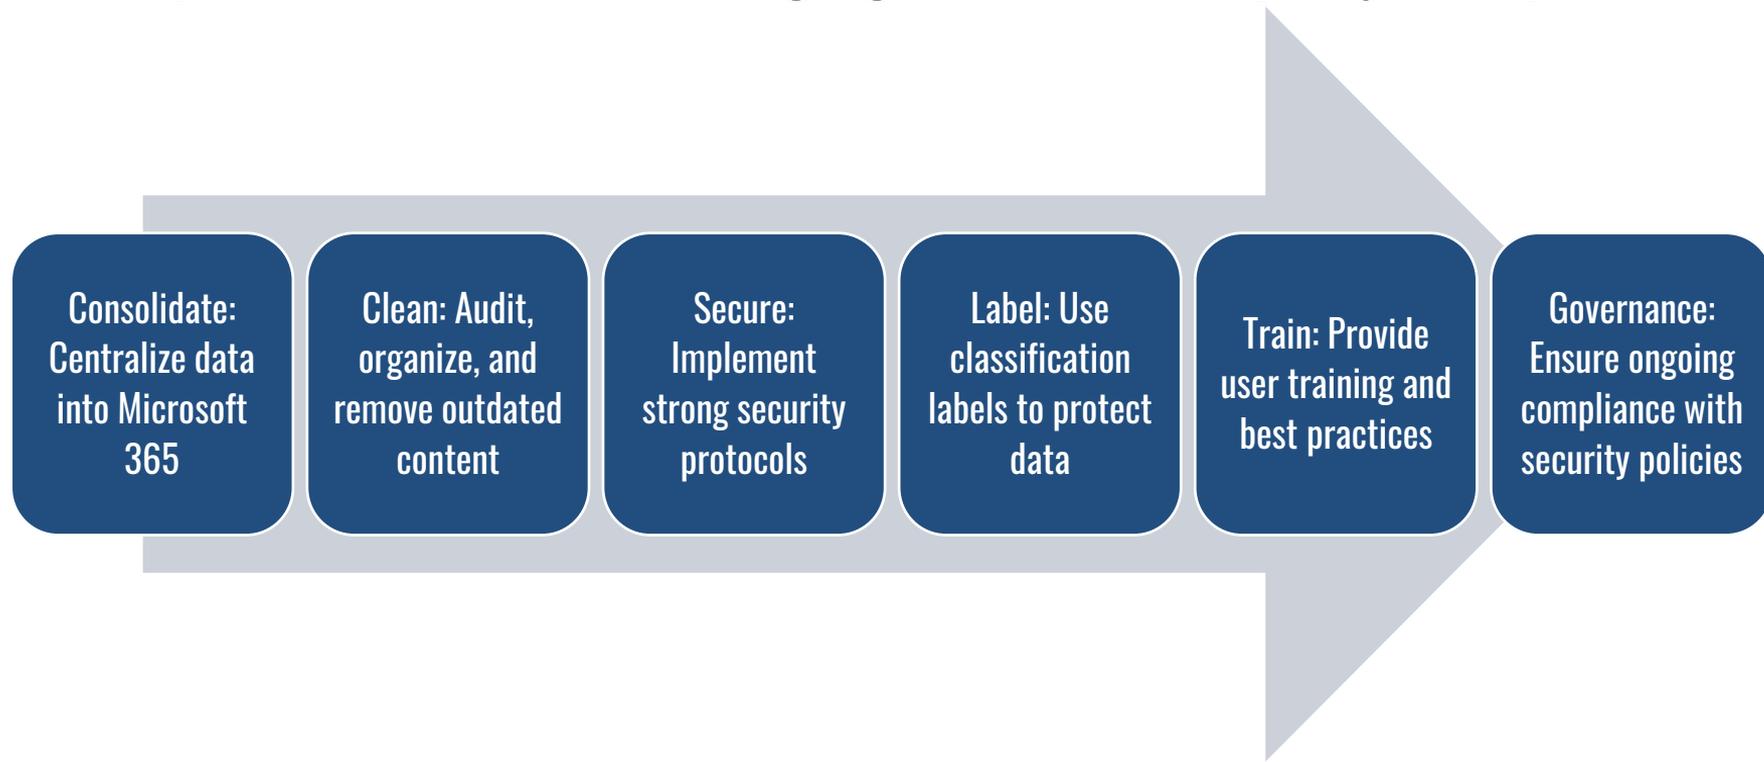
Cost  
Savings

Enhanced  
Customer  
Experience

Employee  
Satisfaction



# Copilot Readiness Engagement: Step-by-Step



# Resources

- [AI Community Conference - Vancouver 2024 | Extranet User Manager](#)
- [Azure AI Search and Open AI Webinar | Extranet User Manager](#)
- [Get started with the chat using your own data sample for JavaScript - JavaScript on Azure | Microsoft Learn](#)
- [Azure AI Search documentation | Microsoft Learn](#)
- [Simple query syntax - Azure AI Search | Microsoft Learn](#)
- [Pricing - Azure AI Search | Microsoft Azure](#)
- [Azure OpenAI Service - Pricing | Microsoft Azure](#)
- [GitHub - Azure-Samples/serverless-chat-langchainjs: Build your own serverless AI Chat with Retrieval-Augmented-Generation using LangChain.js, TypeScript and Azure](#)
- [Building RAG apps with LangchainJS - YouTube](#)



# Speaker Evaluations

Please complete an evaluation for each session you are attending. This will help us in providing our speakers with valuable feedback. You will also increase your chances of winning a raffle prize



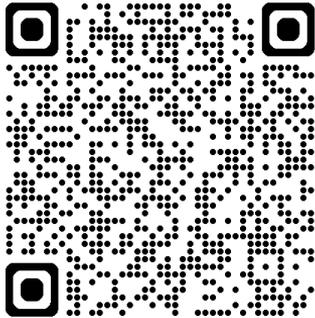
# Upcoming Event



**Webinar:** Azure AI Search and OpenAI

**When:** October 30<sup>th</sup>, 11-11:45AM EDT

**What:** Learn how you can make Microsoft365 search results more relevant and accurate by implementing AI and supercharge your organization's search results.



Register now!



# Thank you for attending!



AI Community Conference Vancouver 2024  
# AICommunityConference