Flightplan: Getting from Enterprise Search to Cognitive Intelligence

Lucidworks

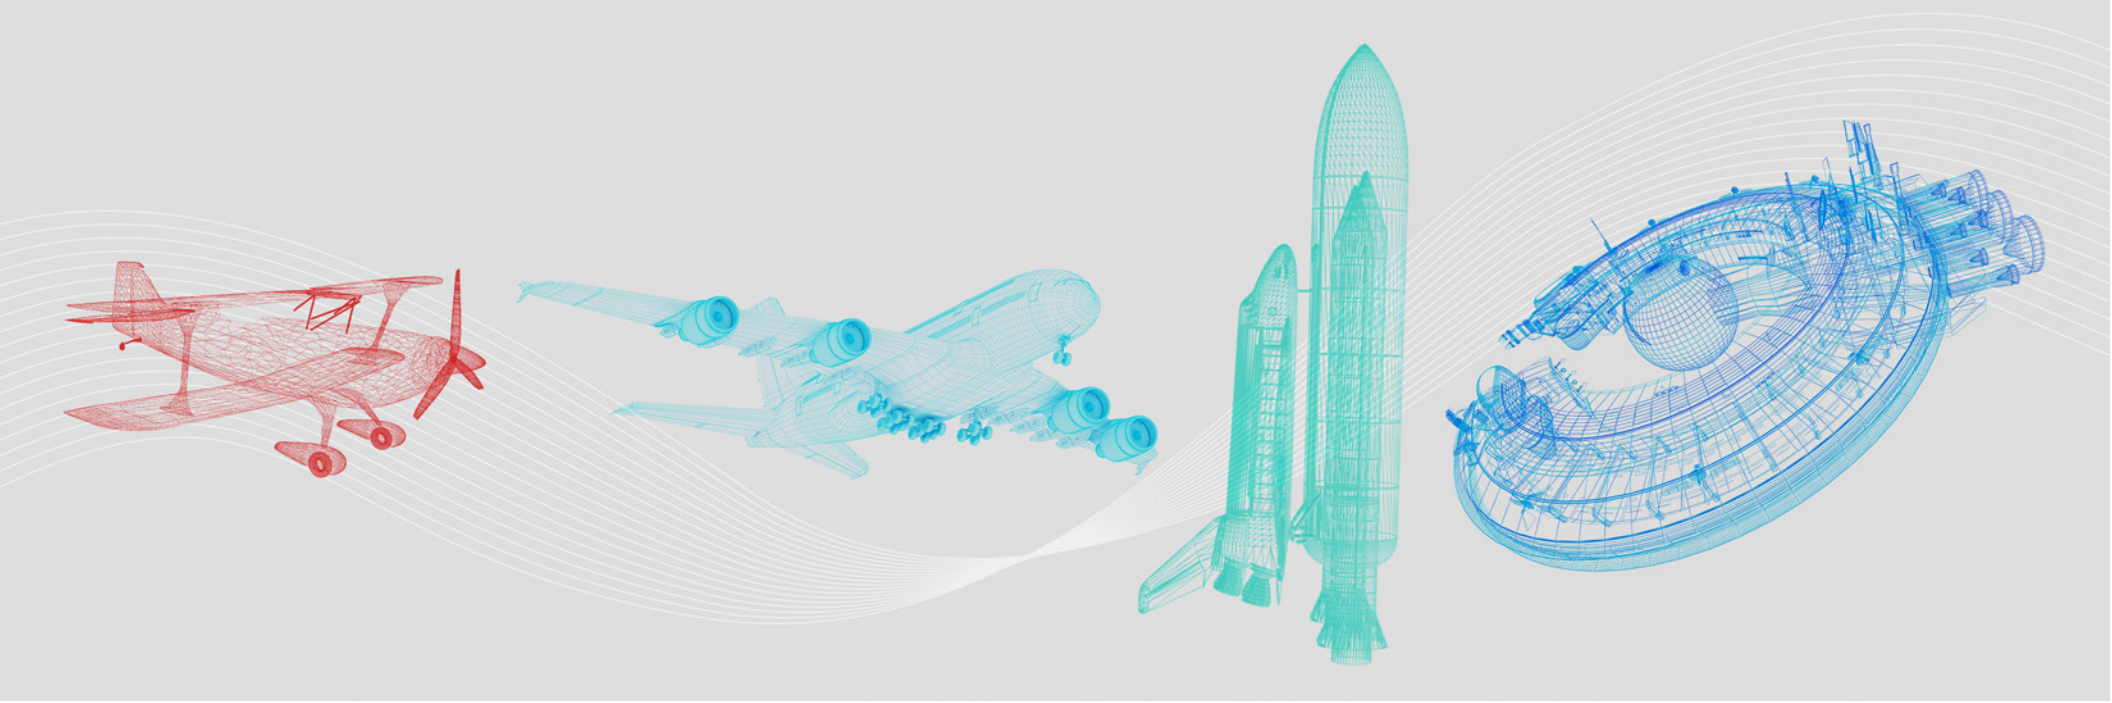

Introduction

You walk into your office and say “What are the sales projections for this quarter.” A small glowing orb talks back, “Do you prefer revenue or profit projections and would you like that broken down by territory?”

Chances are you’re nowhere near there. Chances are your phone is closer to organizing your personal life than your IT department is to organizing your business life (and the UI on your phone is lightyears ahead of what’s on your laptop).

Frankly, the voice interface thing is cool, but a bit of a gimmick. The real power in the above example is the intelligent technology that was able to hear that question in plain English and then respond with clarifying questions. Most organizations want that level of sophistication in their applications, but it remains elusive.

The rise of big data solved the problems of scaling and storing. Search evolved right alongside and become the best way for companies to make sense of the massive amounts of information they had. But the build-and-launch days of search apps are done. Your end users and customers want more. They require more. You don’t just need more data or better data, but better search to find the things you already have.

There is (of course) a shiny new name for this: “cognitive search.” Cognitive search means applications with smarter search that “gets” the context of the data, and “gets” you.

New Search (Same As the Old Search)

Cognitive search marks a sea change in how we think about search and big data. Instead of databases and keywords, we now have questions and answers and accompanying analysis. Think of it as Siri or Alexa for your business.

With the maturing of technologies like machine learning, artificial intelligence, and natural language processing, the opportunity for truly intelligent search and data applications is here and ready for deployment. But when you put aside the buzz and the hype cycle and all the market-y speak, you still need to know where you're going and how to get there.

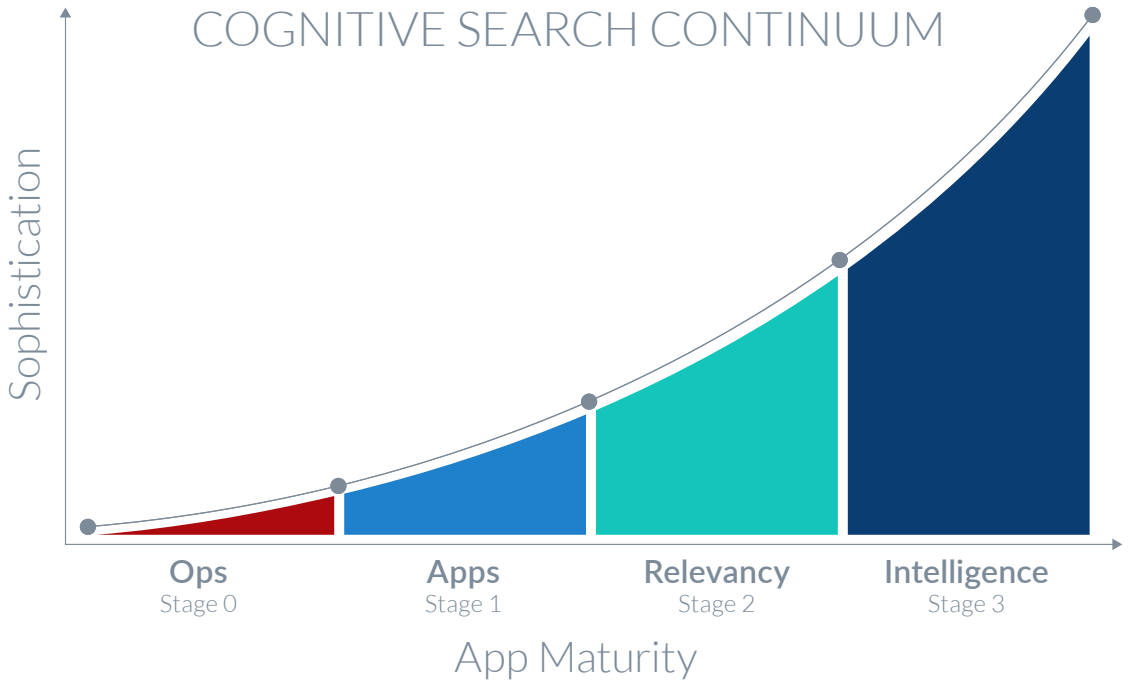

A Maturity Model for Search

Lucidworks developed the Cognitive Search Continuum to help companies understand where they're going, and the steps that need to be taken to achieve better search. The Cognitive Search Continuum is a four part model showing how organizations can start to evolve their search capabilities from simple keyword search applications to more powerful capabilities like cognitive search and conversational apps.

Don't worry, it still goes "up and to the right."

The horizontal axis shows the maturity of your search applications over time. The vertical axis charts the increasing technical sophistication of your search and big data initiatives.

Lucidworks has worked with some of the world's largest companies over the years. In our work, we've identified some common milestones and methodologies that define the growth and expansion of search capabilites at an organization.

Since each phase builds on the technology, expertise, and results gained in the previous one, it's difficult to skip phases.

Being in a particular phase is neither good nor bad. They simply show the state of an organization's search initiative and where it is going. Short of some sort of disastrous implementation or a mass exodus of staff or stakeholders, companies generally don't go backwards. Organizations usually stay where they're at—or start to move and expand. The speed of expansion is up to them and the resources they put into their search initiatives.

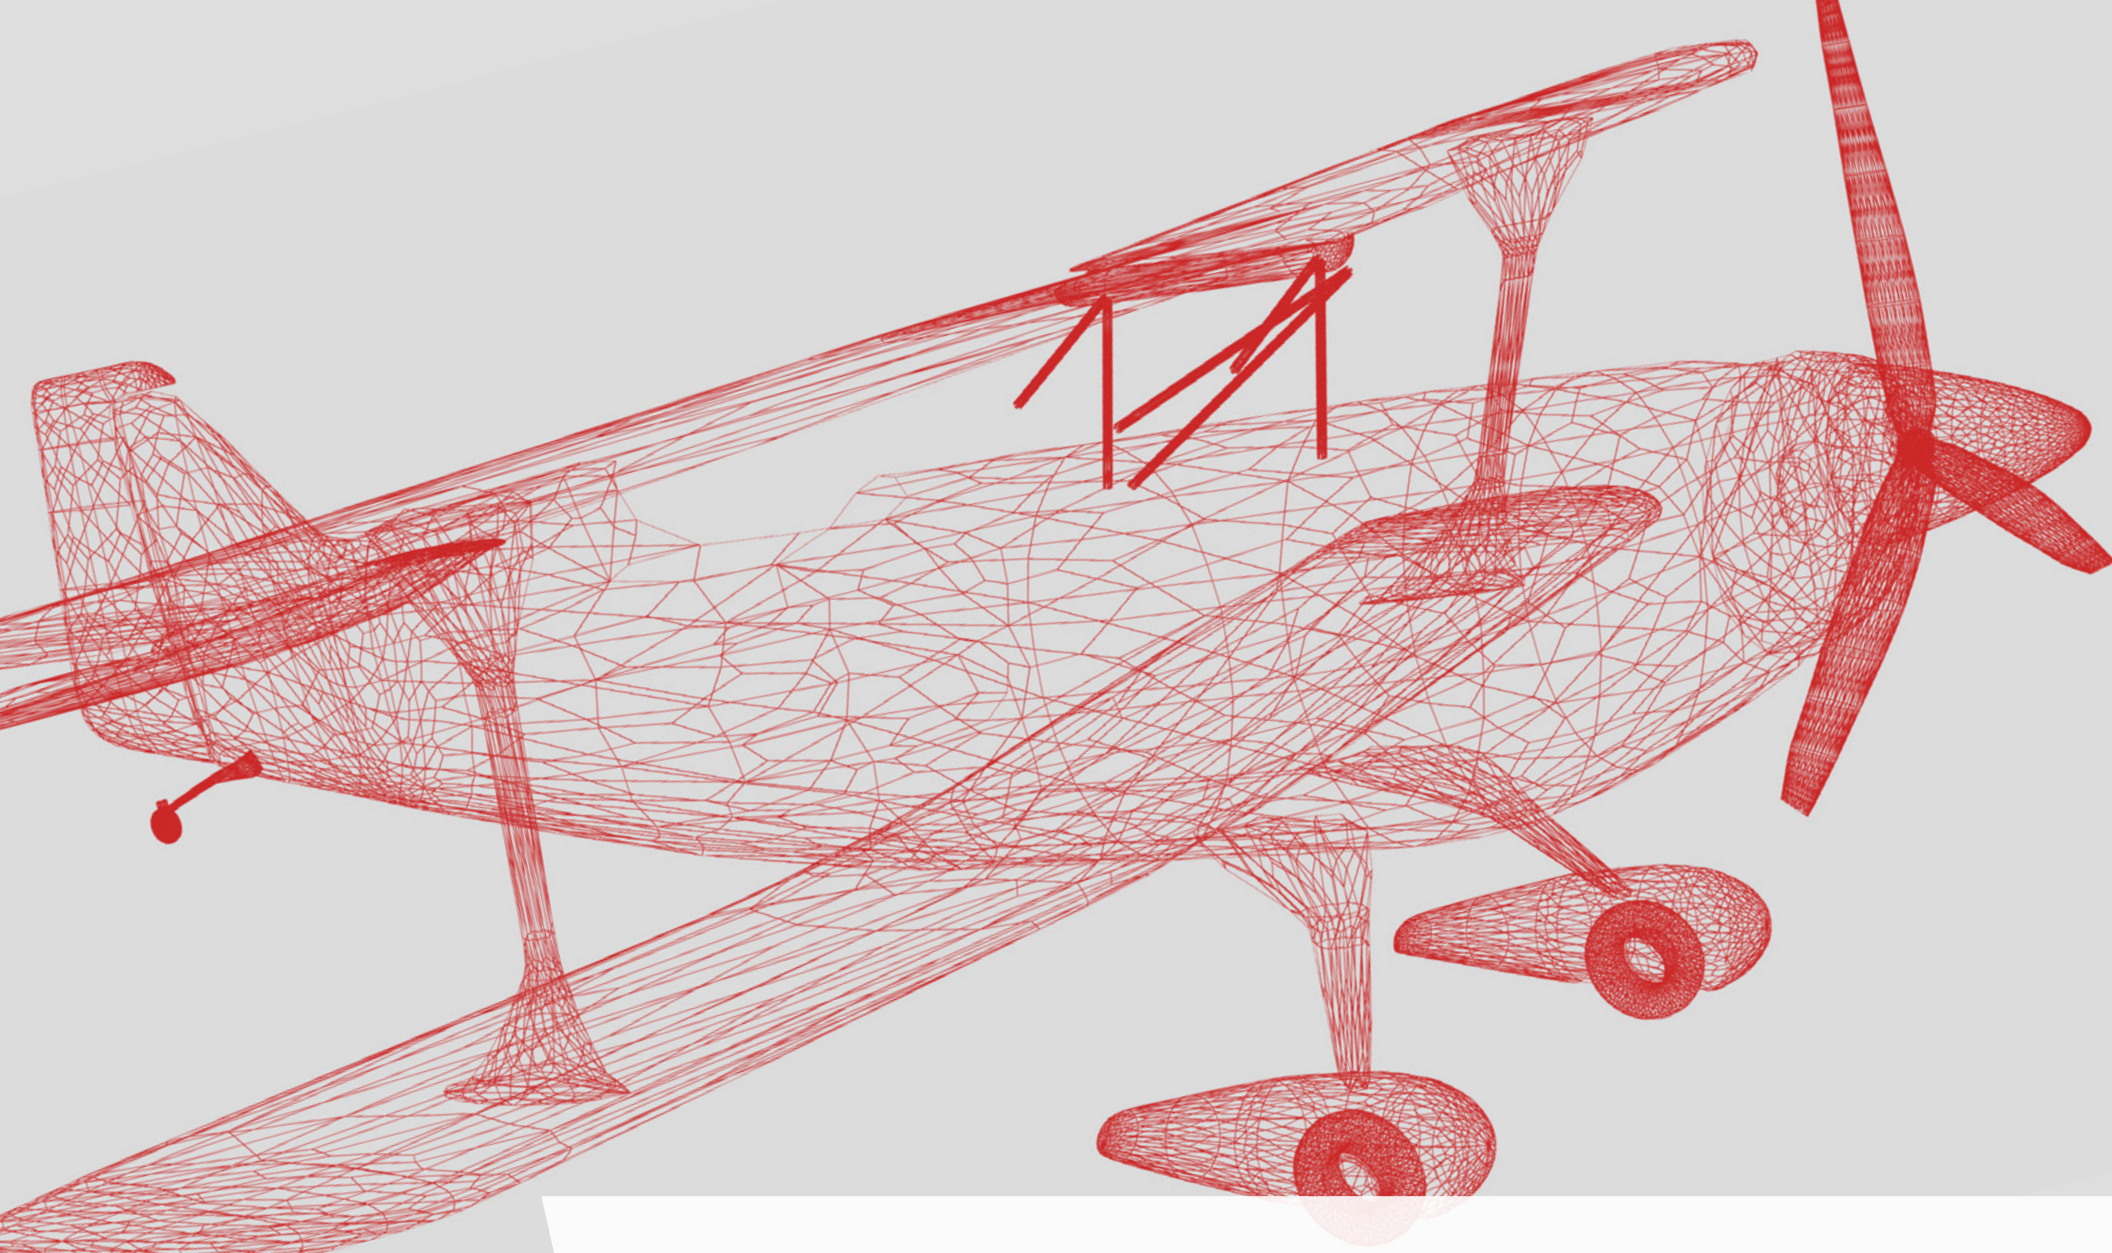

Stage 0
Ops

We don't start at one. We start at zero. The baseline. The inciting moment here is usually "We have all of this data. We can store it and scale it. But we can't search it!" Stage 0 of the continuum is Ops and encompasses the following milestones:

Stage 0 : Ops

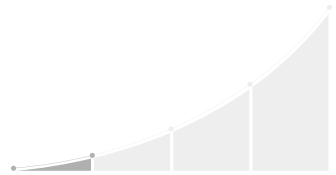

Milestones

You have selected a search platform. You've gone through procurement, purchased, built, tested a search platform and—after some internal gnashing of teeth—it's in production.

You're searching just a few data sources. The platform is probably indexing just a handful of collections of data. Maybe a SharePoint server over here or a Google Drive in the cloud. Or just the company intranet.

You have one generic search application. There is probably either one very generic application—like an internal search that employees use for the entire company—or one very super-specific application that searches a data source and is used by just a few users.

Long story short: **Getting data in, and getting data out.**

Metrics

Stage 0 is measured in the performance of the application and platform architecture including query volume, number of concurrent users, volume of data indexed, and overall size of the index.

Examples

- A risk management company has a simple search for their intranet.
- A textbook publisher offers search of all of their books.
- HR officers can search the HRMS for hiring and policy information.

What's Next

As the application gets used by more end users, gaps start to show. One app doesn't fit all. You're going to need more apps if you want to serve more end users across the organization.

Warning

Don't make the data lake mistake! If you build it, they won't come. They don't know where the data is, what it means, and it has no context. Plus, not everybody is a data scientist. Dumping all the data in a big fat datastore and hoping someone sometime comes along and somehow makes sense of it is doomed to fail. Your existing systems have the context and security information. You need a way to aggregate information into intelligence but a big fat copy of the data on another unstructured filesystem isn't necessarily the way to go.

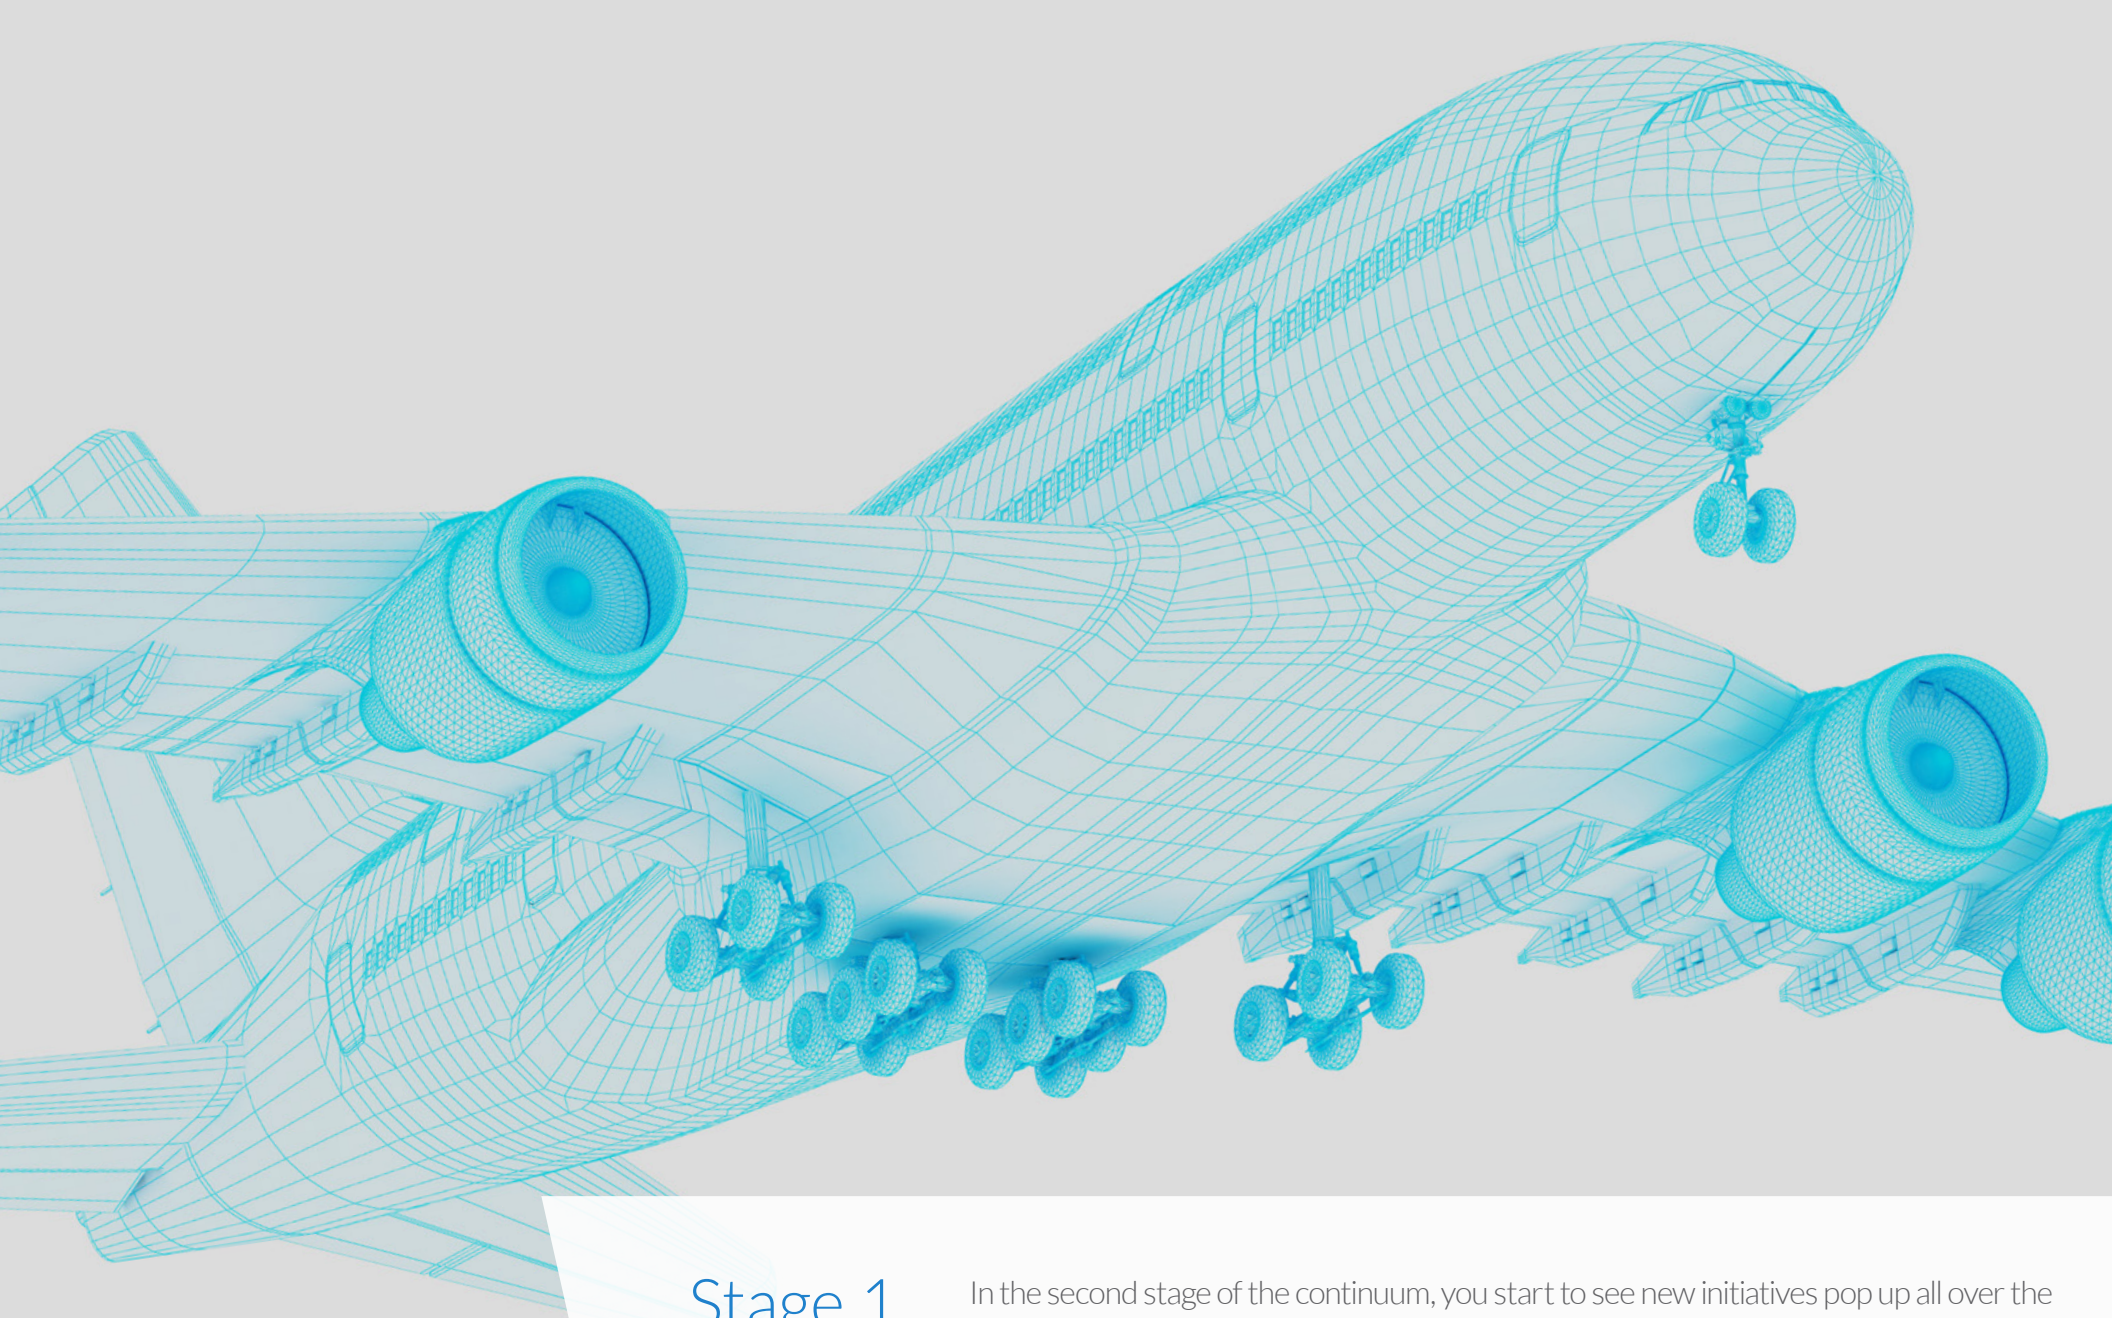

Stage 1
Apps

In the second stage of the continuum, you start to see new initiatives pop up all over the organization. Everyone starts asking when they can get their own apps for their teams that will search their data sources to solve their specific problems.

Stage 1 : Apps

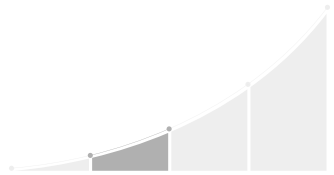

Milestones

More data sources, better indexing. The backend of the search platform is able to index and ingest more data sources including ones that are specific to just certain business units. Document parsing is better so search results are more precise.

More apps. You've gone past the one-size-fits-all apps and started building apps for specific divisions or business roles or functions. This is easier to do if everyone is building on the same platform.

Data silos eliminated. With comprehensive access to data, there are fewer barriers for the end users to get the data they need to do their jobs.

Metrics

Stage 1 is about measuring impact and includes engagement metrics like active users, queries per user, user satisfaction scores, and utilization of search versus just browsing.

Examples

- A sales team has an app that blends CRM data, marketing systems, and account activity into a customer 360 view of every account—no need to look in five places.
- A laptop manufacturer's support team has comprehensive search of all their support documents, tickets, and knowledgebase articles—plus a version for customers to use as well.
- Marketing mixes campaign, collateral, social media data, event, and website utilization data into a searchable timeline allowing them to see the holistic view of their campaigns and user engagement.

What's Next

Things are humming along pretty good. Lots of apps, lots of users—and they're actually using the apps. But, as more users come onboard, they're finding the ranking of search results doesn't really quite fit what they're looking for. Time to increase relevancy.

Warning

Don't make the data warehouse mistake! As you build out multiple apps it's easy to shoehorn user requirements into a handful of apps that are supposed to serve everyone. You should be able to offer various views of data sets that fit a particular user's role. Moreover, data warehouses are great at analytics to well-known questions but terrible at finding new questions or full-text search.

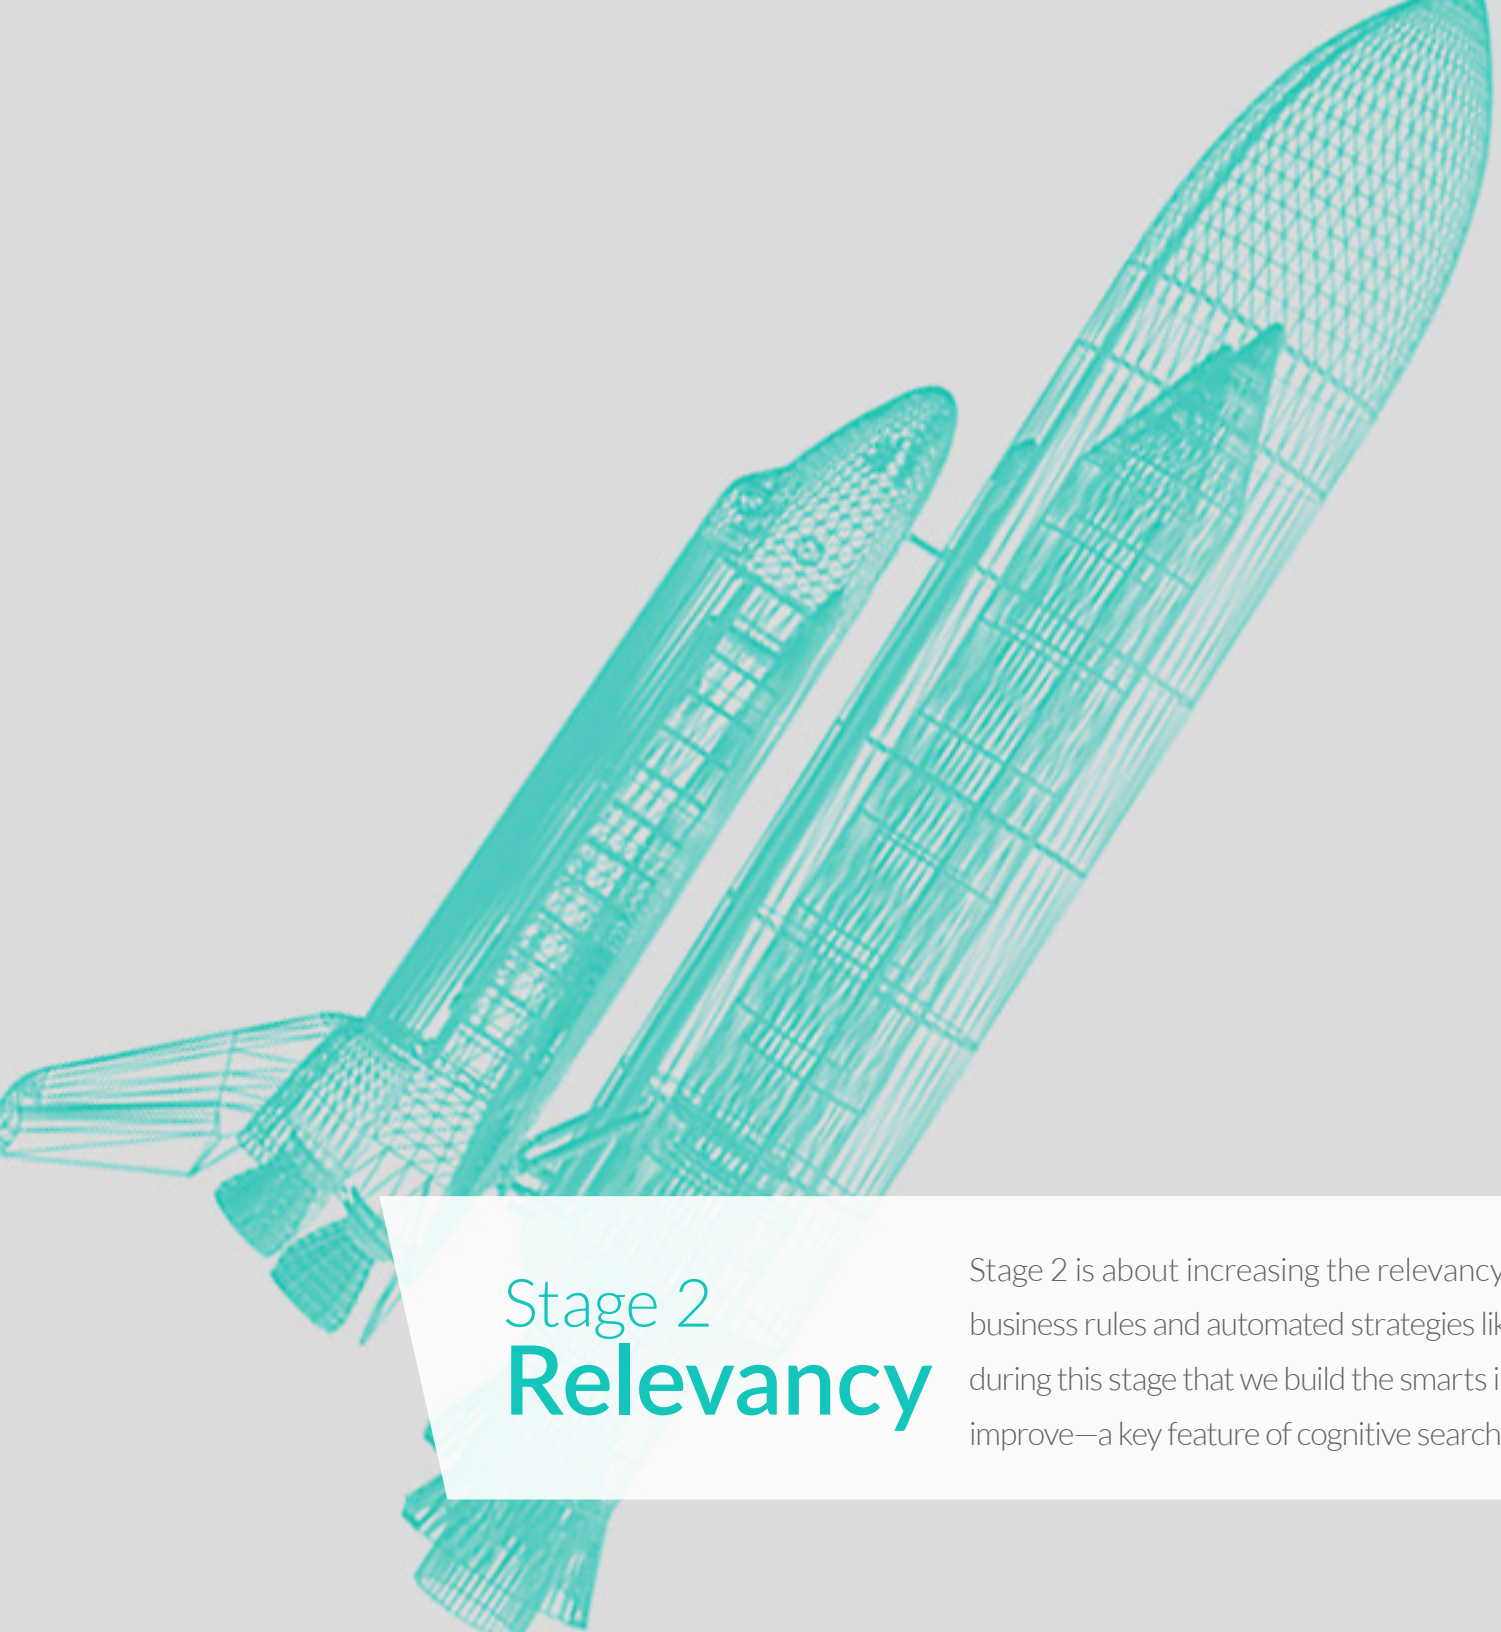

Stage 2
Relevancy

Stage 2 is about increasing the relevancy through both manual adjustments like business rules and automated strategies like signal capture and personalization. It is during this stage that we build the smarts into the system that allow it to continuously improve—a key feature of cognitive search.

Stage 2 : Relevancy

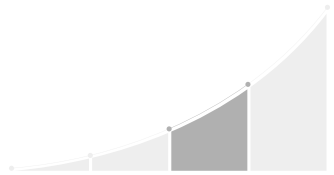

Milestones

Admins add boost and blocks. Business owners are able to add boosts and blocks for different documents, products, or certain criteria.

Business rules are applied. Business rules can be implemented to fulfill policies and practices.

User behavior is captured and aggregated. The system captures events like clicks, views, ratings, past queries, and other signals to constantly refine relevance.

Every end user's experience is personalized. Patterns in user behavior are aggregated to create customized views of data and search specific to each user.

Metrics

Success in Stage 2 means measuring conversion metrics like clicking to view a document, completing an online purchase, or submitting a review. Also important is constant A/B testing and experimentation to see what configurations work best for particular sets of users.

Examples

- A fitness clothing brand's store scans a user's purchase history and sees they've bought two pair of the same shoes in two different colors—and recommends they round out their collection with a new pair in the third color.
- An employee searches the intranet for the company's holiday schedule. The one for last year is still showing as the top results. They click on the current year in the list of results. That document is then boosted to the top the next time any user is looking for the holiday schedule.
- A research portal that gives researchers instant access to data from clinical trials in the field to the latest results from laboratories around the world.
- An insurance company is in the middle of discovery for legal proceedings. Researchers can search millions of documents, and get a zip of the key docs sent to them directly.

What's Next

Most companies stop at Stage 2. Until now that was fine. Most end users weren't demanding anything more. With the continuous fine-tuning of the user experience and all the data gathering that entails, the door opens to offer even more advanced capabilities. In today's global competitive climate, this is the edge we've all been looking for.

Warning

Don't skip over relevancy. It is easy to launch apps and be done with it. You hit your milestone right? But if the results are low quality—users will just up and rage-quit your app—and never look back.

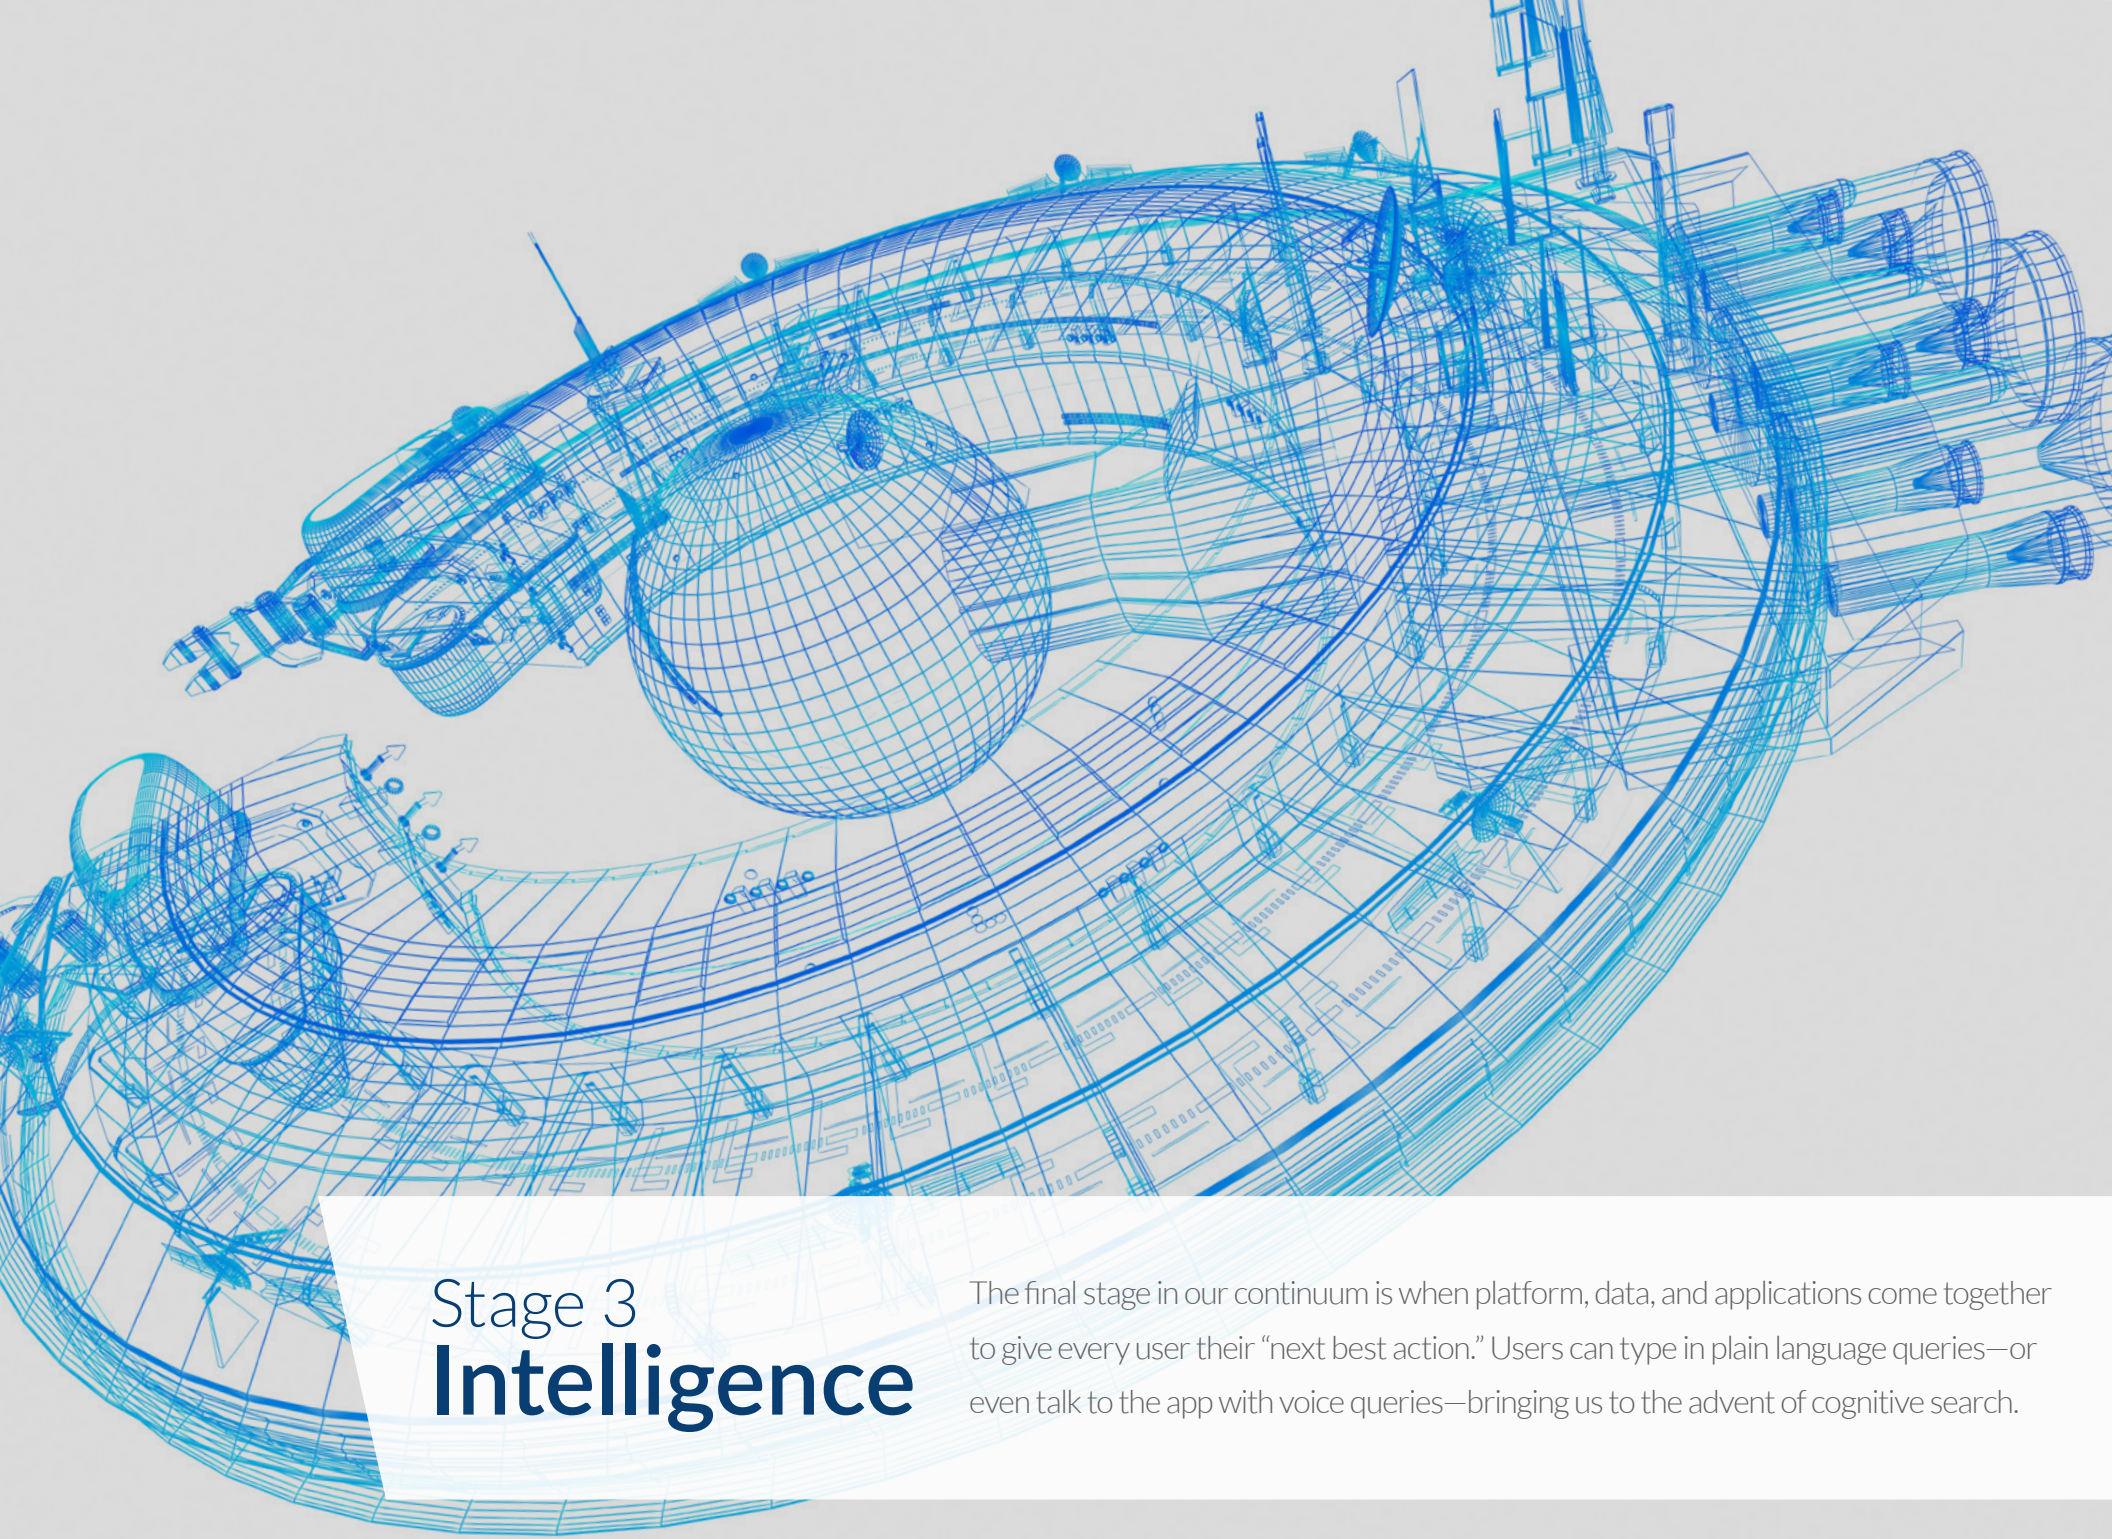

Stage 3
Intelligence

The final stage in our continuum is when platform, data, and applications come together to give every user their "next best action." Users can type in plain language queries—or even talk to the app with voice queries—bringing us to the advent of cognitive search.

Stage 3 : Intelligence

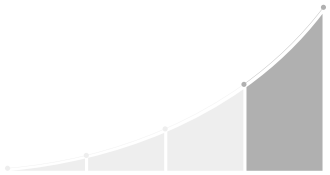

Milestones

You can talk to your apps. Apps have gone beyond typing into a search box and offer input through voice, image upload, or media file. When user's don't ask specific enough questions the system asks follow up questions.

Results reflect user intention. Search results take into account who and where the user is right now, what they did last, and what they're about to do.

Machine learning models deliver intelligence. The search system goes beyond signal capture and aggregation and constantly learns from user interactions, incoming data streams, and external inputs. The system doesn't always wait for users to ask the question but sometimes finds the questions in addition to the answers.

Metrics

Many of the metrics are just extended versions of what was measured at the last stage. Neural networks and pattern recognition require different kinds of tuning but their effectiveness is measured by "did I find what I'm looking for?" Predictive trends and patterns are found in user behavior, and business metrics are applied to those to understand their utility.

Examples

- Automated virtual assistants for customer support that ask clarifying questions before serving up the right help article or routing to a support agent.
- A financial services firm scans all inbound and outbound emails and attachments for possible signs of collusion or bad actors for escalation and investigation.

What's Next

This is the final stage of getting to the promised land of cognitive search. What happens next is further integration of platform, app, and end user to a system that can predict what a user wants to know before they do.

Warning

You can't start here. This is often where companies want to start. They see a segment on Bloomberg and run down to their IT department and breathlessly declare: "Alexa, Siri! AI! Yes! Now!" Until you've got a performant, secure system (Stage 0) that offers unified data access for all your data with multiple applications (Stage 1) and constantly fine-tuned relevancy and enrichment (Stage 2) you're probably putting several horses before the cart.

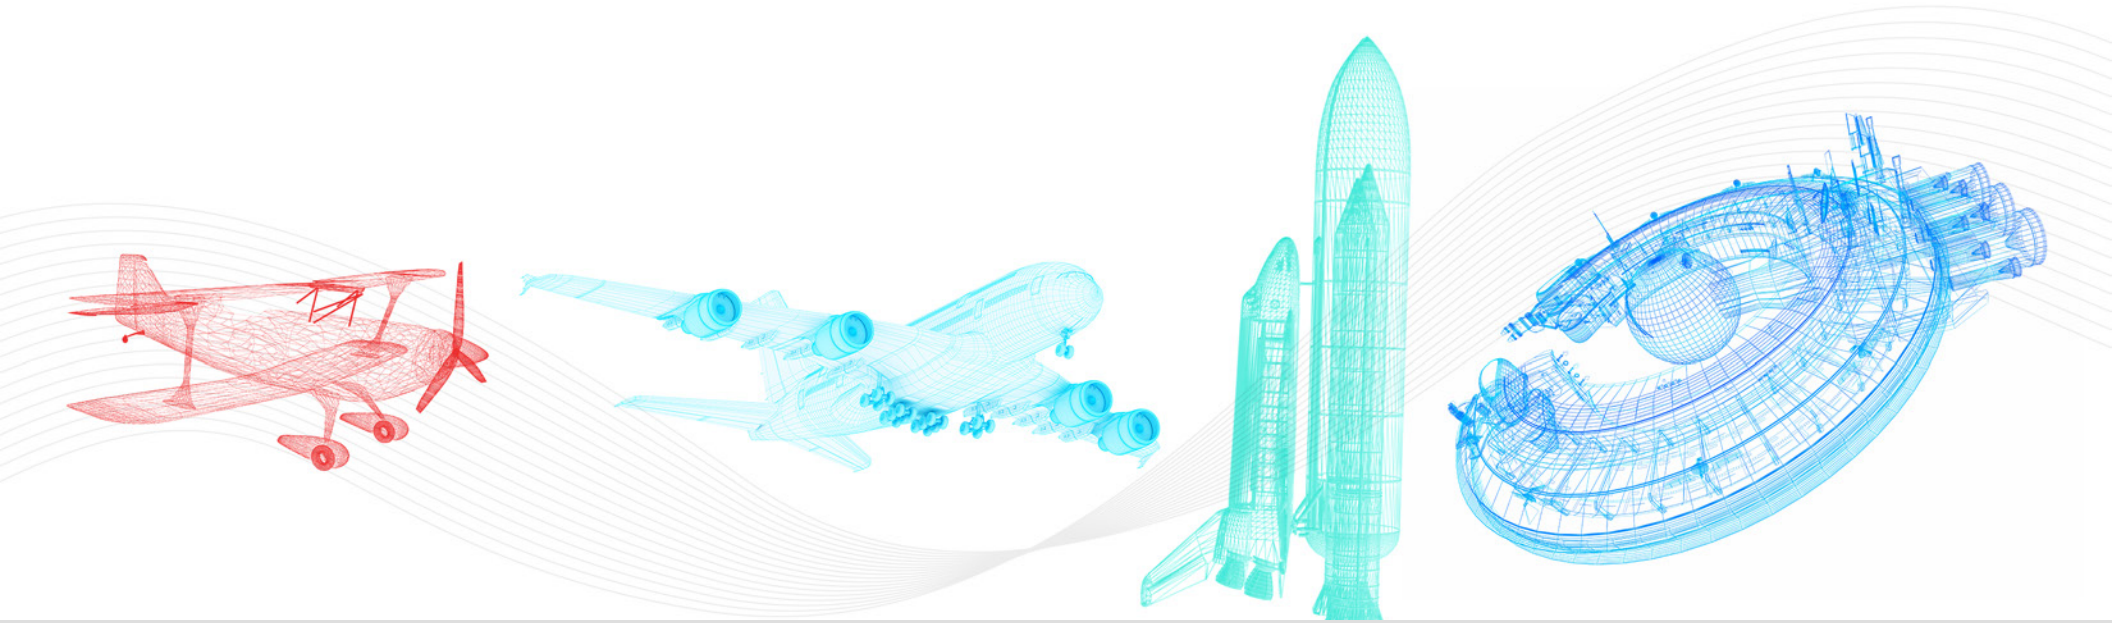

Conclusion

So that's how it happens. How your organization can chart the course from search box to automated assistants and everything else in between.

Lucidworks has been there and done that. We can help you navigate these trends as well as architect and deploy a solution that is scalable, relevant, and future-proof. If you are wondering where to start, contact us at **lucidworks.com/contact** or give us a call at **415-329-6515**.