

The background of the slide is a composite image of space. On the right side, a large portion of Jupiter is visible, showing its characteristic bands of white, orange, and brown clouds. On the left side, there is a vibrant nebula with swirling clouds of gas in shades of yellow, orange, and red, punctuated by numerous bright stars of varying sizes.

The Jupiter Release

Galileo® Goes Enterprise

galileo®
the joshbersin company


Overview

Imagine if every employee, manager, and leader in your company could consult with the world's most knowledgeable HR expert, one who understands global practices and benchmarks, and can be accessed alongside your company's data to solve your most pressing challenges.

Welcome to the Jupiter release of Galileo®, the world's most intelligent AI agent for HR and management.

Introduced in 2023, Galileo is an AI agent built on 30+ years of research, benchmarking, and industry case studies spanning all areas of HR, management, and workforce technology. Today, Galileo is used by more than 1,000 organizations worldwide.

The Jupiter release is the most significant advance in Galileo's history—making it more accurate, open, and deeply connected to your enterprise data and systems than ever before, delivering trusted HR and management intelligence wherever your people work.



Major New Features

- 1 Galileo Jupiter uses the new, highly optimized Agent Ready Corpus (ARC),** making Galileo faster and more accurate than Claude, ChatGPT, Gemini, and other frontier solutions for HR, management, and regulatory questions, with every answer grounded in cited, traceable sources.

We benchmarked our 30 “golden prompts”—a set of comprehensive, real-world HR questions designed to simulate the most demanding queries our clients bring us—against these tools and found that Galileo is more than 200% faster, while these systems generate hallucinations and incorrect answers in more than 50% of the benchmarks. More details follow.

- 2 Galileo runs on any AI agent, available as an MCP server.** This “Galileo everywhere” strategy means you can plug Galileo’s HR and management intelligence into any system and access the trusted expertise of The Josh Bersin Company (JBC) alongside employee, team, or other enterprise data for a wide range of applications. Galileo provides universal, enterprise-grade connectivity to any AI agent and is certified ISO 27001 and SOC2 type 2 compliant.

- 3 Galileo Jupiter is specially optimized for use in Microsoft Copilot,** able to access any information in a company’s Microsoft Graph, including SharePoint, Outlook, Office, People Skills, Viva, and other Microsoft-authored content.

This lets Copilot users ask and answer questions about leadership, management, goal-setting, or any other employee topic with rigorous answers, benchmarks, and examples. Companies can also build custom Galileo-powered agents directly from the Copilot tool set.

- 4 Galileo Jupiter is available today on Sana and through MCP on Claude, Gemini, ChatGPT, Glean, and any other MCP-ready platform—with embedded integrations into Workday Sana Enterprise, Microsoft Copilot, ServiceNow, HiBob, and Gloat arriving with those partners through 2026.**

This lets companies easily deploy Galileo across their entire organization under a simple licensing model with no confusing token fees.

- 5 Galileo Jupiter is engineered to bring HR expertise and benchmarks to enterprise HCM workflows.** When a customer enables the integration, Galileo calls the host platform’s published Workday, SAP SuccessFactors, and Oracle HCM APIs as the signed-in user to retrieve only the data that those endpoints expose to that user. It then benchmarks that data against Galileo’s research corpus. Deployed as an enterprisewide solution, Galileo gives all users access to that guidance without requiring them to leave their platform.

A Closer Look: Key Questions and Answers

What Are the Business Benefits of Galileo Jupiter? 5

.....

What Is Galileo ARC and How Does It Work? 6

.....

Why Does This Architecture Work Exceptionally Well for Galileo Content? 8

.....

**How Is Galileo Jupiter More Accurate, Actionable,
and Reliable Than AI Models on the Open Web?** 9

.....

**How Does Galileo Jupiter Perform against ChatGPT
and Claude on Real-World HR Benchmarks?** 10

.....

Summary Comparison: Benefits of Galileo Jupiter 12

.....

What Are the Business Benefits of Galileo Jupiter?

Single Pane of Glass for Every Employee

When Galileo is embedded in the tools employees, managers, and HR use daily, every user gains a “single pane of glass” for actionable information and advice about their jobs, careers, goals, or management issues. The following examples illustrate what this means in practice across a range of roles and industries.

A sales manager may ask to compare their team’s payroll and benefits cost against their budget, then benchmark that expense against pay data in Galileo’s benchmark database. They may then ask Galileo for advice and best practices about how to better manage bonuses and rewards.

A nursing supervisor may use Workday Decision Intelligence to analyze the skills of their 25 nurses, drawing on Workday skills, learning, and activity data, and leverage Galileo’s expertise and industry-standard nursing skills models to produce a custom annual development plan for each nursing role. They may also ask Galileo for best practices on how to deliver training to meet those needs.

An HR leader may ask Galileo to provide the framework for succession, the skills and experience criteria that define readiness for a critical role, and recommend what each candidate’s development plan should include.

An HR business partner may take their team’s engagement and retention data, compare it against Galileo benchmarks, and use Galileo to recommend research-based interventions and a custom engagement plan for the team and each individual.

An individual using Microsoft Copilot may build an agent with Galileo that reads their emails, analyzes their meetings, and delivers a weekly assessment of strengths and development areas—grounded in Galileo’s corpus and the company’s skills model.

We have mapped more than 400 use cases, many of which are almost impossible or take weeks of effort in today’s systems.

You can also embed Galileo within your current employee chatbot so it simply “makes everything smarter.”

Global Benchmarks and Trusted Data

Galileo Jupiter maintains global benchmarks and data about roles, skills, salaries, regulatory rules, and turnover—drawing on trusted content from Lightcast, SHL, Oyster, Visier, Equilar, and others. This includes global employment rules and practices, leadership models, global employment rules and practices, leadership models, as well as maturity models from JBC. Galileo also has access to 500+ case studies and client examples so users can learn about technology solutions for their HR or management needs. New content for global salary benchmarks and senior C-level pay is coming later in 2026 and 2027, and Galileo’s corpus is maintained on a regular release cadence.

Efficient by Design (No Expensive Tokens)

Unlike most LLM-based solutions, Galileo is highly efficient: compact ARC retrieval keeps token consumption to a small fraction of conventional RAG (retrieval-augmented generation). We license Galileo through standard enterprise pricing, and regardless of the platform you use today—MS Copilot, Teams, Sana, or another agent—Galileo can plug in and deliver its research corpus at minimal token cost.

Easy to Deploy

Any nontechnical professional can turn on Galileo in platforms like Sana or Microsoft Copilot without any engineering. Galileo operates within the host platform’s security model, so for the APIs and skills it calls, Galileo sees only what the signed-in user is already entitled to see. No separate permissioning of Galileo is required.

What Is Galileo ARC and How Does It Work?

The Technical Foundation

Over our three+ years of engineering, we have learned how to make Galileo portable, highly efficient, and extremely accurate and effective. Our proprietary ARC engineering is the foundation.

AI agents use LLMs and RAG to index and analyze documents (e.g., policies, emails, videos, etc.) so the LLM can answer questions, analyze information, make decisions, and support users.

The indexing process uses neural network technology to build “embeddings,” which are long lists of numbers that capture what a piece of text means, and position them in a shared semantic space. A question is converted the same way, and the system retrieves the passages whose meaning sits closest to it.

While this works reasonably well for small data sets, embeddings become confused and “blurry” as the volume of data grows.¹ Give an LLM a 200-page document and it may draw on the first few pages for an answer, never reaching commentary later that further explains or even contradicts what was stated earlier.

The LLM also has to guess what words mean. If the word “performance” is used in a *Car and Driver* article referring to the speed of a car and then in a *Wired Magazine* article to refer to PC performance, the LLM must resolve that ambiguity—consuming computing cycles and tokens, with results that can be hallucinations.

A recent study by the BBC found that 45% of all queries about news result in errors²—erroneous information, hallucinations, or just incorrect numbers or dates. These incorrect results also burn valuable and expensive tokens as they appear.

In Jupiter we use world-class techniques to “chunk” our entire database into smaller representations, tagged with consistent metadata developed by our analysts, semantic tags capturing what words and phrases mean, and authoritative statements of intent to direct the LLM toward accurate answers. Every chart, graphic, maturity model, case study, and podcast is chunked and tagged in the same way, so answers cite accurate examples and correct visuals.

This combination of semantic and contextual chunking, coupled with extensive tagging and semantic intelligence, is what makes ARC so effective. Rather than guessing what something means, Galileo ensures the agent LLM understands the topic, industry, user type, and intent behind every piece of content it uses to answer a query.

With ARC, Galileo serves as a delegation layer that calls the published APIs and skills of Workday, SAP SuccessFactors, and Oracle HCM as the signed-in user, retrieves the information that those endpoints expose to that user, and presents it alongside its benchmarks and expertise.

Because Galileo calls these APIs and skills as the signed-in user, it inherits that user’s security rules, privileges, and access rights. If an employee or manager asks Galileo for salary information, it surfaces only what those endpoints already permit them to see.

1 [“A Systematic Investigation of Document Chunking Strategies and Embedding Sensitivity,”](#) arxiv.org, March 2026, and [“Chunking Strategies to Improve LLM RAG Pipeline Performance,”](#) Femke Plantinga and Victoria Slocum/weaviate.io, September, 4, 2025.

2 [“Largest study of its kind shows AI assistants misrepresent news content 45% of the time – regardless of language or territory,”](#) BBC Media Centre, October 22, 2025.

How ARC Chunking Works

When you chunk a large document into smaller pieces, each piece gets its own embedding—its own focused vector representation. Embedding an entire document as a single block produces a single vector that is an average of everything inside it.


A fifty-page report covering compensation, leadership development, and AI governance, collapsed into one embedding, yields a generic representation that does not strongly match any specific query—producing hallucinations, errors, and high token usage.

Chunking gives each piece a sharper, more focused semantic signature. When someone asks a specific question, the LLM retrieves the exact paragraph or section that's most relevant, rather than searching a whole document for an answer buried somewhere inside.

This delivers three concrete benefits:

1. **Precision**, feeding the model only relevant information rather than noise.
2. **Scalability**, as you can't load an entire corpus into a prompt.
3. **Accuracy**, ensuring a level-playing field in which the LLM sources and identifies the best content based purely on relevance—regardless of how long a document is or where the relevant information appears within it.

There's also a subtler but important effect: with Jupiter's intelligent chunking, irrelevant or contradictory chunks are far less likely to get pulled in alongside relevant ones. Mixed or conflicting context is one of the primary drivers of hallucination, so cleaner retrieval directly reduces that risk.



Why Does This Architecture Work Exceptionally Well for Galileo Content?

A Domain-Bound and Structured Corpus

Chunking and tagging are standard practice in modern AI systems. What is not standard is the knowledge system behind ours. Every chunk in ARC is labeled across six dimensions — topic domain, content type, semantic role, evidence type, audience applicability, and concepts—using a controlled vocabulary our analysts built from thirty years of HR practice, anchored to a proprietary ontology of HR concepts, maturity stages, roles, industries, and company sizes.

Each chunk is additionally enriched with the questions it answers, so retrieval matches the way HR professionals actually ask. This is not English being indexed; it is HR knowledge being represented — and it is why a query about skills-based organizations retrieves our point of view, not a lexical coincidence.

Galileo's corpus is also domain-bounded and consistent. Every chunk is written with the same terminology, the same level of rigor, and reflects The Josh Bersin Company's point of view. When someone asks about skills-based organizations, there's no competing noise from unrelated fields—every relevant chunk in that space is genuinely relevant.

Metadata tagging adds another layer on top. Because every chunk carries consistent tags—report title, publication date, topic taxonomy, author—the system can filter before it performs a semantic search. So instead of asking the whole vector space “find me anything about compensation,” you can ask, “find me salary and hiring information about nurses in emergency room roles located in the Midwest,” and the LLM can traverse the tags to find it immediately.

There's also a subtler benefit: consistency. A generic corpus is written by thousands of different people with different structures, styles, and opinions. Galileo's corpus follows repeatable formatting, consistent headers, consistent report structures, and an authoritative voice, so Galileo can be tuned to reliably produce clean, coherent answers every time.

How Is Galileo Jupiter More Accurate, Actionable, and Reliable Than AI Models on the Open Web?

To understand the difference, let's walk through what happens when an AI system uses the open internet as a corpus without Jupiter's optimizations.



On accuracy

Without ARC a model either embeds entire documents as single vectors or dumps raw unstructured text straight into the context window. Either way, it loses the ability to point to a precise passage as evidence. It works from a blurry, averaged sense of what a document is roughly about rather than an exact grounded statement. When it generates an answer, it fills gaps with general knowledge or pattern completion instead of drawing from the actual source, and that's exactly the condition that produces hallucinations.



On actionability

Retrieval against web content means every similarity search runs against huge, semantically muddy vectors. A 10,000-word web page covering five different subtopics produces one washed-out embedding that neither strongly matches nor strongly rejects most queries. The top results become mediocre matches across the board rather than a few excellent, tightly relevant ones, and the model reasons over irrelevant filler mixed in with the useful content, degrading answer quality.



On reliability

Because ARC is optimized for LLMs, it can be dropped into various products, context windows, and retrieval pipelines without loss of fidelity. Open internet content, by contrast, is stuck in its original format—full of ads, navigation text, inconsistent structure, wildly different lengths—making it brittle, hard to move between systems, and hard to reuse beyond the exact context it was originally scraped from.

How Does Galileo Jupiter Perform against ChatGPT and Claude on Real-World HR Benchmarks?

Here's the bottom line. While there is no perfect answer to a question about management, leadership, or HR, we do have benchmarks. As part of our effort to test the Galileo ARC corpus, we developed 30 "golden prompts." These prompts simulate the real-world, comprehensive questions our clients ask us directly. (They fall into all areas of HR, including recruitment, training, vendor analysis, turnover benchmarking, salary benchmarking, leadership development, and HR organizational design.)

We use an AI "grader" to evaluate these 30 golden prompts on a five-point scale as we test Galileo against multiple agent LLMs.

When we ran the 30 golden prompts against public chatbots without any access to licensed JBC or partner research, more than 50 hallucinations, exaggerations, and outright fabrications appeared. When the identical models were grounded in ARC, those fabrications disappeared — replaced by cited, traceable answers. These are the same frontier models Galileo itself runs on; the difference is not the model, but the grounding.

Examples include:



ChatGPT

- Invented pricing for vendor products
- Fabricated incorrect salary information without attribution
- Invented unsourced statistics for turnover rates by industry
- Cited McKinsey as the source for maturity models developed by JBC
- Fabricated survey data on employee engagement, time to hire, and other metrics



Claude

- Hallucinated incorrect answers about recruiting CRM
- Fabricated salary data without attribution
- Invented statistics about JBC proprietary talent density research
- Speculated about JBC position on 9-box grids
- Reproduced JBC maturity models without attribution
- Attributed real proprietary JBC research to non-JBC consultants

(continued)

These results have three implications.

- 1 These systems are unreliable for important human capital decisions, benchmarks, and vendor data. They reproduce proprietary frameworks without attribution and cite incorrect recommendations—leaving users with no way to verify or defend an answer.
- 2 Such unreliable results create risk. Hallucinated answers may encourage managers or HR professionals to buy the wrong vendor product, underestimate the impact of a particular benefit or intervention, or even budget and plan incorrectly for a major transformation project. By contrast, the Galileo corpus is based on 500+ real-world interviews and millions of survey responses from validated companies. It is industry-specific, validated, and attributed. Much of what we observed has no source, no industry perspective, and no depth.
- 3 The lack of attribution in these systems leaves users with no ability to hold themselves or their companies accountable for AI-driven decisions. While many HR decisions involve judgment calls, if someone in a bank or large retailer asks, “why did we buy product X” or “why did we decide to pay someone at this level” and the HR professional has no source, both their job and their company is at risk.

This is why we focus so heavily on the word “trust” when we talk about Galileo Jupiter. All the decisions, advice, benchmarking, and education provided by Galileo are reliable, authenticated, and backed up by our company, our consultants, and the hundreds of real-world examples in our research.


Summary Comparison: Benefits of Galileo Jupiter

Benefit	Off-the-Shelf Agent	Galileo Jupiter
Freedom from vendor lock-in	Using one agent's internal or external corpus locks you their content and token costs.	Galileo frees you from any possible limitations in an AI agent's results, giving you an intelligent expert to validate information from others.
Trusted intelligence for decision-making	Lack of attribution, hallucination, and fabrications leave important questions with questionable answers.	All answers are attributed to JBC research, real world case studies, and validated research.
Precision of answers	Washed-out, averaged embeddings from long, multitopic pages produce generic rather than specific answers.	Sharp, focused embeddings per chunk produce answers specific to each user's industry and company scenario.
Retrieval relevance	High false-positive rate—unrelated content sits near relevant content in vector space, causing incorrect or fabricated answers.	Domain-bounded vector space ensures nearly everything retrieved is genuinely relevant, producing specific, attributed answers.
Hallucination risk	Model fills gaps from blurry, unattributed context.	Model is grounded in a specific, traceable passage and viewpoint.
Traceability and citation	No consistent source or authorship structure; many attributions incorrect or missing.	Every chunk traces back to a specific report, author, and section.
Metadata filtering	Little to no reliable metadata; can't prefilter before search.	Consistent tags (topic, date, author, content type) narrow the pool before similarity scoring.

Summary Comparison: Benefits of Galileo Jupiter


Benefit	Off-the-Shelf Agent	Galileo Jupiter
Chunking consistency	Inconsistent formatting across millions of authors—chunk boundaries are guessed and often wrong.	Consistent report structure lets chunking be tuned and reliable.
Contradictory content	Common—conflicting or outdated claims retrieved together.	Timeliness and authority tagging lets retrieval prefer current thinking.
Efficiency (cost/speed)	Retrieves whole documents or truncates arbitrarily, wasting tokens and increasing cost.	Retrieves only the relevant few hundred tokens, creating an order-of-magnitude reduction in tokens per answer—with no per-token agent fees.
Portability/reuse	Brittle—tied to original messy format, including ads, navigation text, and inconsistent length.	Modular, self-contained chunks reusable across products and pipelines.
Authoring quality	Written for human narrative flow, not machine consumption—LLM must guess author intent with no way of validating experience or industry knowledge.	Authored with claims up front, explicit references, and locally defined terms, reflecting JBC's authoritative position on all major topics.
Case study integration	Narrative content isn't linked to conceptual or topical content. No in-depth case analysis.	Case studies dual-indexed and cross-tagged with the same taxonomy as analytical reports, with hundreds of industry case studies throughout.
Evaluation and testability	Hard to build a reliable test set—no ground truth structure.	Structured corpus designed for systematic evaluation—validated through the 30 golden prompts suite.

The Galileo® Story So Far



Galileo Jupiter is available immediately via subscription.
Click [here](#) for more details.

About the Authors



Josh Bersin

Josh founded Bersin & Associates in 2001 to provide research and advisory services focused on corporate learning. He expanded the company's coverage to encompass HR, talent management, talent acquisition, and leadership and became a recognized expert in the talent market. Josh sold the company to Deloitte in 2012 and was a partner in Bersin by Deloitte up until 2018.

In 2019, Josh founded the Josh Bersin Academy, a professional development academy that quickly became the *"home for HR."* In 2020, he put together a team of analysts and advisors who are now working with him to support and guide HR organizations from around the world under the umbrella of The Josh Bersin Company.

Josh launched Galileo®, the world's first AI-powered expert assistant for HR, in 2024, followed by Galileo Learn®, the hyperpersonalized professional development solution and successor to the Josh Bersin Academy, in 2025. He is frequently featured in publications such as Forbes, Harvard Business Review, HR Executive, The Wall Street Journal, and CLO Magazine. He is a popular blogger and has more than 900,000 followers on LinkedIn.



Amy Farner

Amy is the executive vice president and head of product for The Josh Bersin Company. In this role, she helps align the company's solutions to HR leaders' most critical needs. Amy has more than 20 years of experience developing insights and solutions that drive more effective HR and business management amid a rapidly changing environment. Prior to joining The Josh Bersin Company, she was a leader in Deloitte Consulting's survey research and analytics practice and developed and launched a range of innovative HR solutions as part of CEB's Corporate Leadership Council (now Gartner for HR). Amy holds a BA from Wesleyan University's College of Social Studies and lives in Phoenix, AZ.



Ram Subramanian

Ram is head of AI and ecosystems at The Josh Bersin Company, bringing Galileo's HR intelligence into enterprise HR platforms, AI agents, and the partner ecosystems. He works with customers and technology partners globally to design agentic architectures that combine orchestration, context, and governance to deliver JBC's research to leaders, managers, and employees. His focus spans corpus engineering, AI agent evaluation, and intelligence layers that integrate with HR systems to make research-grounded reasoning governable and trustworthy within agent ecosystems. Ram brings more than 30 years of experience in enterprise software, consulting, AI, and product engineering, with leadership roles at Ernst & Young, Bank of America, Tata Consultancy Services, and Opkey. He holds a master's degree in computer science and lives in Chicago.

The Josh Bersin Company Membership

The Josh Bersin Company provides a wide range of research, tools, and advisory services to help HR leaders and professionals address the ever-evolving needs and challenges of today's workforce. We cover all topics in HR, HR technology, talent, and corporate learning, with a special focus on the professional development of HR teams.

Our corporate membership program provides HR leaders and teams with the skills, strategies, benchmarks, and insights to build cutting-edge HR and people strategies through research, assessments, professional development, exclusive events, and community. Corporate membership also includes access to Galileo®, the world's first AI-powered expert assistant specifically developed for HR. Trained on 25 years of The Josh Bersin Company's research, insights, and expertise and enriched by carefully curated material from our trusted content partners, Galileo unlocks information from over 50,000 verified assets to answer any HR-related question with timely and meaningful answers. Galileo Suite also includes Galileo Learn®, which provides hyperpersonalized, up-to-the-minute learning in the most dynamic, interactive HR professional development solution available today. Built on the legacy of Josh Bersin Academy—and its mission to build skills and capabilities through the power of shared stories—Galileo Learn brings advanced AI to the world's most comprehensive and ever-growing set of business insights, best practices, industry trends, and case studies.

To license Galileo Suite for individuals and small teams, please go to getgalileo.ai, and visit joshbersin.com/membership for more information about our corporate membership featuring Galileo and advisory support from Josh Bersin Company experts.



**The most trusted human capital
advisors in the world ■**