KMWorld
Content, Document and Knowledge Management

KMWorld 2024

SOUND OFF: CONFERENCE HIGHLIGHTS AND 2025 FORECAST

TABLE OF CONTENTS

Photo Credit by Steve Nathans-Kelly, and Gavin Willison

Introduction to the KMWorld Conference Collected Presentations

The most recent KMWorld conference was held in Washington DC, November 18-21, 2024, with the theme "KM & Enterprise Intelligence: Human or Artificial?" The four co-located conferences also had interesting themes.

- Taxonomy Boot Camp: "Agents of Change"
- Enterprise Search & Discovery: "Bringing Our Search and Discovery Superpowers to Work"
- Text Analytics Forum: "New Synergies and New Solutions"
- Enterprise AI World: "Crossing the Chasm in the New Era of AI"

The KMWorld conference team decided to try something new this year. We asked speakers to submit an article or paper that summarized their talks or put forward thought leadership ideas. Those that responded are included in this special publication. They highlight some of the important trends and topics that emerged at the event.

Talks often revolved around AI technologies, tempered with insights regarding human involvement and use of knowledge. They touched on knowledge sharing to aid customer service, information management, search capabilities, and strategic planning. The practical side of integrating GenAI into KM projects was a constant throughout the conference.

The importance of AI technologies for knowledge management is a trend that is likely to continue into 2025. It is also a major focus of our new conference, the Knowledge Management and AI Summit (https://www.kmworld.com/KMAISummit/2025), to be held March 17-19 in Scottsdale, AZ. Heralding a new era of KM, the summit intends to provide numerous opportunities for informal collaboration along with the more traditional conference presentations.

We hope you enjoy the contents of this special publication from the KMWorld team.

Marydee Ojala
Editor-in-Chief, KMWorld

Revolutionizing Knowledge Management with Knowledge Graphs and Neuro-Symbolic AI

By Jans Aasman, CEO, Franz Inc.

In today's data-rich world, organizations are increasingly challenged by the need to manage vast amounts of information efficiently. Traditional knowledge management systems often struggle to keep pace with the complexity and volume of enterprise data. Enter Knowledge Graphs and Neuro-Symbolic AI—two cutting-edge technologies that are rapidly transforming the landscape of knowledge management. By integrating these tools, organizations can significantly enhance their ability to structure, retrieve, and reason about data, leading to more informed decision-making and improved business outcomes.

THE GROWING ROLE OF KNOWLEDGE GRAPHS

Knowledge Graphs have emerged as one of the most promising solutions for managing and integrating data across diverse systems. At their core, Knowledge Graphs are a way to represent information in a semantic, interconnected format. They allow organizations to structure data in a way that reflects real-world relationships, making it easier to link and query information across multiple domains.

Gartner has recognized Knowledge Graphs as a critical trend in data integration for 2024, highlighting their role in addressing the complexity of enterprise data. By providing a framework that connects disparate datasets, Knowledge Graphs enable businesses to uncover hidden insights, enhance data discovery, and improve the accuracy of their analytics. For industries like healthcare, finance, and technology, this ability to semantically integrate data is proving invaluable.

One of the key advantages of Knowledge Graphs is their flexibility. They can represent highly complex and dynamic relationships, which are essential for capturing the intricacies of today's enterprise environments. For example, in the healthcare sector, a Knowledge Graph could link patient data with medical research, treatment histories, and drug interactions, allowing clinicians to make more informed treatment decisions.

INTRODUCING NEURO-SYMBOLIC AI: THE NEXT FRONTIER

While Knowledge Graphs provide an essential framework for organizing and linking data, Neuro-Symbolic AI takes their capabilities to the next level. Neuro-Symbolic AI is a hybrid approach that combines the learning power of neural networks with the reasoning capabilities of symbolic AI. This merging of technologies enables more sophisticated knowledge representation, reasoning, and adaptability.

Gartner recently added Neuro-Symbolic AI to its AI hype cycle, recognizing its potential to bridge the gap between data-driven machine learning models and rule-based reasoning systems. In traditional AI models, neural networks excel at pattern recognition and learning from data, but they often struggle with explainability and complex reasoning tasks. Symbolic AI, on the other hand, is better at logical reasoning and inference but lacks the ability to learn and adapt from new data.

By integrating Neuro-Symbolic AI with Knowledge Graphs, organizations can create systems that not only understand the relationships between data points but can also reason about those relationships. This is a significant leap forward in AI technology, as it enables systems to make complex decisions based on both learned patterns and logical rules. In practical terms, this means that organizations can build more robust and adaptable knowledge management systems capable of handling both structured and unstructured data.

REAL-WORLD APPLICATIONS AND IMPACT

The combination of Knowledge Graphs and Neuro-Symbolic AI is already addressing complex problems across various industries. In healthcare, these technologies are enhancing diagnostic accuracy and enabling more personalized treatment plans. By integrating patient data with medical research, treatment histories, and genetic information, Knowledge Graphs provide clinicians with a comprehensive view of a patient's condition. When combined with Neuro-Symbolic AI, these systems can reason about potential treatment outcomes, helping doctors select the most effective interventions.

To make this more concrete, consider the work we do at Franz. We process large volumes of patient EMRs and clinical notes, organizing them into an entity- and event-based patient graph. This allows us to view a patient as both a holistic graph and a sequence of events. We feed these events into a rule-based, Prolog-driven system to make predictions about likely outcomes. These event sequences also train machine learning models, generating data-driven predictions. Finally, we can present patient events to a GenAI model to generate predictions or explain a machine learning model's output. Integrating logic-based, machine learning, and GenAI methods lies at the core of Neuro-Symbolic AI.

Since these three prediction approaches might occasionally differ, we also create a meta-reasoning system to evaluate and decide on the most plausible prediction—another essential aspect of Neuro-Symbolic AI.

In the financial sector, Neuro-Symbolic AI and Knowledge Graphs are being deployed to enhance fraud detection and risk management. By integrating data from diverse sources—such as transaction histories, social media, and legal records—these systems can detect patterns indicative of suspicious behavior and assess the likelihood of fraudulent activity. This approach not only improves the accuracy of fraud detection but also enables financial institutions to respond to potential threats more swiftly.

Similar to how we present patients in a Knowledge Graph, we also represent each individual in the graph as both a node in a network and as a series of events. In fraud detection, we utilize four distinct types of AI to determine whether an activity is fraudulent. First, graph analytics helps uncover interaction patterns that might indicate fraud. Second, specific patterns are hard-coded using logic (often in Prolog). Third, machine learning predicts likely next events for individuals and classifies them into relevant categories. Lastly, generative AI examines sequences of events as narratives, suggesting the most probable next step in each story. This approach exemplifies the power and versatility of Neuro-Symbolic AI.

CHALLENGES AND OPPORTUNITIES

Despite their significant potential, implementing Knowledge Graphs and Neuro-Symbolic AI is not without challenges. Data integration remains a critical issue, as organizations often have to deal with diverse, siloed datasets that lack consistent formatting or structure. Scaling these systems to handle large volumes of data while maintaining performance and accuracy is another obstacle. Furthermore, the interpretability of AI models remains a concern, particularly as systems become more complex and data-driven.

However, innovative solutions are available to address these challenges. Tools and methodologies are constantly improving data integration and scalability, ensuring that Knowledge Graphs can handle the growing demands of modern enterprises. Advances in explainable AI are also making it easier to interpret the reasoning behind Neuro-Symbolic AI systems, helping organizations to trust and act on the insights generated by these technologies.

THE FUTURE OF KNOWLEDGE MANAGEMENT

As organizations continue to grapple with the complexity of their data ecosystems, the integration of Knowledge Graphs and Neuro-Symbolic AI offers a powerful solution for enhancing knowledge management. By providing a semantic framework for data integration and enabling advanced reasoning capabilities, these technologies are transforming the way organizations manage and leverage their knowledge assets. Moving forward, businesses that adopt these tools will be better positioned to optimize their operations, improve decision-making, and gain a competitive edge in their respective industries.

Jans Aasman (jans.aasman@franz.com) is the CEO of Franz, Inc.

The Challenges of Building Generative AI Applications in the Enterprise: From POC to Production

Zeta Alpha

By Fernando Rejon Barrera, CTO @ Zeta Alpha

INTRODUCTION

The rapid advancements in Generative AI have opened up new horizons for enterprises, enabling them to enhance operational productivity and augment the value of their product offerings. Modern Large Language Models (LLMs) possess broad internal knowledge, good reasoning abilities, and can be easily steered via human instructions. Furthermore, LLM vendors have significantly lowered the barriers to entry, empowering internal teams to quickly iterate on innovative AI-enabled projects and ideas. However, transitioning from a Proof of Concept (POC) to full-scale production presents numerous challenges, which this article explores drawing from my experience while working with enterprise clients. By addressing these challenges, enterprises can better navigate the path from ideation to implementation, ensuring that their Generative AI initiatives are both successful and sustainable.

EVALUATION

As organizations begin to explore the potential of Generative AI, establishing clear metrics for success becomes paramount.

ESTABLISHING EVALUATION CRITERIA

Defining success criteria often requires collaboration with domain experts to outline meaningful and realistic interactions that the solution should support. For example, a customer service application might be evaluated on accuracy, appropriate information use, and minimal human intervention. Defining these goals ensures the solution aligns with business needs.

CREATING DATA FOR EVALUATION

Expert-created benchmarks may lack sufficient coverage for complex use cases. Synthetic data generation can help expand the evaluation dataset, ensuring a broader range of scenarios is tested. For instance, an AI-driven HR assistant could benefit from synthetic examples that represent diverse employee queries, making the evaluation more comprehensive and effective. These can be generated by LLMs with the guidance of domain experts to ensure the generated cases still align with the solution's intended use cases and real-world context

EVALUATION VOLUME

Continuous iterations require extensive manual evaluations, which can quickly overwhelm domain experts due to the many tweakable aspects of the solution (e.g., retrieval strategies, prompt variations). LLM-as-a-judge techniques, like RAGElo, streamline this process by automating comparisons and rankings across variations, allowing experts to maintain oversight without excessive workload.

DATA

Once evaluation criteria are set, the next step is to ensure that the right data is available to ground the Generative AI application.

FEEDING DATA TO THE GENERATIVE AI APPLICATION

Integrating data effectively is crucial to success, though it's often more challenging than LLM prompting. Selecting the highest-quality sources can streamline efforts and prevent teams from getting bogged down with less essential data early on. It's also vital to determine how the Generative AI solution will access the data. For instance, we faced a challenge when designing a product support chatbot: understanding user intent was made feasible by using a product ID resolver at retrieval time. This increased the accuracy of the data fed to the LLM without a costly, complex indexing process.

DATA PROCESSING QUALITY

Enterprises often manage large volumes of data across multiple formats and systems, making consistent data processing a challenge. Building a flexible data ingestion pipeline with tailored extraction and transformation routines for different formats is crucial for data relevance. Progress can be made without striving for perfect data cleanliness, as LLMs can often manage unstructured text well enough for the initial POC phase. As teams iterate the solution, the evaluation loop will point to places where data processing needs to be improved.

SCALING AND AUTOMATION

As projects transition to production, the need for a dynamic data pipeline grows. Handling the sheer volume and diversity of data sources can quickly outstrip the capabilities of open-source tools. Automated pipelines are essential to keep data ingestion and transformation efficient and up-to-date, minimizing manual interventions as data changes. Creating workflows that can scale seamlessly allows the Generative AI application to rely on current, high quality data continuously, which is vital to long-term reliability.

DATA ACCESS CONTROL

Access control is critical for the enterprise. During POCs, teams often use simplified data samples without proper access restrictions

Register Now to Save $200!

Use Code **KMWU** for Early Bird pricing.

(Available Through February 14)

KM & AI SUMMIT

MARCH 17–19, 2025

JW Marriott Camelback Inn Resort & Spa

Scottsdale, AZ

EMPOWERING AI-DRIVEN KNOWLEDGE MANAGEMENT

Join us to:

- Learn from Industry Leaders
- Future-Proof Your Strategy
- Explore Real-World Case Studies
- Network with Innovators

Unlock the future of innovation and knowledge management at the KM & AI Summit in Scottsdale, Arizona. This premiere event, brought to you by the organizers of KMWorld, brings together thought leaders, technologists, KM practitioners, and industry experts to explore cutting-edge strategies for managing organizational knowledge and leveraging the power of AI to drive transformation.

Premier Sponsor: eGain

Diamond Sponsors: Bloomfire, coveo™, NICE

Platinum Sponsors: RWS Tridion, RightAnswers, VERINT

Brought to You By the Organizers of: KMWorld

Media Sponsors: AIPressRoom, database TRENDS AND APPLICATIONS, ENTERPRISE AI World, KMI, KMWorld, UNITE.AI

Organized & Produced By: Information Today

kmworld.com/kmaisummit

CONNECT WITH US! #KMAISummit

nor user authentication. However, integrating secure access controls during production is not only a best practice but often a requirement to meet company policies. Ensuring data privacy means building access controls into the Generative AI solution's architecture from the start. Without these protections, teams can face significant delays or have to overhaul systems later. Unfortunately, most open-source solutions lack support for the level of security needed in enterprise environments, so tailored approaches are often necessary to ensure compliance and secure data handling.

EXPECTATION MANAGEMENT AND MONITORING

With a robust data strategy in place, it's essential to manage stakeholder expectations and implement effective monitoring mechanisms.

SETTING REALISTIC EXPECTATIONS

To succeed, Generative AI initiatives must align stakeholder expectations with current capabilities and limitations of the technology. For instance, if developing a customer support assistant, stakeholders should know that while the AI can handle many routine inquiries, complex cases may still require human intervention. This understanding fosters realistic expectations and helps prevent future frustrations.

INVOLVING STAKEHOLDERS IN PILOT TESTING

Allowing stakeholders to participate in pilot testing builds confidence in the solution and provides valuable feedback. Their input can reveal specific needs, help to uncover performance gaps, and inform further iterations. By including stakeholders in pilot phases, teams can collect diverse perspectives and ensure the Generative AI solution aligns with user needs and is equipped for real-world application.

ESTABLISHING PRODUCTION MONITORING

Once in production, robust monitoring is essential to track the performance of the Generative AI application, gather feedback, and ensure it continues to meet expectations. Regular monitoring enables teams to identify issues early and track key metrics—such as accuracy, response time, user satisfaction, and cost. Without continuous monitoring, the Generative AI solution risks drifting from its intended purpose, potentially leading to costly rework or reduced user confidence.

CONTINUOUS IMPROVEMENT THROUGH MONITORING DATA

Data collected from monitoring and feedback loops is invaluable for refining the Generative AI application over time. Regular reviews of this data help prioritize updates, refine workflows, and align the Generative AI's performance with evolving business needs. As the solution matures, the monitoring framework itself may need to adapt, incorporating new metrics or user insights to support more advanced use cases and maximize long-term reliability.

CONCLUSION

As Generative AI continues to evolve and find its place within enterprise environments, organizations must navigate a complex landscape filled with both opportunities and challenges. By strategically prioritizing evaluation, data management, expectation alignment, and monitoring, enterprises can unlock the full potential of this transformative technology. Embracing these best practices will not only enhance the reliability and effectiveness of Generative AI applications but also empower organizations to innovate and stay competitive in an increasingly AI-centric landscape.

Fernando Rejon Barrera (rejonbarrera@zeta-alpha.com) is CTO, Zeta Alpha.

Optimizing Information Management: How to Harness the Power of Generative AI

By Laura Carson, CMO, M-Files

Since the debut of ChatGPT, both businesses and individuals have been eager to tap into the capabilities of generative artificial intelligence (GenAI). In just 10 months, from 2023 to 2024, the adoption of GenAI within companies doubled, with 65% of business executives reporting their organizations use of GenAI tools regularly. Expectations for GenAI's impact remain high, and this number is likely to continue growing.

While the potential of GenAI is vast, successfully integrating it into an organization requires a solid foundation, effective governance, and organizational readiness. As with any emerging technology, it's important to balance enthusiasm with a realistic approach.

THE BENEFITS AND CHALLENGES OF GENAI IN THE ENTERPRISE

GenAI is already making a measurable difference across various business processes. Notable examples include: capturing and processing information from meetings and calls in the form of meeting transcripts and summaries; generating first drafts or providing a brainstorming spark for content creators; or helping users digest and extract insights from lengthy reports or research papers. But despite its promise, GenAI's effectiveness in managing enterprise content faces several challenges. One major obstacle is information chaos—businesses often struggle with data silos, duplicate information, version control issues, and poor data classification. This disorganized information landscape makes it difficult for GenAI to produce accurate and reliable insights.

To overcome these challenges and unlock the full potential of GenAI in content management, companies must focus on three key areas: connectivity, confidentiality, and curation.

CONNECTIVITY

For GenAI to provide accurate and comprehensive insights, it must have access to all relevant data within the organization. This requires breaking down data silos and ensuring smooth integration across various information systems. Organizations should develop an information management strategy that includes defining data flows, setting governance rules, and creating a unified information layer that GenAI can understand.

CONFIDENTIALITY

GenAI processes vast amounts of data, making robust security measures a top priority. AI tools must take into account access controls to ensure that sensitive data remains protected. Businesses need to establish clear policies for managing data access and ensure that AI-generated outputs do not unintentionally reveal confidential information.

CURATION

For GenAI to deliver high-quality outputs, it needs access to the most relevant and up-to-date information. Content curation goes beyond organizing data; it involves ensuring that important information is properly documented and retains its business context. Regular content audits, metadata tagging systems, and clear content management guidelines helps optimize how information is classified and maintained, allowing both humans and AI systems to find and use it more effectively.

SUCCESSFULLY IMPLEMENTING GENAI: ORGANIZATIONAL READINESS AND GOVERNANCE

Effective GenAI implementation requires more than just technological tools. Businesses need to evaluate their organizational readiness and establish robust governance structures. The following framework, based on the capability maturity model (CMM), outlines the steps to guide this process:

STEP 1: DEVELOP A COMPREHENSIVE STRATEGY FOR INFORMATION MANAGEMENT AND AI INTEGRATION

This strategy should align with the organization's broader goals, outlining how GenAI will contribute to achieving these objectives. Leaders must recognize the strategic importance of information management and communicate a clear vision for the integration of AI. Shared responsibility is critical: stakeholders from various departments should collaborate to ensure the strategy meets diverse needs and fosters ownership across the company.

STEP 2: IMPLEMENT EFFECTIVE DATA GOVERNANCE

Organizations need clear data governance policies that ensure consistency and compliance. Governance should address data retention, organization, classification, and access control. It's also crucial to define which types of content are suitable for AI processing, considering both business value and legal/ethical implications.

STEP 3: ASSESS IT INFRASTRUCTURE'S READINESS

Businesses must ensure that their IT systems and infrastructure can support GenAI implementation. This includes ensuring that data systems are integrated and capable of supporting connectivity, confidentiality, and curation. When evaluating AI vendors, factors like trustworthiness, pricing models, and system compatibility should be carefully considered.

STEP 4: IDENTIFY OPPORTUNITIES FOR EXPANDING AUTOMATION IN CONTENT PROCESSES

Many businesses already have some degree of content automation in place. The goal should be to evolve from basic automation focused on consistency and compliance to more sophisticated solutions that support decision-making and offer real-time insights. GenAI should be viewed not just as a tool to speed up existing workflows, but as an opportunity to reimagine how work gets done.

STEP 5: EVALUATE END-USER NEEDS AND CHANGE MANAGEMENT REQUIREMENTS

Understanding how GenAI will impact different user groups is critical to successful adoption. Businesses should assess employees' readiness to embrace AI tools and invest in training programs to

ensure smooth transitions. Building enthusiasm and addressing concerns through clear communication strategies will be key to ensuring widespread acceptance.

REALIZING THE POTENTIAL OF GENAI IN INFORMATION MANAGEMENT

GenAI is a transformative tool for reshaping how organizations manage, analyze, and use their information. However, unlocking its full potential requires a comprehensive approach that goes beyond just adopting new technology. Companies must address underlying challenges in information management—particularly connectivity, confidentiality, and curation.

By following this strategic framework, organizations can lay the foundation for successful GenAI adoption, positioning themselves to thrive in an increasingly data-driven, AI-powered business landscape. While the journey to effective GenAI implementation may be complex, the potential benefits—including improved productivity, better decision-making, and increased innovation—make it a worthwhile pursuit for forward-thinking leaders across industries.

As we continue into the era of AI-enhanced information management, businesses that successfully navigate this transition will be well-positioned to capitalize on new opportunities and maintain a competitive edge in an evolving digital landscape.

Laura Carson (laura.carson@m-files.com) is CMO, M-Files.

1 https://www.mckinsey.com/featured-insights/mckinsey-global-surveys

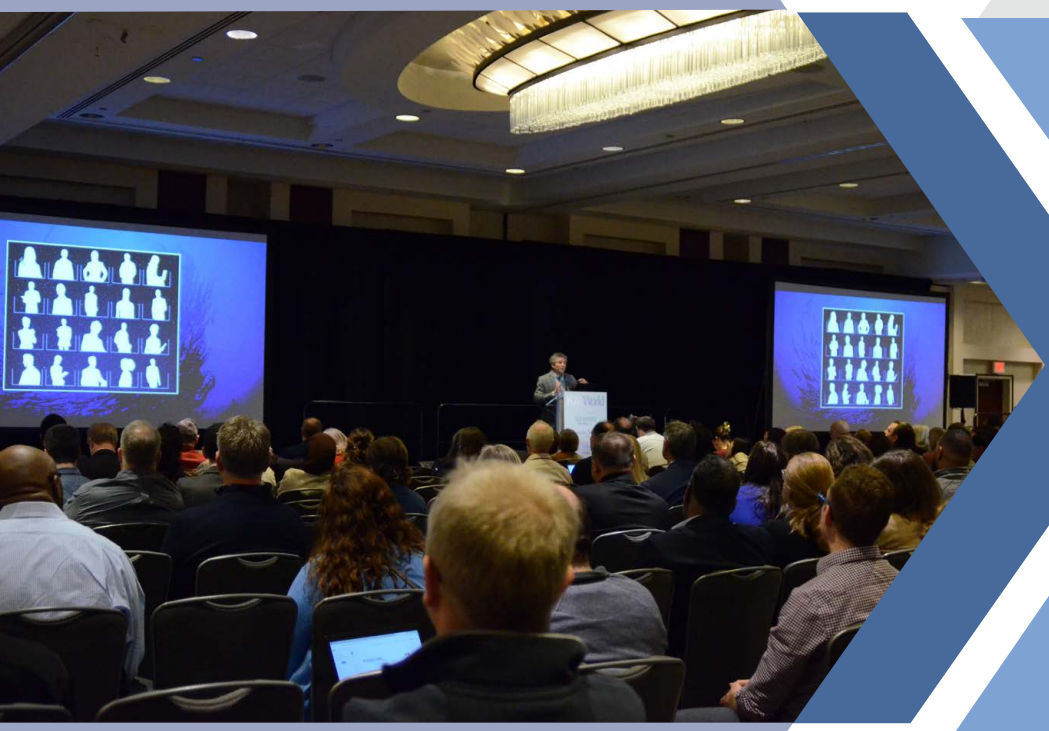

A Knowledge-Powered Customer Service Transformation

By Amy Durst, Assistant VP, Internal Support, Rogue Credit Union

Rogue Credit Union has navigated a significant evolution in knowledge management, driven by growth, change, and a commitment to enhancing both member and employee experiences. Founded in 1956 as a small cooperative of 10 teachers, Rogue Credit Union has expanded to 23 branches and serves over 200,000 members across Oregon, Idaho, and California, managing assets exceeding $3.5 billion. With a mission to provide exceptional experiences and foster mutually beneficial relationships, Rogue faced a transformative period during and after the COVID-19 pandemic.

THE PERFECT STORM OF GROWTH AND CHANGE

The challenges that led to a knowledge management overhaul stemmed from unprecedented net asset growth, an influx of new employees, a core system change, and a major merger. In just four years, Rogue's net assets grew from $1.5 billion to $3.5 billion, and its workforce almost doubled—a growth that reflected new positions rather than replacements. Additionally, the core system change and expansion into Idaho through a merger further complicated internal processes. These changes underscored the need for a robust internal information system, as employees across departments required quick and reliable access to essential knowledge. To address these needs, Rogue established a new Internal Support department and sought an innovative knowledge solution.

JOURNEY TO EGAIN

Rogue initially attempted to optimize its previous knowledge base, which was originally set up for IT workflows and lacked accessibility for all employees. The system also had limited analytics capabilities, tracking only basic usage counts. After engaging 17 consultant, we concluded the legacy system could not be improved further. We entered the RFP process and began evaluating 15 knowledge-base vendors. The selection process included hosting three vendors onsite, with eGain standing out due to their team's effort to understand Rogue's culture and needs. Rogue concluded only eGain met the stringent 74 business requirements.

The implementation phase began in January 2023, involving a collaborative project team of eight individuals from Rogue and eGain. By the live launch in June 2023, the team had transferred over 3,000 articles, including multimedia content, and established customized groups and permissions tailored to various job descriptions. The launch timeline was completed two weeks ahead of schedule, showcasing the team's effective collaboration.

FEATURES AND BENEFITS OF EGAIN

The current system boasts powerful features such as:

- Advanced search capabilities and personalized content by role.
- A global 'Find All & Replace' function.
- Change audits and a feedback system to close the loop.
- Usage analytics, dashboards, and guided workflows.
- Scheduled publishing and expiration options.
- Custom dictionary and synonym list tailored for Rogue-specific terms.

These features ensure that employees can quickly find relevant information, enhancing efficiency and reducing wait and handle times. eGain's user-friendly, personalized portals have also increased knowledge base usage by 40% over the previous system. Currently, 24 active portals exist to support the unique needs of different employee positions. The layout includes trending, new, and recently updated articles, ensuring employees stay informed. Additionally, the knowledge base allows for article suggestions and ratings, with 256 suggestions fully resolved and an average article rating of 4.56.

SEARCH EFFECTIVENESS AND PERFORMANCE IMPROVEMENTS

eGain's intelligent search, powered by natural language processing (NLP) and a custom dictionary of Rogue-specific terms, provides contextual recommendations and self-learning capabilities. This has significantly improved search outcomes and reduced the average handle and wait times for calls. AI-powered search enables users to find answers even if they do not know the exact keywords or spelling, making it easier for new employees to navigate the system. Metrics highlight marked improvements from December 2022 (pre-eGain) through October 2024, showing progressive benefits as features like Instant Answers were piloted and refined. For example, the average wait time to reach an agent dropped by 25 seconds within six months post-launch.

ENHANCING EFFICIENCY WITH ASSISTGPT

AssistGPT has been a game-changer for content authors within Rogue. It acts as a collaborative partner, ensuring content maintains consistency in tone and verbiage. Authors report increased efficiency and the ability to focus on more complex projects while still ensuring comprehensive and understandable content. One author noted that AssistGPT helps in not overlooking small details and allows them to produce higher-quality content at a faster pace. Internal support specialists have also praised the Instant Answer feature for organizing relevant information efficiently and providing quick access to supplementary details through hyperlinks. This feature has enabled employees to self-serve answers to common queries, further reducing the number of calls and support requests.

ACHIEVEMENTS AND AWARDS

The implementation of eGain has contributed to substantial achievements. Notably, the internal Net Promoter Score (NPS) increased by 20%. Additionally, Rogue's knowledge management initiative led to the formation of a dedicated team, including roles such as Knowledge & Content Manager and Content Writers. This dedication to knowledge management has fostered a culture of continuous improvement and collaboration. In 2023, the department earned significant accolades, including Internal Department of the Year, Internal Employee of the Year, Rookie of the Year, and the prestigious President's Award. These awards not only recognized effective deployment

but also highlighted a cultural shift within the organization toward prioritizing knowledge sharing and empowerment.

LOOKING FORWARD

Rogue Credit Union continues to advance its knowledge management strategy. Future plans include full integration with Talkdesk, an external knowledge base, and expanded language capabilities, such as Spanish translations. These upcoming features are designed to enhance both employee efficiency and member satisfaction. The external knowledge base will empower members to find reliable information independently, improving their overall experience. Additionally, introducing Spanish translations underscores Rogue's commitment to inclusivity and better service for its diverse member base.

Rogue's journey with eGain exemplifies a proactive approach to meeting the demands of a rapidly evolving operational landscape. The credit union's commitment to continuous improvement and leveraging innovative technology positions it to better support its members and employees, staying true to its mission and purpose.

Amy Durst (adurst@roguecu.org) is Assistant VP, Internal Support, Rogue Credit Union.

The Importance of Knowledge Management in the Vastly Expanding and Intricate Universe of Generative Artificial Intelligence

By Doron Gower, Chief Solution Architect, KMS Lighthouse

In our highly advanced and rapidly evolving world, where technology is increasingly taking a front seat, GenAI is truly shifting the paradigm. The advent of models such as GPT and it's subsequent versions are fundamentally transforming how we generate, interpret and utilise all forms of content, be it text-based, visual imagery or even complex programming codes. Despite all this cutting-edge development, one truth remains consistent—the enormous success of GenAI systems isn't possible without an impactful and well-structured knowledge base, and that's where effective knowledge management comes in. Simply put, in the absence of efficient KM, AI models will end up producing content that is far from valuable, or even completely off-target. Let's explore why KM is regarded as the integral backbone that supports and powers the success of GenAI.

One of the key principles of AI is that its competence is directly proportional to the quality of the data it is trained on. Robust knowledge management practices ensure that the content utilized by AI isn't a chaotic jumble but is instead structured, consistent and high-quality, thereby setting the stage for superior decisions and insights. By effectively consolidating, cleaning, and organizing the vast ocean of information that is available, KM paves the path for AI models to create and provide solid, reliable and valuable outputs.

Take healthcare as a case in point. Within this critical field, effective KM guarantees that AI models have access to the most recent and properly vetted processes, methodologies and guides. The end product of this meticulous process? Accurate, useful and clinically sound content that professionals can trust and rely on. This is a key focus of KMS Lighthouse ensuring leading organizations, operating under strict rules and regulations are able to provide accurate and consistent knowledge.

In an era where AI is disrupting traditional norms and swiftly rewriting the rules of the game, it is crucial for individuals and teams to keep pace. This is another realm where KM masterfully steps in, supplying resources such as comprehensive training materials, in-depth guides and insightful case studies to enable teams to swiftly adapt to and master these transformative tools. GenAI doesn't merely aid in the learning process; it turbocharges the speed of learning, helping users get up-to-speed at a rapid pace. Moreover, by following a copilot approach, KMS Lighthouse is empowering users to benefit from GenAI as a tool to improve their efficiency rather than a technology as a replacement for their skills and talent.

Doron Gower (dorong@kmslh.com) is Chief Solution Architect, KM Lighthouse

The Complete Guide for Sourcing Terms

By Heather Hedden, Taxonomy Consultant, Hedden Information Management

Taxonomies improve the findability of content over search engines alone. This is done by means of organized, semantic concepts that bring together varied names, labels, or synonyms for the same thing. To create a taxonomy, terms are collected and refined into concepts, combining synonyms or other terms whose meanings are sufficiently similar for the context of the content and the users and differentiating other similar terms, and whose meanings is sufficiently distinct for the context of the content. A single concept may be represented by more than one label or term. Definitions and/or scope notes may be added to the concepts to clarify any potential ambiguity.

The taxonomy and its concepts serve as an intermediary between users and content, whether by matching user search strings or queries to concepts or by displaying concepts for users to select from select in browsable hierarchies or lists grouped by facet-filters. Thus, in serving as an intermediary, the taxonomy concepts need to represent both the terms (words or phrases) occurring in the content and the terms that the users choose.

The Taxonomy Boot Camp presentation "The Complete Guide for Sourcing Terms," explains the importance and methods of utilizing both the body of content that is to be searched for information retrieval and input from sample users for obtaining terms for the taxonomy. In addition to direct input from users, users can serve as an indirect source of terms, if we look at search logs and uncontrolled keyword tagging. Other, external sources for taxonomy terms may also be consulted, but these provide a secondary, supplemental source for terms after the primary sources of the content and the users.

Sometimes the sourcing of terms for a taxonomy is described as either "top down" or bottom up," with a combination approach recommended. To a certain extent "top down" refers to input from users, and "bottom up" refers to term sources from the content. But following this dichotomy may lead to omitting some term sources, such as search logs and uncontrolled keyword tagging. "Top down" and "Bottom up" could also refer to the hierarchical approach to building a taxonomy, which could be either by (1) starting at the "top" of the hierarchy with the broadest categories or facets and extending the hierarchy downward to add narrower terms at each level, or (2) starting with the most specific terms, which have been gathered through various means and grouping them into hierarchies. Regardless of the directional approach, the terms still need to be sourced from somewhere, and that is the focus of this conference presentation.

Sourcing terms from content can be manual or automated. Automated methods can extract far more terms from a greater number of documents (perhaps all documents and pages) in a shorter period of time, but then you might be left with an extremely long list of terms, that, even if sorted by relevancy scores, can be difficult to deal with. Manual content analysis is time-consuming and can only deal with a representative sample of documents, but the smaller number of concepts identified may be of better value. Using generative AI and LLMs to identify terms in a body of content is useful by enabling the use of a more focused and targeted approach through formulating restrictive prompts. This is best done for parts or individual branches of the taxonomy and not the entire taxonomy at once.

Sourcing terms from users is just as important, because this way the taxonomy aligns with user needs. Methods of obtaining user input (interviews, focus groups, interactive group activities) go beyond just obtaining terms, but also provide useful information for the taxonomy structural design and user interface implementation. Sometimes, however, users are not available for input into terms, and that's why search logs and tagging reports and tagging reports can be especially valuable.

Heather Hedden (heather@hedden.net) is an independent taxonomy consultant.

How Atrion stopped their retirement brain drain and made smarter strategic decisions with AI

By Vanessa Liu, Co-Founder & CEO, Sugarwork

Atrion (now a part of Nordson Corporation) was a leading supplier of medical devices and components to niche markets in the health care and medical industry. It had around 750 employees. Two thirds were highly-skilled blue collar workers, and the remaining one third—the engineering and R&D department—were college-educated in engineering or related fields.

THE CHALLENGE: RETAINING NICHE TACIT KNOWLEDGE HELD BY HARD-TO-FIND EMPLOYEES

As its blue-collar workforce began to retire, fewer younger workers were interested in learning the skills Atrion needed. The organization wanted to de-risk a situation in which a small number of employees held knowledge critical to the success of the business, with a structured approach to retaining this specialized knowledge in-house.

COLLABORATING WITH SUGARWORK FOR TACTICAL KNOWLEDGE TRANSFER

Atrion chose to partner with Sugarwork, a solution that leverages AI to capture undocumented knowledge and create instantly-usable formats, to ensure they retained the valuable tacit knowledge held by their blue collar workforce. However, they soon widened the scope across other departments, including engineering, manufacturing operations, and talent development.

VISIBILITY INTO KNOWLEDGE CAPTURE AND TRANSFER FROM A MANAGEMENT PERSPECTIVE

Knowledge transfer is a task that is important but not urgent, and is frequently pushed back when emergencies crop up. However, with Sugarwork, senior management is able to log in to track progress. This means the initiative was treated as a priority through the organization.

> *"Before Sugarwork, I could communicate to HR the need to capture knowledge, but it's hard to understand how thorough the initiative is. Sugarwork changed this by making it easy for me to understand the depth of knowledge we were capturing and where we were in the process," said* David Battat, *former CEO of Atrion.*

LEVERAGING SUGARWORK TO CONNECT STRATEGY AND TALENT

From a management perspective, there was an extra challenge around talent. In this highly skilled and regulated industry, finding the right hires takes time, and onboarding them can take up to a year. One or two key people leaving the company can result in it losing the race to file a patent, which means that months or years of investment and future potential profits are lost.

Therefore, while having the right skills on hand to execute on strategy is critical, it is easier said than done.

> *"I've gone into meetings and said 'I think we should do X', and received enthusiastic responses. However, the reality can look very different when you dive down several layers deeper and ask 'Do our people possess the right skills? Should we invest millions in this specialized equipment just yet? What if this key person leaves?'" said David Battat.*

That's where Sugarwork outputs which provide insights into the skills across the organization were invaluable in terms of providing a tool for the management team to quickly and accurately understand what their internal capabilities were.

> *"With Sugarwork, all of a sudden I had this map of where we were, and it was very clear to me that some of the products we wanted to pursue were not possible unless we added a lot of talent. Given the precious time and investment it takes to find and onboard skilled talent, this was hugely valuable," said Battat.*

RESULTS: RETAINING TACIT KNOWLEDGE FROM SKILLED EMPLOYEES, AND ALIGNING STRATEGY WITH TALENT

Sugarwork delivered two key outcomes to Atrion:

It ensured the organization was able to retain tacit knowledge from hard-to-find employees across a range of departments, in a structured and efficient manner that gave management insights into the process and depth of knowledge captured.

It enabled the management team to align strategy closely with talent management, ensuring they could really execute on their plans, while reducing the risk of being late to file patents.

> *"When you get to the strategy level, you really need to draw the connection between growth and talent, and what you can execute on is what matters for growth. Sugarwork was valuable for us in terms of knowledge transfer, but also in terms of executing our strategy, by making sure we had the right talent," said Battat.*

Vanessa Liu (vanessa@sugarwork.com) is Co-Founder & CEO, Sugarwork.

GenAI and AI Depend on Good Search & KM Strategies

By Juanita Olguin, Senior Director, Product Marketing, Coveo

We're in the "show me" phase of generative AI, where people need to see real results to believe the hype. Every day, there's another article questioning if GenAI can actually live up to its promises. Sure, a flashy demo or proof of concept is nice—but what about handling the real needs of an enterprise? Complex, secure, and scalable solutions are what enterprises truly need.

Here's the reality: GenAI and AI rely on good search, which relies on well-organized content and data, which ultimately depends on good human input. We're at an interesting point where two often-overlooked elements, search and knowledge management, are suddenly crucial. Yet, they're still widely misunderstood.

By now, it should be clear that solid search and knowledge management are the backbone of any effective AI implementation. But they're not getting enough attention. Why? Because organizing data and content is hard work, and it's a lot easier—and more glamorous—to talk about AI. But without investing in the "unsexy" parts, companies will keep struggling with massive amounts of siloed information.

This is the moment for search technologies to shine. Foundational capabilities like term weighting, TF-IDF, BM25, and advanced filtering for big databases (think Google Drive or SharePoint) all help make sense of structured and unstructured data. Proven AI tools, refined over the last decade, have already been using these techniques to learn from user behavior and deliver relevant recommendations.

As Dan Shapiro, Associate Director of Enterprise Knowledge at Organon said in a recent webinar, "When we talk about AI nowadays, it's really just a glorified search engine."

"AI Search" as it is now being called, has been empowering companies with secure, scalable, and transparent tools for years. Generative Answering is the next evolution (not revolution), enhancing productivity and boosting revenue—but only if supported by strong knowledge and content management.

Frameworks like Knowledge-Centered Service (KCS), backed by search technology and collaboration across departments, can drive huge benefits. But it takes a coordinated effort—people, process, and technology all need to work together. IT and business teams have to join forces, with subject matter experts and knowledge managers empowered to make the most of the tools. As Shapiro put it in a recent webinar, "Search is a mirror… GenAI is a funhouse mirror." Keeping content fresh and relevant is essential for tuning both search and GenAI as they evolve.

TAKING AN AGNOSTIC APPROACH: BREAKING DOWN SILOS, NOT LOCKING PEOPLE IN

Today's enterprise challenges stem from decentralized decision-making, departmental budgets, and siloed platforms. While every platform vendor is looking to be the one-stop shop, content and knowledge rarely resides in one place; rather it resides everywhere, is ever changing, and ever growing. An agnostic approach means leveraging AI search and knowledge discovery solutions that work across siloed platforms, with the end user in mind, finding the best information to solve that user's need—whether it's browsing, learning, getting help, or onboarding. It's not about tech lock-in, it's about tech interdependency and unlocking knowledge silos that have persisted for far too long.

THE RESULTS SPEAK FOR THEMSELVES

For businesses that get this right, the outcomes are real and measurable:

- SAP Concur lowered its cost-to-serve by €8 million
- Salesforce avoided 20K cases per year
- Xero improved self-service success by 20%

These are the kinds of challenges, opportunities, and solutions we talk about with our clients every day. Let's keep the conversation going and explore how we can help you get the most out of AI, GenAI, search and knowledge management practices.

Juanita Olguin (jolguin@coveo.com) is Senior Director, Product Marketing, Coveo

THE 7 KNOWLEDGE MANAGEMENT TRENDS SHAPING 2025

Discover the **Top 7 Knowledge Management Trends**, from AI integration to cross-functional sharing, shaping the future of knowledge management.

Explore key trends transforming knowledge management, including leveraging AI and data governance, integrating seamless technology ecosystems, and focusing on productivity to maximize growth. Learn the importance of clean data for AI success, consumer-grade enterprise tools, and aligning investments with measurable business outcomes. By adopting these trends, businesses can drive strategic innovation and achieve sustainable success in a competitive landscape.

Featured in:

Forbes tdwi

CFO DIVE

Read the full article here!

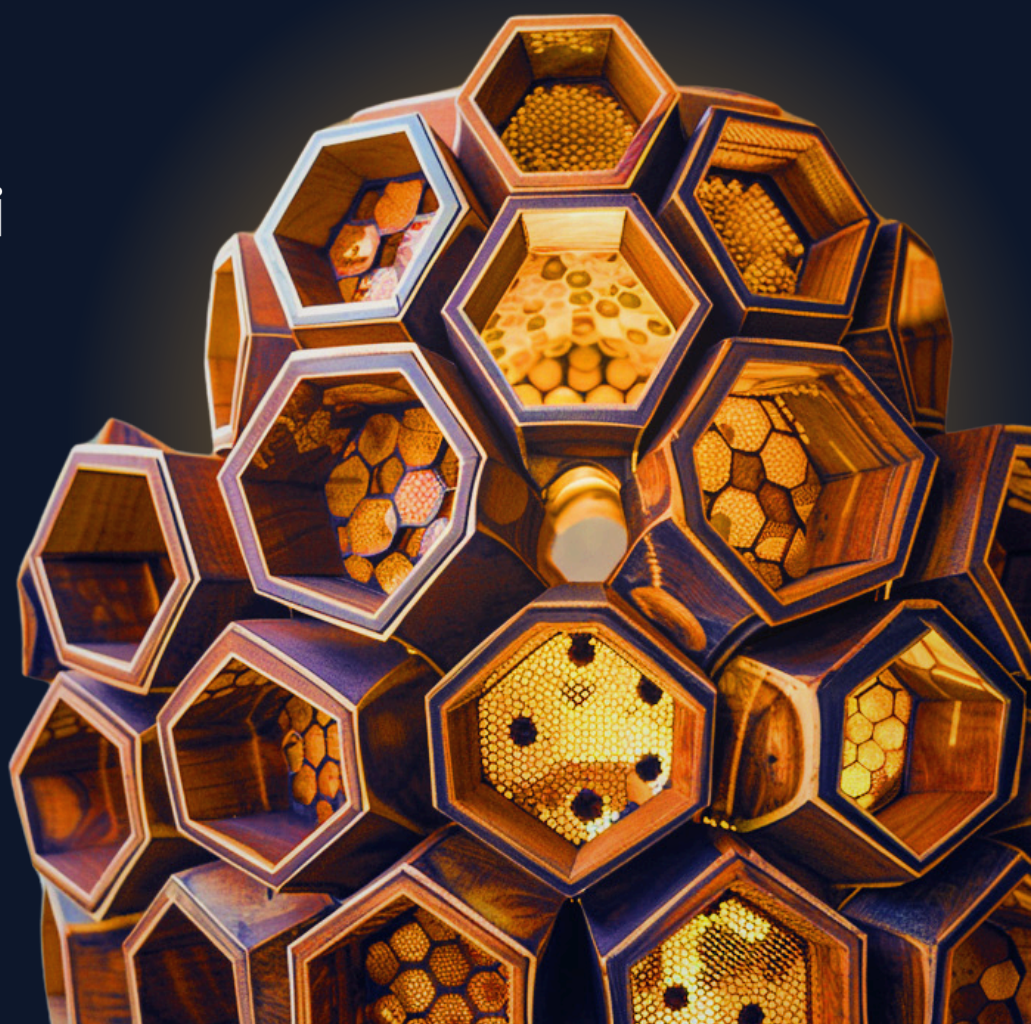

Bloomfire

GenAI Readiness and Why It Matters for Your Company

By **Charity Queret,** Chief GenAI Consultant, GenIQ Advisors

Whether we like it or not, Generative AI, or GenAI, is becoming an integral part of modern business. While some companies are already leveraging it, others remain unconvinced by the hype. Regardless of how companies feel about Generative AI, it will change how business is done. Companies wanting to compete and stay relevant will need to integrate it in some capacity.

However, adopting Generative AI is not as simple as upgrading to GenAI-ready software or letting your employees use ChatGPT. Before starting any Generative AI initiatives, companies should be asking themselves "**Are we GenAI ready?**" Answering this question requires a clear understanding of the company's current capabilities and existing gaps. It also involves identifying the steps needed to align AI efforts with business goals.

WHY GENAI READINESS MATTERS

GenAI holds the potential to create new value for companies across all industries. However, the successful adoption of GenAI requires more than just hiring data scientists to build a LLM or acquiring the latest technology; it requires a readiness to align AI initiatives with a company's business goals, data resources, and people.

A GenAI readiness assessment serves as a foundation for companies to evaluate where they currently stand, what gaps exist, and how to bridge those gaps. Without such a structured approach, AI projects will fail, leading to wasted resources, low adoption, and missed opportunities.

1. GENAI LEADERSHIP AND STRATEGY

GenAI adoption begins at the leadership level. Companies must recognize that implementing Generative AI merely for its own sake will not lead to success. GenAI initiatives provide value only when they support strategic business goals.

Establishing a cross-functional steering committee that includes leaders from technology, data, operations, and business functions helps to ensure alignment across the company. This leadership group plays a critical role in defining strategic objectives, identifying business priorities, and setting timeframes for GenAI adoption.

2. TECHNICAL PREPAREDNESS

Technical preparedness is another crucial element of GenAI readiness. A comprehensive assessment should test the compatibility of existing systems with GenAI tools.

This includes evaluating data cleansing capabilities and ensuring integration with current platforms. It also involves identifying the technical support needed for ongoing AI operations. Ensuring technical readiness reduces the risk of unexpected failures during deployment.

3. DATA QUALITY AND GOVERNANCE

Data quality and governance are at the heart of GenAI's effectiveness. As companies move to implement GenAI, it is critical to have data that is dependable, consistent, and well-governed.

An effective GenAI readiness assessment includes evaluating data sources for quality, assessing data governance frameworks, and ensuring data compliance. A focus on data management ensures that the insights generated by GenAI are trustworthy and actionable. Ensuring compliance with data privacy regulations, like GDPR, is not just beneficial but often legally required, adding another layer of importance to data governance.

4. TALENT AND TEAM STRATEGY

The successful implementation of GenAI requires the right mix of talent. A GenAI readiness assessment identifies talent gaps and establishes a plan for addressing those gaps. This may involve hiring data scientists, AI engineers, and subject matter experts or partnering with AI consulting firms or specialized GenAI vendors.

Upskilling current employees, to ensure they understand AI ethics, prompt engineering for language models, and basic data analysis ensures that the entire organization is ready to adopt and leverage GenAI, building a culture that embraces AI-driven innovation.

5. PILOT PROJECTS

To effectively implement GenAI, companies should begin with pilot projects that allow them to learn through execution. Pilot projects offer a controlled environment to test AI capabilities, assess performance, and gain insights into areas for improvement. By choosing a well-defined internal use case, companies can validate AI readiness, measure outcomes, and make informed adjustments.

A successful pilot project not only validates GenAI readiness but also helps inform the best practices for scaling AI across other departments.

CONCLUSION

Preparing for the GenAI revolution requires an intentional, structured approach. By focusing on leadership alignment, technical readiness, data quality, talent strategy, and pilot projects, companies can ensure they are set up for success.

Companies that take the time to understand their current state, identify and address readiness gaps, and align GenAI initiatives with strategic objectives are more likely to see GenAI adoption. This, in turn, will help them gain and maintain a competitive advantage.

THIS IS WHY GENAI READINESS MATTERS FOR YOUR COMPANY!

Charity Queret (charity.queret@geniqadvisors.com) is Chief GenAI Consultant, GenIQ Advisors

Enabling Thoughtful Conversations at Massive Scale

By Louis Rosenberg, CEO, Unanimous AI & Author, Our Next Reality: How the AI-Powered Metaverse Will Reshape the World

Imagine if a large enterprise team of 400 people could hold a thoughtful real-time meeting in which they brainstorm ideas, share knowledge, debate alternatives, prioritize options, and converge on optimized solutions that leveraged their diverse skills, expertise, and insights. After all, conversational deliberation is the single most important collaboration method for problem solving, strategic planning, and information exchange among enterprise teams.

Unfortunately, the ideal size for a productive real-time discussion is only 4 to 7 people, whether the group gathers in-person or remotely. In small groups, each individual gets a good amount of airtime to express their views and has relatively low wait-time to respond to others. But as group size grows, airtime drops, wait-time rises, and by a dozen people devolves into just a series of monologues. Above 20 people, forget about deep deliberation—it's pure chaos.

So how can 40 people hold a thoughtful conversation, or 400, or even 4,000? This is now possible using a new generative-AI technology called Conversational Swarm Intelligence (CSI). It was inspired by the unique dynamics of fish schools, which can easily have thousands of members, each with a different view of the world. The amazing thing about a fish school is that it can quickly share knowledge across the full population and almost instantly make optimized life-and-death decisions.

Fish can do this because they have a special organ that allows them to track the speed and direction of neighboring fish based on pressure changes in the water around them. Using this organ, they can *deliberate* in small local groups of neighbors that collectively decide on the direction that portion of the school should go. This is interesting but does not explain how rapid global decisions are made. The magic happens because every fish in the school "deliberates" with a different group of neighbors. This means there are many "overlapping conversations" happening at once which enables information to quickly propagate across the full school.

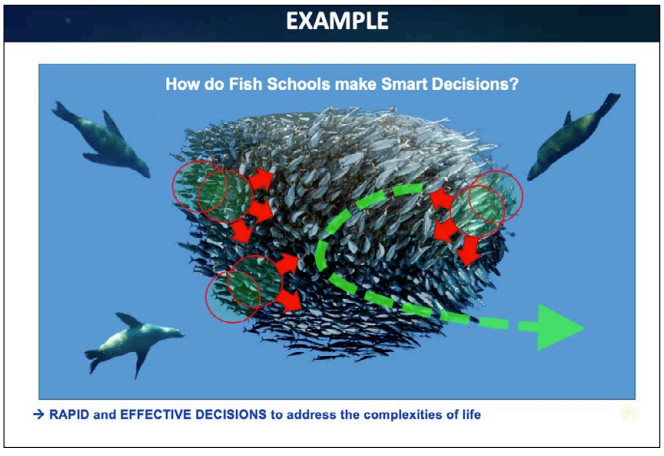

CSI emulates fish schools by empowering large human groups to hold small overlapping conversations, but it's not as simple as it sounds. That's because we humans did not evolve with the ability to be in multiple conversations at once. This is commonly called the "cocktail party problem" because it happens often when small groups gather within earshot of each other. If you try to pay attention to a neighboring conversation you immediately lose track of the discussion you are in.

SO HOW CAN WE OVERCOME THIS HUMAN LIMITATION?

The technology of CSI works by first breaking any large group into a set of parallel subgroups. It then adds an AI-powered conversational agent into each subgroup. We call these "Surrogate Agents", and they are tasked with distilling the real-time human insights within its local group and sharing those insights with other groups. This enables all the local groups to overlap, weaving the local conversations into a single large conversation among all the members. And it works, enabling hundreds of people to brainstorm, prioritize, and problem solve in real-time.

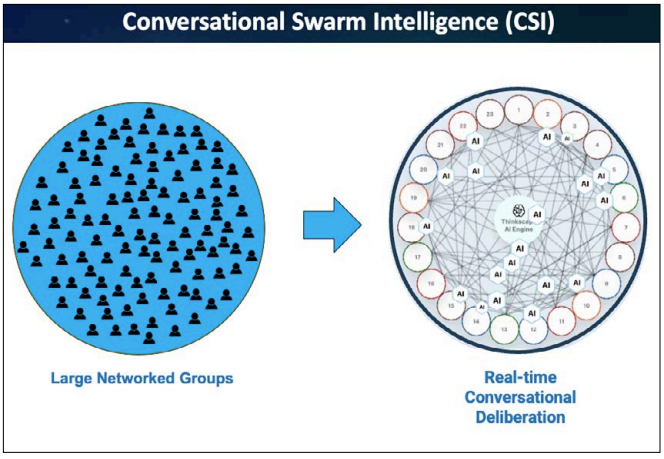

Even more exciting, CSI technology has been built into a commercial platform called Thinkscape from Unanimous AI. It allows real-time deliberations among groups as large as 400 people. In collaboration with researchers at Carnegie Mellon, rigorous testing has been conducted using Thinkscape to assess the value of CSI among large, networked teams. In a 2023 study, researchers tested 50 person groups in Thinkscape (using CSI) vs traditional methods. They found the CSI groups were able to hold more coherent conversations and that each individual was found to contribute 50% more content than traditional methods.

BUT DOES CSI AMPLIFY TEAM INTELLIGENCE?

To explore this, a 2024 study by researchers at Carnegie Mellon and Unanimous AI tested the ability of networked groups to take IQ tests. Results showed that groups of 35 people, who averaged an IQ of 100 (the 50th percentile), could score an effective IQ of 128 (the 97th percentile) when deliberating together using CSI-powered Thinkscape. Although this study used groups of only 35 participants, other recent studies have tested groups up to 250 and confirmed amplified intelligence. Overall, Conversational Swarm Intelligence (CSI) is a unique and powerful Generative AI technology with the potential to revolutionize communication, collaboration, and collective intelligence across large groups many professional settings from enterprise

collaboration and market research to civic engagement and deliberative democracy. Already some of the world's largest market research firms are using Thinkscape to conduct qualitative focus groups with hundreds of people at a time. And the US Air Force is currently exploring the use of Thinkscape for brainstorming and need-finding across large groups.

Considering that the average product team in a Fortune 5000 company has over 200 members, CSI technology can finally enable real-time deliberative meetings that allow participants to share and amplify their knowledge, expertise, and insights.

Louis Rosenberg, PhD *is a well-known researcher in the fields of artificial intelligence and immersive media. He is known for founding Immersion Corporation (IMMR: Nasdaq), Outland Research, and Unanimous AI, and for developing the first functional augmented reality system at Air Force Research Laboratory.*

Beyond AI: Low Tech for High Stakes:

By Pascal Saura, Senior Knowledge & Learning Officer, The World Bank

Picture yourself as a water specialist joining the World Bank. On your first day, everything feels unfamiliar and a bit unsettling. You are suddenly representing a leading development institution, and are expected to deliver top-notch solutions grounded in 80 years of global operations. Your first assignment may be in a country office. You don't know any of the other World Bank Water experts, and you find yourself several time zones away from most of them anyway. Still, you are the sole interlocutor of government officials who expect you to be the voice of the World Bank. How will you live up to that? Fortunately for you, the Water Global Practice has a service desk. You can send an email that will trigger a chain reaction of research, contacts, and answers to your questions. An entire community will come to your rescue and will help you speak with the confidence that your own expertise has been augmented by the knowledge of the institution you are representing.

How is such a service created? How does it work? Is AI the silver bullet? This is about low tech coming to the rescue of high stakes.

ASKWATER

The World Bank's Water Global Practice has been operating a service desk named "AskWater" since 2016. *AskWater* accompanied the creation of the first National Water Institution in Papua New Guinea, informed the dialogue with Cape Town through the "Day Zero" crisis, and helped deploy water information systems in Nicaragua, the Dominican Republic, the Lake Victoria basin countries, and Moldova to name a few. It provided benchmarking and lessons learned for the establishment of countless water supply and sanitation national strategies. Connecting challenges and solutions around the world, the *AskWater* Service Desk linked a project on urban stormwater management in Djibouti with findings from Brazil and approached fecal sludge management challenges in Lusaka using past successes in Sri Lanka. Beyond answering questions, the Service Desk supports experimentation: engagement with private sector operators in Bangladesh, involvement of women in water resources management bodies in Peru, wastewater surveillance methodologies to detect new waves of COVID-19 infections in Ecuador and Uruguay, and hotspots in Kenya, among many others.

IN THE LONG RUN: ACCRUING BENEFITS.

Considering why and how *AskWater* persisted as a successful service desk, four basic principles of success stand out: strong managerial endorsement, a light and nimble structure, a deep culture of collaboration and a keen attention to individual recognition. But benefits are accruing.

A well-designed service desk helps unpack tacit knowledge and codify lessons learned. It provides an "always-on" connection between those who know and those who use knowledge for implementation, closing what is at times called the "knowing-doing gap".

Unpacking tacit knowledge, codifying good practices and curating critical information. The answer to a complex question is a powerful driver for institutional knowledge and performance: it comes with data, information, lessons learned and an expert point of view that tailors the answer to its context, potentially alerting the requestor to sensitivities created by context while echoing strategic directions critical to the organization.

Linking knowledge to action. There is a classic and irritating gap between what an organization knows and what it does. Service desks activate a dialogue between knowledge and implementation. The collective and open resolution of challenges ultimately becomes a formidable learning machine, anchored in praxis and just-in-time problem-solving. It also originates from an immediate and present need, if not desire, for learning, which brings the finishing touch to a convenient alignment with the principles of effective adult learning.

Building an expert community. Successful service desks develop a symbiotic relationship with communities of practice. To create all the value it is capable of, the service desk cannot be a black box. By being very deliberate about attributing answers to the people who contributed their expertise, it recognizes experts, incentivizes knowledge sharing and helps build a knowledge community.

GO-TO KNOWLEDGE VS. KNOWLEDGE TO GO FURTHER.

Channeled through identifiable people and a visible product, the knowledge management function offers a strong *pull* (a practice activated by daily questions from the frontline) rather than a tedious *push* (*"Will you please consider using this library of lessons we've learned so far?"*). Funding, behaviors, and support for knowledge activities follow this pull, for the same reason that adoption follows satisfaction.

Service desks take advantage of the *pull* of questions in order to *push* good practices and new priorities. A good answer not only addresses the question at hand but also elevates it, and the best answer may lead to reshaping the initial thinking and changing the conversation. It can also recognize that the question is new and that the search for a solution will lead to experimentation. In transitioning from providing "*go-to knowledge"* to leveraging "*knowledge to go further"* the *AskWater* service desk was paired with a small grant facility that sources external expertise, funds innovative pilots. While relentlessly addressing the needs of the client-facing teams, it also promotes forward-looking organizational goals, ideally helping to solve today's emergencies in light of tomorrow's directions.

INTO THE MAELSTROM: *MOBILIS IN MOBILE*

Service desks have more data, better answers and faster responses powered by search engines and AI. Manned helpdesks may disappear. But the people who run service desks and orchestrate the answers to knowledge questions should remain. They will see their own intelligence augmented by technology. The World Bank is investing in AI to make sense of 80 years of development. Through interdependent economies, far-reaching supply chains and mass migrations, development challenges, climate-related disasters and health emergencies reach a global scale at an increasingly fast pace. Good service desks must have eyes everywhere and connect experts around the word at a moment's notice, making everyone smarter and reducing seemingly undecipherable complexities. Embracing knowledge for development is learning how to move within a moving environment. What is required to navigate this maelstrom of questions and answers are steady, ingenious, and intelligent captains who know something about knowledge management.

Pascal Saura (psaura@worldbankgroup.org) is Senior Knowledge & Learning Officer, The World Bank.

Enterprise Approach to AI for CX

By Robin Oswald Schwartz, Senior Knowledge Platform Manager & Strategist, Specialized Global Rider Care

Specialized would like to take you on our customer service knowledge journey. You'll see where we were, where we are now, and where we're headed.

WHERE WE WERE:

In 2023, the Specialized Global Rider Care team identified a crucial vulnerability. The platform we used for our customer self-service website was going out of support and at risk of failure. At the same time, our customer service agents were fielding record numbers of calls and messages. Our content team had created more than 1,200 knowledge articles available in 22 languages; however, our customer service agents struggled to find the relevant answers to customer queries. We saw an opportunity to revolutionize our customer service through knowledge management to provide a premium experience for customer self-service and improve agent interactions.

We started with discovery to understand who was creating content, where it lived, and who used it. We also surveyed and met with our customer service teams to understand their pain points. Next, we built a business case for our new investment that quantified our risk of continuing the status quo vs. the opportunity for cost savings and improved customer satisfaction. We used those metrics to guide us in writing an RFP and evaluating different knowledge platform vendors to choose the one that best fit our needs.

WHERE WE ARE:

We onboarded eGain as our knowledge platform in early 2024. The first phase consisted of migrating our customer self-service website content to the new platform and recreating the web portal. We completed this in May 2024, just as our busy product launch season began.

Our team is using the analytics capabilities in the platform to identify areas for improvement, such as infrequently used content, poorly rated content, or content that is not easy to find. We are developing templates for new content and beginning to unite different content authors into our single source of truth. AI tools built into the platform help us improve our consistency and efficiency in content creation, migrate content from legacy sources, and improve the accuracy of our website search. We also added an online troubleshooting guide for our retailers and riders to guide them through steps in resolving problems with e-bike systems.

WHERE WE'RE GOING:

We've created a multi-year plan to integrate additional content sources and teams into our platform as well as surface our knowledge in different areas. Our goal is to have content authors create knowledge once that can be used many times for different applications and audiences. We are learning how to use the AI tools within our knowledge platform to identify content that is outdated or inconsistent, summarize longer forms of content to provide quick answers to consumer queries (as Google does), and provide translated internal knowledge for our customer service agents worldwide. We also anticipate using AI to help us triage customer needs into simple requests that can be answered with self-help content vs. more detailed conversations that require a personal touch.

Robin Oswald Schwartz (robin.oswaldschwartz@specialized.com) is Senior Knowledge Platform Manager & Strategist, Specialized Global Rider Care

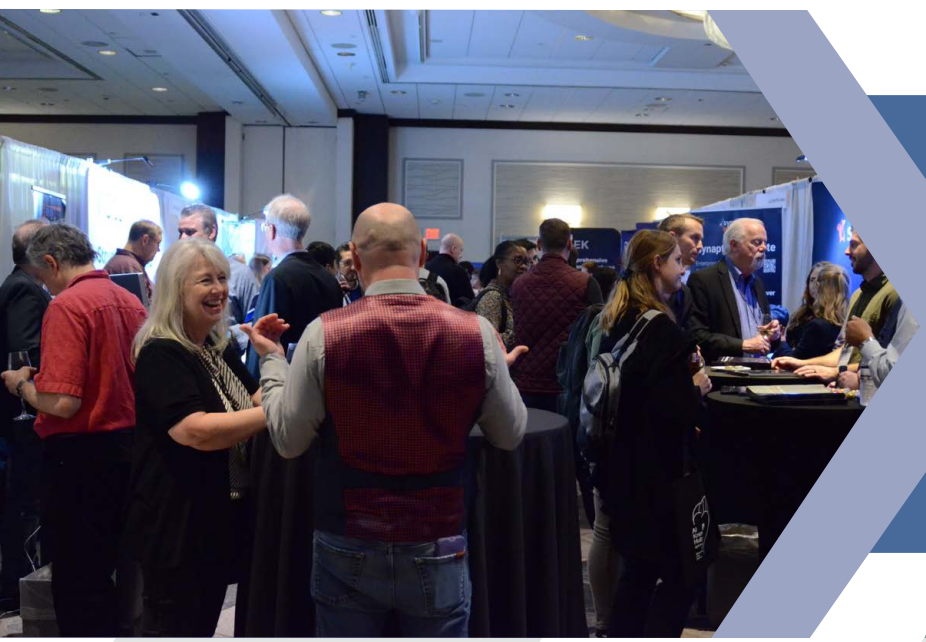

The Rise of Next Enterprise Search: Why Information Architects Matter

By Marianne Sweeny, Daedalus Information Systems

The SEO community's current struggle with Google AI Overviews, it AI answer feature, offer a stark warning for those working in enterprise search. Like King Lear raging against the storm, SEO practitioners watch helplessly as their traditional methods crumble in the face of intelligent search algorithms. Their resistance to change, rooted in hubris, serves as a cautionary tale for the enterprise search community as we navigate our own shifting landscape of AI-driven information retrieval.

Enterprise search has adapted somewhat better to these changes, though significant challenges remain. Recent research from Pryon AI confirms that enterprise AI adoption is not merely a trend but represents a permanent shift in how organizations manage and access information (https://www.pryon.com/resource/how-ai-will-transform-enterprise-knowledge). Major corporations like Costco, Walmart, Amazon, and Adobe are leading the way by implementing knowledge graphs built on proprietary ontologies and taxonomies to produce more highly relevant search results. However, they are the exception and not the norm. Many of these sophisticated AI feature's potentials are not fully realized.

Figure 13. What still needs to be done to fully implement an enterprise AI solution? (Multiple responses permitted)

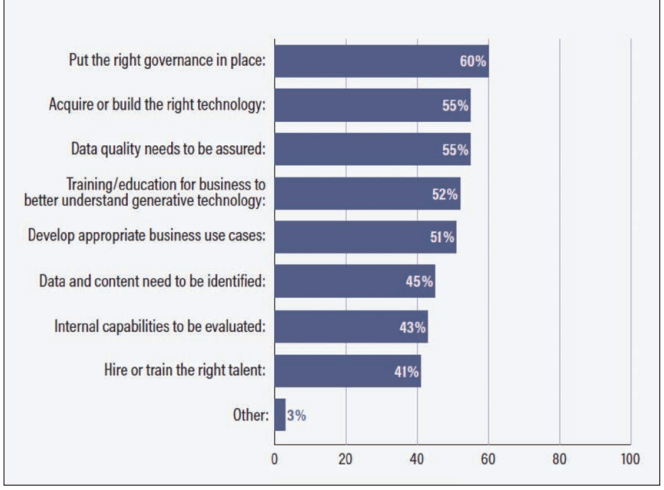

To avoid the SEO community missteps, enterprise search must forge a new path that builds on Next Information Architecture to deliver Next Enterprise Search. This integration demands breaking down organizational silos and embracing AI technologies more holistically. Success requires us to understand that machines are now users alongside humans in our information ecosystem.

Systems thinking provides the foundation for this transformation. Rather than focusing on individual components, this approach examines the relationships and interactions within complex systems. By understanding these connections, feedback loops, and leverage points, we can design more effective search solutions that serve both human and machine users.

Peter Checkland's Soft Systems Methodology (SSM) offers a proven framework for the cross-disciplinary collaboration needed in AI development. This approach, developed in the 1970s, remains remarkably relevant for today's challenges. SSM emphasizes creating teams that include representatives from all key disciplines and establishing clear communication channels between different specialty areas. A project can only begin once the team has developed shared vocabularies, frameworks, and a unified vision for success.

Information architects bring crucial expertise to AI development that extends far beyond traditional information organization and should be part of an AI SSM project. Their deep understanding of information structure and experience with cross-platform consistency provides essential insights for AI system design. Their ability to comprehend complex relationships between system components and create detailed data maps can be invaluable when developing AI applications.

Most important, information architects excel at aligning information organization with users' mental models, thereby enhancing system usability and trust. This human-centered perspective becomes increasingly crucial as AI systems grow more opaque to end users.

This path forward requires bilateral learning and growth. Information architects must expand their knowledge into computer science and AI development requirements. Resources like Google's People + AI Research (PAIR) provide valuable tools and instruction to help human factors professionals develop a deeper understanding of AI complexities. Meanwhile, development teams must learn from information architects about the principles of information structure, cross-domain consistency, and the crucial importance of user mental models.

This transformation won't be easy. It requires sustained commitment from all parties involved and a willingness to step outside traditional role boundaries. However, the alternative—continuing to develop AI applications without proper information architecture input—risks creating systems that fail to meet user needs or achieve their full potential.

The time for Next Enterprise Search has arrived. By learning from past mistakes and embracing cross-disciplinary collaboration, we can create AI systems that truly serve both human and machine users while avoiding the fate of our SEO predecessors. The storm of technological change may be upon us, but with proper preparation and collaboration, we can achieve more that fist shaking anger on social channels.

Marianne Sweeny (marianne.sweeny@gmail.com) is Principal Consultant, Daedalus Information Systems.

KMWorld 2025

CO-LOCATED WITH

Enterprise Search & Discovery

Text Analytics Forum

Enterprise AI World

taxonomy BOOT CAMP

FIVE Events. **ONE** Location.

Unlocking Powerful Knowledge-Sharing Solutions by Harnessing Search, Discovery, Analytics, Taxonomy, & AI

November 17–20

JW Marriott | Washington, DC

ORGANIZED AND PRODUCED BY

Information Today

#KMWorld | #TaxoBC | #ESDDC | #TextAnalyticsDC | #EnterpriseAIWorld

SAVE THE DATE!

DIAMOND SPONSORS

eGain M-Files NICE VERINT

GOLD SPONSORS

ENTERPRISE KNOWLEDGE

factor Information Architecture & Taxonomy Consulting

LUCIDEA

Progress Semaphore

MEDIA SPONSORS

Enterprise AI World

KMI KM INSTITUTE

KMWorld

KMWORLD.COM/CONFERENCE

Enabling the AI-Driven Future of Enterprise Technology: Why Unification Is the Gateway to Success

By Andy Tolton, VP of Marketing, MangoApps

Over the last 12-24 months, the enterprise technology landscape has been crossing a pivotal moment. While artificial intelligence (AI) brings a promise to revolutionize our work, a fundamental challenge threatens to undermine its potential: the fragmentation of our digital workplaces.

As organizations race to implement AI solutions, many seem to be overlooking a crucial truth—without unification of information and systems, AI may only amplify existing chaos.

For over a decade, digital transformation has aimed to streamline operations and boost productivity. However, this shift inadvertently created a maze of disconnected systems and siloed data that seems to be compounding. Moving away from paper-based processes was necessary, but it led to a proliferation of disparate applications, each with its own data repository and workflows. The cost of this has been passed down to employees, who are having to make sense of this mishmash of systems and data on their own.

Consider these eye-opening statistics:

- Employees now juggle an average of 11 applications daily, up from 6 in 2019. (Gartner[1])
- Workers spend up to **four hours each week** switching between apps to find information. (Harvard Business Review[2])
- **47%** of digital workers **struggle to find the information** they need. (Gartner[1])
- **Only 32% of workers are engaged at work**—a mere 6% increase since 2000. (Gallup[3])

Frontline workers face even greater challenges due to limited access to digital tools and enterprise systems. Without foundational digitization, organizations miss out on critical data that AI and analytics require to deliver meaningful insights.

This sets the stage for what we call the "**AI Paradox.**" As enterprise digital workplaces become increasingly fragmented, organizations have rushed to implement AI-powered solutions in a piecemeal fashion—deploying chatbots and virtual assistants within their existing, siloed applications. While these specialized AI tools aim to streamline workflows, they often have the opposite effect.

Each interaction of AI operates in isolation, accessing only the data within its own silo. For example, an HR chatbot lacks visibility into customer service history stored in the CRM system. Employees must navigate multiple AI tools, each with incomplete information, leaving critical insights locked in silos. Instead of simplifying work, this fragmentation forces employees to navigate a maze of AI tools with limited understanding of their needs.

Ironically, AI was supposed to be the great unifier and enabler—a technology to break down silos and augment human intelligence. But without an enterprise-wide approach to AI integration, we're amplifying the chaos we sought to eliminate.

This is where **unification** becomes critical. Not just simple integration of data on the backend, but a true integration of experience all the way to the frontend. Imagine a single, integrated interface that brings together every enterprise system, tool, and resource an employee needs—all personalized to their role and context. This doesn't just provide a facade of mere convenience; it establishes the fundamental conditions required for AI to truly excel.

At MangoApps, we've developed a unifying digital hub that provides this critical layer of integrated experience. By creating a single, intelligent interface that unites all of an organization's systems and data sources into one place that employees can access, we're giving you a foundation for a truly AI-driven future. This unification unlocks unprecedented levels of contextual awareness, data accessibility, and seamless user experiences—the three pillars essential for unleashing the transformative potential of enterprise AI.

1. **Comprehensive Data Access:** Our unified platform allows AI to access data across all enterprise systems—from HR and CRM to project management and knowledge bases. This comprehensive visibility empowers AI assistants to provide insights and recommendations impossible within siloed systems. By breaking down barriers between data sources, we enable new levels of contextual understanding and cross-functional intelligence.
2. **Contextual Intelligence:** We give AI a deep understanding of each employee's unique role, permissions, and needs. This contextual awareness allows AI assistants within our platform to deliver truly personalized support. Instead of navigating a maze of disconnected AI tools, employees receive tailored interactions that anticipate their requirements.
3. **Seamless Experience:** Employees interact with all their enterprise tools and AI capabilities through our single, intelligent interface. They can switch between AI assistants as needed, without disrupting workflow or losing context. This unified experience eliminates productivity drains from switching between systems.

These pillars form a fabric that will transform workflows and empower employees across an entire organization all the way to the frontline. For example:

- An **HR AI Assistant** can access policy documents, employee records, and support tickets to provide comprehensive guidance and powerful self-service.
- A **Company Knowledge AI Assistant** can pull relevant information from knowledge bases and repositories, delivering a cohesive experience to information and knowledge discovery.
- An **IT Support AI Assistant** can diagnose technology issues, provide step-by-step guides, and automate tasks—all while aware of the employee's permissions and history.

Looking ahead, unification will become even more crucial as organizations adopt new technologies and AI capabilities. The key to harnessing AI's full potential lies not just in integrating systems and data, but in connecting the AI assistants themselves. By developing API-based integrations that enable specialized AI assistants to communicate and collaborate, we can break down barriers that limit AI's

scope and impact. At MangoApps, we're exploring strategies to realize this vision of interconnected AI assistants.

By embracing unification across systems and AI capabilities, organizations can unlock AI's transformative potential and build sustainable advantages. The path forward may not be simple, but the rewards are immense. Unification isn't just the gateway to AI success; it's the foundation upon which the future of enterprise technology will be built.

Andy Tolton (andyt@mangoapps.com) is VP of Marketing, MangoApps

1. Gartner. (2023, May 10). Gartner Survey Reveals 47% of Digital Workers Struggle to Find the Information Needed to Effectively Perform Their Jobs [Press release].
2. Harvard Business Review. (2022, August 17). How much time and energy do we waste toggling between applications?
3. Gallup (2024, June 12). State of the Global Workplace: 2024 Report.

The story of the +

By Victoria Ward, Co-director Jigsaw Foresight and Caitlin McDonald, Jigsaw Research & Innovation Strategist

Close encounters of a new kind: humans + AI

"Despite the incredible impact of AI recently, the world is still struggling to appreciate how big a deal its arrival really is. We are in the process of seeing ***a new species*** *grow up around us. Getting it right is unquestionably the great meta-problem of the twenty-first century."*
(Mustafa Suleyman, CEO of Microsoft AI, writing on X)

The abrupt arrival of a young new species bursting into knowledge working and reinventing many aspects of it, is and is not, a surprise.

"The future is coming at us too fast and shimmers with patterns that are for now invisible.
Going forward, who knows what will happen?
People are nervous."
(Nora Bateson—International Bateson Institute)

Historians of work and workplace cast back to previous, watershed moments when a new technology arrived on the scene and upended assumptions about how work works: pen and paper, the printing press, the lightbulb, the spreadsheet, desktop and laptop computers, the internet and email, the mobile phone, the pandemic. We have been here before, only this time, somehow, feels different. Or is it? Is it more dangerous, more democratic, more demanding, more delighting? Or is that just how we are feeling because it's all so frenetic, so wild, and so fast?

Knowledge workers face an imperative to engage with the volatility and acceleration of change (right now, it's Generative AI (Gen AI), but in general it's any changes), and to find ways to meet the moment and embrace it, without being glamoured into rushing, then regretting, hunkering down and hiding in the hope it will blow over. The rush to adopt, thirst for productivity, magically addicting fizz of fast change, fear of missing out, being passed over, or being seen to naysay at the expense of your job—these can all encourage workers to press on rather than press pause. This riptide of change is what we are seeing play out at furious pace with Gen AI. And it can be very easy to be sucked into the adrenaline, excitement and perhaps anxiety, rather than pursue the steadier path of a more sober reality of sustained and sustainable developments.

HOW CAN FORESIGHT + METHODS (FORESIGHT + SHIFT + CHANGE + STORY) HELP HUMAN + MACHINE AND HUMAN + HUMAN TEAMING?

Jigsaw argues for intentionally structuring places in which to slow down, reflect, wonder about things and rethink relationships, both human and with and through technology: think of collaboration infrastructure that making space for change in day-to-day knowledge working realities. We have summed this up as THRIVE (Tools for Heightening Resilience with Innovation in Volatile Environments), methods that can be combined to provide the following, in containers for collaboration:

- **Provocative prompts.** Sets of questions that interrupt assumptions, slow things down a bit, and provoke a different take.
- **Place-making qualities.** Visual and beautifully hosted, meeting places which explore, entertain and get the work done, and in which common and uncommon ground can be found. Memories of time well spent together leave a lasting trace and fertilise the soil of social fabric. Trust ripens over time. Memories are something a machine does not have, so making our memories work for us matters more than ever in an era characterised by rapid and abrupt change.
- **Flexible strings.** They string easily together to give light structure and momentum to conversations—a kind of pattern language for collaboration.
- **Plurality and power sharing.** They intentionally make room for plurality—points of view, timelines of change. This in turn is useful in shifting the implicit knowledge power structures, making room for diversity and difference in a respectful and safe way.
- **Inhabited places.** They combine to develop a rich, felt, inhabited sense of what's going on today, and what might go on tomorrow, and where we might intervene intentionally between today and tomorrow to shape the future a bit differently together.

In Jigsaw learning programmes we teach people to reconfigure their ways of knowledge working, and collaborate with each other, through engaging with disruption and emerging changes in several ways:

1. **Scanning for emerging change.** Where and how do we look for and note change, in likely and unlikely places, externally and in our own experiences? Scanning beyond one's usual sources in a systematic way broadens and deepens insight into patterns of emerging change. For example, in one recent strategic workshop, when considering un-bundling and re-bundling a series of tasks to incorporate Gen AI into regular rhythms of knowledge production(for example research reports to pass on to users) a participant noticed out loud how prompting Gen AI to act as an intern had led both to faster production, and to a different way of producing outputs—to shorter bursts of prompting the Gen AI—thus more rounds of iteration and more collaboration and conversation with colleagues along the way. That kind of insight could easily be skipped over if you weren't paying attention, yet it profoundly illuminates emerging new micro-processes for knowledge management.
2. **Watching for emerging terminology and how it shapes human-machine connections.** Of particular interest are the emerging new terms, languages, metaphors and imagery arising from new technology trends that can quickly become shorthand that misdirects, sometimes rather dangerously, sometimes in a way that illuminates things we do not yet understand. Does a Large Language Model (LLM) actually learn? One recent article proposes use cases for Gen AI where the machine operates as a god (super-intelligent, autonomous), an intern (a copilot supervised by a human, sometimes operating as a toy) or a cog (functions well at doing one task, a component). Depending on how you are building those conceptual frameworks into your knowledge working patterns, these cases can either be a useful shorthand, or dangerously seductive. After all, if you have a god in mind, or you are lazy in supervising your interns, you might end up with unexpected or undesired consequences rather than truly making the most of your potential AI collaborator.
3. **Creating spaces for social learning and experimentation,** where people can puzzle together data (including sensory data, vague hunches, feelings and stories, as well as facts and apparently firm evidence). Experimentation is a way of creating new pattern languages in the face of evolving data. By making it a day-to-day practice, and a social learning habit shared among

colleagues, new knowledge can be constantly evolving to changing conditions. A world where humans + machines is a closer coexistence than ever needs closer humans + humans too. These kinds of social learning spaces (a structured flow of synchronous and asynchronous collaboration, as we have learned in the pandemic) also sets conditions in which informal trust can ripen and provide a kind of human mycelial-esque intelligence network helping humans collaborate to shape the best human + machine relationships they can.

4. **Fostering interdisciplinarity** on equivalent terms. Counterintuitively this benefits from being de-professionalised. Credentialled experts have a role to play, but collectively being open to uncertainty, and to making room for many different ways of knowing, is what can lead to truly transformative new understandings. Interdisciplinarity is also cross-cultural, intergenerational, intersectional, and interested in putting philosophers, artists and poets, and those whose fine motor skills represent embodied knowing, on a level footing with scientists, and operational experience on a level footing with domain-specific qualifications and expertise.
5. **Being open to standing still for a while** and wondering 'what if'. Being willing to stand still, look around, and be immersed with others, and challenge assumptions about what is going on under the surface and what kinds of futures are emerging. Sometimes the most effective reaction is not to react, but to wait: the Gartner hype cycle shows us the extreme highs and lows of change which comes from being an early adopter. Allowing the change to marinate a little and noticing its effects can encourage a deeper understanding of how different players are reacting to change: who are the heroes, who are the villains, and who are the sidekicks? What opportunities are there to engage with change in different and perhaps unexpected ways?

We'd argue that these are the conditions of effective knowledge working (including human + machine knowledge working)—a vital way to repurpose time and space, bringing new attitudes to create more effective collaborative working spaces.

"When you create a Human + AI team, the hard part isn't the "AI". It isn't even the "Human".
It's the +" (Nicky Case 'How to become a centaur,' JoDS, January 2018.)

With this 'new species' that is growing up around us, we as humans must grow up a bit too. Perhaps need to accept an invitation to evolve as a species. Joseph Aoun, in 'Robot-Proof' argues for humanics, a new set of human disciplines—deepening in our interior worlds those things we can do that machines cannot easily emulate—mental elasticity, critical and design thinking, empathy, imagination, human to human relationships. By deepening humanics, we meet each other and machines differently. And It's the plus that we're aiming to work on. Learning to embrace the pluses that new kinds and species can bring is key for leaning forward into new possible futures. In ancient Greek mythology, the god Poseidon brought forth the horse, a new and wondrous species of great power—but the goddess Athena invented the bridle, enabling humans to work alongside this new species as companions and coworkers. Jigsaw's learning programmes allow knowledge workers to harness the power of the plus.

Jigsaw Foresight, originated by Victoria Ward and Wendy Schultz in 2006 is a futures, foresight and change management collective equipped to help you solve problems from the personal to the planetary. THRIVE (Tools for Heightening Resilience with Innovation in Volatile Environments) is a Jigsaw method collection mixing selected foresight, narrative and other methods. Its purpose is to strengthen critical thinking; deepen collaborative capacity, confidence and creativity; provide the mechanisms to structure a pattern language for enduring deep and bold collaboration; and, help users build connected islands of stability in a sea of change.

Victoria Ward (victoria@jigsawforesight.com) is Co-director Jigsaw Foresight, and Caitlin McDonald, (caitlin@jigsawforesight.com) is Jigsaw Research & Innovation Strategist.

Written and edited by humans

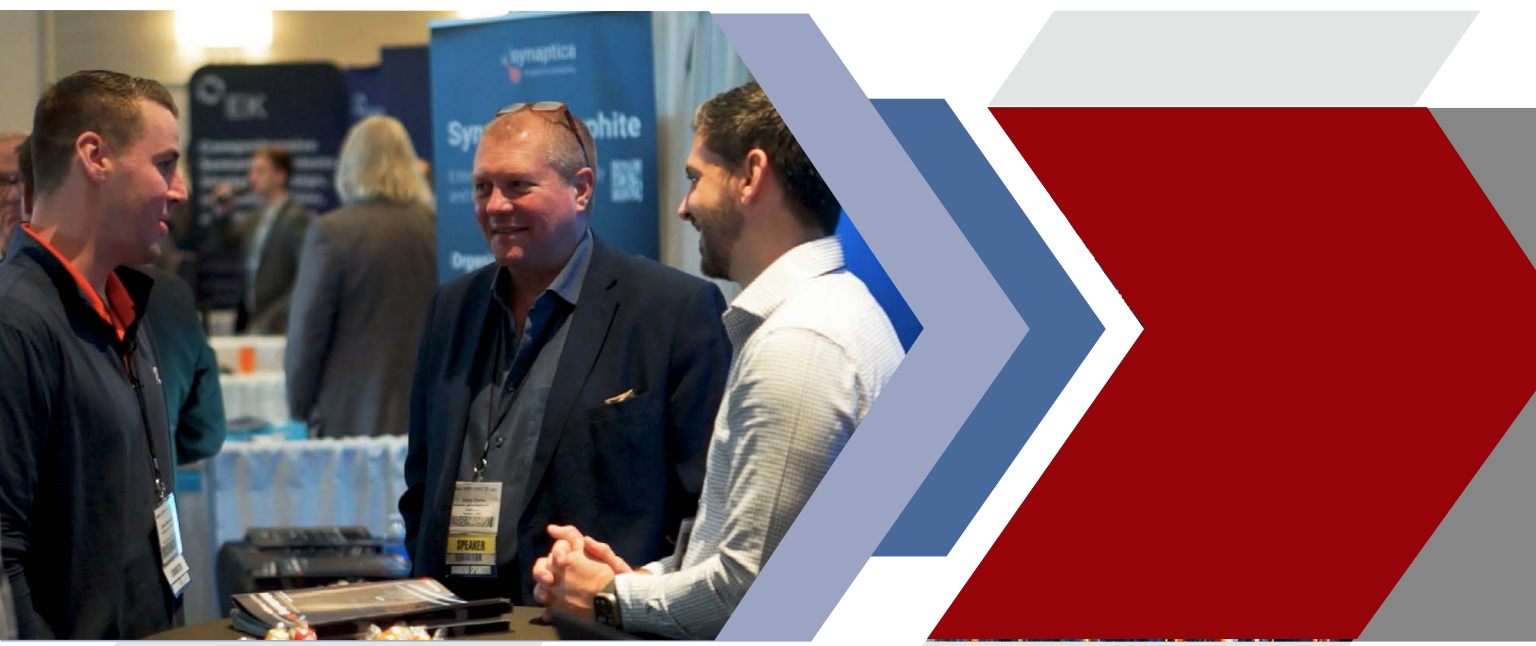

Beyond Chat: How Enterprise AI Agents with New Superpowers are Transforming Knowledge anagement

Zeta Alpha

By Jakub Zavrel, Founder and CEO, Zeta Alpha

As enterprises move beyond the initial excitement of prototyping chat-based generative AI, business leaders are increasingly seeking tangible returns on their AI investments in terms of productivity, cost savings, and growth. One of the most exciting current paths forward is AI Agents (Robison, 2024). For knowledge management (KM) professionals, the key question now is how to develop and deploy AI agents, and go beyond chat interactions and question answering to perform sophisticated knowledge-intensive work as true intelligent assistants.

AI Agents work by chaining together building blocks composed of Large Language Model (LLMs) prompts combined with larger quantities of data than humans can typically digest. AI Agents break down goals into simpler tasks to address sophisticated domain-specific tasks that might take people days of work to complete, a key characteristic of work that generates real business value (Rosenbush, 2024), and interact with existing IT systems through APIs.

Imagine AI agents that not only answer questions and create reports, but contribute to a better tactical and strategic decision-making from existing knowledge in real time. The knowledge management function would finally be able to create a competitive advantage.

THE EVOLUTION FROM CHATBOTS, VIA RAG TO AI AGENTS

Most enterprises have started their AI journey with basic chat-based assistants, which offer a convenient interface for users to interact with LLMs. As LLMs are trained on public information from the internet, they cannot deal very well with proprietary company knowledge. However, they can quite easily be extended to pull information from internal sources through Retrieval-Augmented Generation (RAG). The rapid evolution of LLMs and Generative AI (GenAI) has led to broad adoption of chatbots, allowing employees to quickly produce content, and get help with simple day-to-day questions for which existing Enterprise Search systems have been failing for decades.

Although such RAG chatbots can be prototyped quickly, with open source frameworks like LangChain or LlamaIndex and off-the-shelf AI models, in the end they often fall short of the expectations of domain experts when it comes to tasks that require a precise understanding of domain knowledge. This is one of the key reasons why only around 30% of the AI prototypes end up going in production (Deloitte, 2024).

RAG, while powerful, is not a silver bullet; it relies heavily on the quality of the underlying retrieval system. For enterprises to trust their AI Agents, it's essential to have a well-tuned, domain-specific retrieval system foundation that understands the unique context and nuances of the organization's data, terminology and knowledge domains.

Imagine an analyst investigating emerging market trends. Unlike traditional RAG systems that generate brief summaries, an AI Research Agent divides the task into sub-goals, performs multiple searches, assesses the relevance of each source, synthesizes insights, and organizes key trends into a cohesive report. Rather than overwhelming the analyst with raw data, the agent distills actionable intelligence from internal and public sources, achieving in minutes what would normally take days.

DOMAIN-SPECIFIC AI USE CASES IN R&D, ENGINEERING, IP AND COMPLIANCE

Research and Development (R&D) departments often review literature to identify scientific trends and validate new innovations. AI Agents can streamline this process by summarizing findings from various studies, identifying gaps, and even proposing further research. But the applications of AI Agents in KM extends far beyond R&D. Here are some domain-specific use cases where AI agents are making a measurable impact:

- **Knowledge Extraction from Technical Manuals:** AI Agents can automate the extraction of essential information from complex technical manuals, e.g. creating customized maintenance schedules and troubleshooting guides based on actual maintenance logs for each machine. This reduces human error and minimizes downtime in production facilities.
- **Onboarding and Offboarding:** For onboarding, AI Agents organize and tailor training materials, ensuring new hires quickly access relevant information suited to their skill levels. During offboarding, these agents capture departing employees' knowledge, preserving expertise and maintaining continuity.
- **Compliance Analysis:** AI agents can review contracts, policies, and procedures with regulatory insight, flagging compliance issues and saving time for legal teams while ensuring adherence to evolving standards.
- **Intellectual Property (IP) Research:** AI agents streamline patent research by identifying relevant patents, comparing claims, and suggesting new opportunities, freeing IP professionals to focus on strategic analysis rather than exhaustive document searches.

THE RISE OF AI AGENT ECOSYSTEMS: INTEROPERABILITY AND RECENT INNOVATIONS

For AI agents to achieve their full potential, they need the ability to interact not only with humans but also with other agents, tools, and systems. Big Tech is setting the stage for agent ecosystems (Ghaffary, 2024), where multiple AI agents can work in tandem. Microsoft's "Copilot" vision, for example, provides for a marketplace of AI Agents, allowing seamless interaction with ubiquitous tools like MS Teams, Excel, and Word. Salesforce, with their new Agentforce product line, offers agents that pull data from across the Salesforce ecosystem to provide direct action on customer support and sales.

In these ecosystems, agents communicate with each other and understand and access APIs of existing business applications, which opens up completely new possibilities for KM workflows. With the introduction of agents that can understand and operate a computer screen such as the recent model launched by Anthropic, the IT integration puzzle also seems largely solved. We are already seeing significant growth in the AI agent startup ecosystem, with hundreds of new companies emerging in 2024, and this number is projected to expand to thousands by 2025 (Owyang, 2024).

While critics of AI Agents are quick to point out the analogy to early AI assistants like Microsoft's Clippy, which only offered very basic, and often intrusive help, today's AI agents have evolved significantly and the underlying LLM technology plays in a different league altogether. In 2024, these agents can operate seamlessly in the background, autonomously managing complex tasks, scheduling, information retrieval and API use. Eighty-two percent of business leaders surveyed said they expect to integrate AI Agents into their businesses in the next few years (Boulton, 2024).

EMBRACING THE FUTURE OF AI-ENHANCED KNOWLEDGE MANAGEMENT

AI agents represent a new frontier in knowledge management, transitioning from simple interactions to sophisticated, autonomous problem-solving for complex tasks. For enterprise leaders, the opportunity is clear: deploying domain-specific AI agents that leverage reliable retrieval systems and operate within dynamic ecosystems can significantly enhance team capabilities. Gartner advises IT leaders to "look for agentic AI in your technology stack" (Coshow, 2024). At the same time, the goal of AI implementation should not be to replace human judgement, but to augment the workforce with superpowers that boost productivity, and increase quality and business growth, in a world where the supply of experts and skilled labor is increasingly limited.

Zeta Alpha's work in this field demonstrates the transformative potential of these agents. With the Zeta Alpha RAG Agents SDK, the recently introduced AI Research Assistant, and many domain-specific applications across knowledge-intensive industries, we showcase how AI agents can be developed and deployed in production at scale. In this way, AI isn't just augmenting human knowledge management—it is redefining the possible in enterprise KM.

As enterprises embrace this shift, leaders will be faced with questions about implementation, integration, and reliability. However, those who take the leap will find that the rewards go beyond incremental improvements—they're unlocking true KM superpowers that will shape the future of their organizations. The evolution from chatbots to AI agents is here, and for forward-thinking leaders, the time to act is now.

Jakub Zavrel (zavrel@zeta-alpha.com) is Founder and CEO, Zeta Alpha.

References:

Clint Boulton, August 19, 2024. AI agents loom large as organizations pursue generative AI value, CIO Magazine. https://www.cio.com/article/3488553/ai-agents-loom-large-as-organizations-pursue-generative-ai-value.html

Tom Coshow, October 1, 2024. Intelligent Agents in AI Really Can Work Alone. Here's How. Gartner. https://www.gartner.com/en/articles/intelligent-agent-in-ai

Deloitte. August 2024, The State of Generative AI in the Enterprise - Moving from potential to performance. https://www2.deloitte.com/us/en/pages/consulting/articles/state-of-generative-ai-in-enterprise.html

Shirin Ghaffary, February 15, 2024. Tech Companies Bet the World is Ready for 'AI Agents'. Bloomberg. https://www.bloomberg.com/news/newsletters/2024-02-15/tech-companies-bet-the-world-is-ready-for-ai-agents

Jeremiah Owyang, September 12, 2024. Plot Your Path: The 2024 AI Agent Ecosystem Map. Web Strategist. https://web-strategist.com/blog/2024/09/12/plot-your-path-the-2024-ai-agent-ecosystem-map/

Kylie Robison, October 10, 2024. Agents are the future AI companies promise - and desperately need. The Verge. https://www.theverge.com/2024/10/10/24266333/ai-agents-assistants-openai-google-deepmind-bots

Steven Rosenbush, October 16, 2024. AI Agents Can Do More Than Answer Queries. That Raises a Few Questions. Wall Street Journal. https://www.wsj.com/articles/ai-agents-can-do-more-than-answer-queries-that-raises-a-few-questions-15009853

Decoding 2025: The Must-Read Report on Knowledge Management Trends

BY Bloomfire

Companies that embrace knowledge management continue to enjoy a competitive advantage. In 2025, KM is quickly becoming a business imperative. Companies that can harness, share, and scale their organizational knowledge will thrive.

Adapting to emerging trends is more important than ever. Based on Bloomfire's research into industry shifts, the 7 Knowledge Management Trends Shaping 2025 highlights key areas where companies must focus, including transforming data into a strategic asset, integrating AI-driven insights, building smarter tech ecosystems, and prioritizing employee experience.

THE 7 KNOWLEDGE MANAGEMENT TRENDS RESHAPING 2025: KEY INSIGHTS

These trends are not just industry predictions–they offer actionable strategies that help organizations drive innovation, collaboration, and sustainable growth.

- Corporate data will continue evolving from passive storage to an actively managed asset that fuels business value.
- Artificial intelligence in knowledge management will continue to expand, enabling smarter, faster decision-making.
- High-performing companies will prioritize data-cleaning technologies to eliminate ROT (redundant, outdated, trivial) data and maximize data-driven outcomes.
- The need for interconnected technologies will push IT departments into a strategic role in business decision-making.
- Enterprise users expect a seamless experience with work tools similar to consumer apps.

These insights provide a glimpse into the transformative potential of knowledge management in 2025. The full report offers deeper analysis and actionable strategies.

Read the Guide: The 7 Knowledge Management Trends Shaping 2025

FROM INSIGHT TO ACTION: TURNING KNOWLEDGE TRENDS INTO LEADERSHIP SUCCESS

Successfully navigating the evolving knowledge management landscape requires more than awareness—it demands action. Leaders must translate trends into strategic plans, making knowledge a force for business growth.

The following are some of the best practices discussed in the guide that appropriately respond to how you can utilize the trends for improving your company's knowledge management.

REIMAGINE KNOWLEDGE AS A STRATEGIC ASSET

Just as companies invest in financial and human capital, they must also manage knowledge capital. Knowledge isn't just information–it's an asset with real business value that, when actively managed, drives competitive advantage.

Consider how knowledge functions on a company's balance sheet. Corporate data, proprietary processes, and employee expertise contribute to operation efficiency, revenue growth, and innovation. But without proper governance, knowledge can lose its value–becoming outdated, irrelevant, or even a liability. Businesses should treat knowledge like other critical assets: maintaining its value through continuous investment, regular audits, and strategic use.

By shifting from passive data storage to active KM, companies can ensure that their knowledge fuels business decisions, creates efficiencies, and drives long-term growth. This approach transforms knowledge from an abstract concept into a measurable business asset.

INVEST IN CUTTING-EDGE KNOWLEDGE MANAGEMENT TECHNOLOGIES

A robust knowledge management system is no longer optional–it's imperative. AI, machine learning, and advanced analytics continue transforming how we capture, analyze, and utilize knowledge.

Embrace these technologies to automate routine tasks, identify hidden patterns, and extract actionable insights from your data. They can empower your teams to anticipate market trends and gain a competitive edge in decision-making.

PRIORITIZE DATA INTEGRITY AND KNOWLEDGE QUALITY

As with any business aspect, technology alone is not enough. The type and quality of knowledge you enroll and circulate in your system influence the value of data and the management system itself.

A polluted knowledge ecosystem can lead to flawed decisions, missed opportunities, and, ultimately, organizational failure. To avoid these consequences, implement robust data governance frameworks and invest in data cleansing tools to ensure your knowledge base is accurate, reliable, and up-to-date.

CENTRALIZE YOUR KNOWLEDGE SYSTEMS

A fragmented knowledge landscape, with information scattered across disparate platforms and silos, hinders efficiency and impedes innovation. A connected tech ecosystem allows knowledge to flow freely across departments and teams. Employees can access the information they need, regardless of

source, and foster a holistic understanding of the organization's knowledge assets.

Eliminate information silos by building an interconnected system where knowledge is easily accessible and actionable.

CHAMPION A CULTURE OF KNOWLEDGE SHARING

A thriving knowledge-sharing culture is the key to materializing the full potential of this technological ecosystem–and leaders must set the tone. This transcends departmental silos and embraces cross-functional collaboration.

Lead by example and encourage a knowledge-sharing culture where expertise flows freely, and insights drive success. Foster an environment where knowledge is valued, shared freely, and recognized as a key driver of organizational success. To do this, you may need a shift in mindset, moving away from knowledge hoarding to knowledge sharing as the norm.

BLOOMFIRE: YOUR PARTNER IN EMPOWERING KNOWLEDGE-DRIVEN ENTERPRISES

Leading the way in 2025 means turning these trends into action. Whether creating a smarter tech ecosystem or unlocking your organization's knowledge value, forward-thinking companies need solutions that empower employees, encourage collaboration, and smeasure success. Explore how Bloomfire's platform can support your KM journey. ▮

Zeta Alpha

BUILD INTELLIGENCE AND DRIVE RESULTS

Advanced AI solutions for the enterprise

Zeta Alpha builds high quality custom AI search and RAG agents for knowledge-intensive companies. Turn your own unique expertise and R&D into the power behind your AI transformation.

Deliver results faster with Zeta Alpha's secure and versatile AI platform. Achieve more reliable outcomes using AI search models finetuned by our expert team. Support the reuse of existing knowledge with user-friendly AI assistants based on your own business data and workflows. Empower teams with modern tools to discover and organize knowledge.

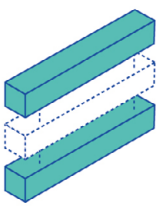

RAG AND AI AGENTS

A secure multi-tenant search foundation for all your AI use cases

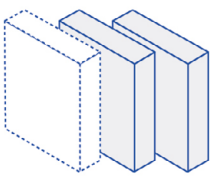

DISCOVER AND ORGANIZE

Get answers, chat, summarize, analyze and visualize your own data

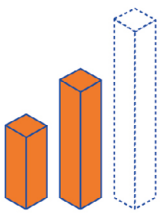

HIGH QUALITY ANSWERS AND SEARCH RESULTS

Customized RAG, finetuned embeddings, optimized ranking

FOR WHOM

Generative AI teams

Move from prototype to production faster. Serve many RAG, AI Agent and LLM use-cases in a unified and scalable multi-tenant platform.

Enterprise knowledge management

Connect all your internal and external data sources in a single semantic search engine. The best tools to discover and organize knowledge, and to build your next Gen AI capabilities.

Research teams, analysts and decision makers

Chat with your own document collections. Use our secure private Generative AI Assistant to help you make better decisions and get work done faster.

GET STARTED WITH ZETA ALPHA

www.zeta-alpha.com