Smarter Analytics: Taking the Journey to IBM Cognitive Systems

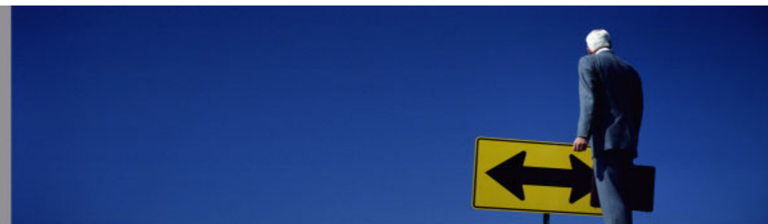

Redguides
for Business Leaders

Mandy Chessell
Eric Jenney
Scott Gerard
Bill Rapp

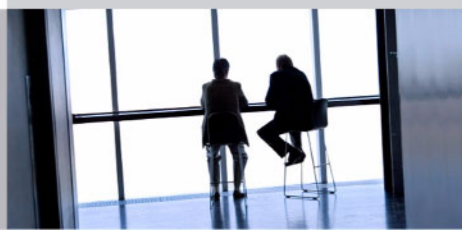

- Understand the importance of managing a wide variety of data
- Evolve skills and processes with the help of analytics
- Apply analytics to enhance decision making

Executive overview

Cognitive systems move information technology (IT) beyond automating standardized administrative tasks. A cognitive system is trained to extract facts from different sources of information. It compares and contrasts this evidence to make a recommendation to a professional who is in the process of making a decision. When using cognitive systems, a professional asks for explanations of how the recommendation was calculated and might be prompted to supply additional facts about the current situation to guide and refine the recommendation.

In this model, the human and machine are in a partnership. The human has the creative thought and intuition to assimilate the facts about a specific situation, develop an understanding of what is required, and make a decision on the next course of action, particularly when the information is limited. The machine has the speed to sift through huge repositories of existing knowledge to locate potentially relevant information and make an assessment of it. Together, they combine the best knowledge available and apply it to a specific real world problem.

Cognitive systems are a revolutionary technology. The challenge is to capitalize on its potential. Looking back in history to the industrial revolution or even the revolution that the IT industry brought to modern commerce, we see that technological advances played a key role. However, the impact of new technology was only felt when it was placed in an environment where it could be used effectively and where the environment was adapted to take advantage of the new capabilities.

This IBM® Redguide™ publication describes the transition that an organization must make to embrace the cognitive systems opportunity. It describes the three aspects of this transition, which are: How information is managed, how the analytics behind cognitive systems are maintained, and how the organization's processes and people adapt to take advantage of this new capability. The guide covers the types of transitions that must occur and how IBM can help with the transformation.

Introduction

Today, even top experts find it difficult to keep up with the growth of new knowledge within their field. Developments in science, engineering, medicine, economics, and business are accelerating as professionals reach out and connect to one another in worldwide communities, sharing insight and collaborating on new discoveries. Much of the knowledge they generate and exchange is designed for human consumption such as, research papers, videos, web pages, and audio podcasts. This knowledge also has the human quality in that it contains contradictions, inaccuracies, and inconsistencies in the use of terminology and background assumptions. It takes expertise and judgment to sift through this material and extract the relevant detail to make a decision.

Add to this explosion of knowledge, the growth in machine-generated data. Sensors and other measuring devices are recording what is happening in the physical world beyond the scale that our human senses can perceive. Some of this data relates to human activity particularly through digital devices such as, cell phones and tablets or from existing machinery and systems such as, automobiles and their associated transportation systems. This data needs summarizing and projecting in different ways in order for an expert to detect and extract new insight.

In the past, organizations have looked to IT to scale up their ability to process information. IT has served them well by automating many repetitive, deterministic, and literal tasks, enabling modern organizations to operate efficiently and effectively at a global scale.

Cognitive systems are the next phase of development for IT. It handles natural language to a level of sophistication well beyond the keyword matching capability that is found in a search engine. Cognitive systems analyze the structure of a sentence to determine the underlying intent of the speaker. It is able to request additional information from its human operator in order to hone its recommendation and eliminate irrelevant information. Cognitive systems provide the kind of exchange you would expect if you had a couple of professionals, each with their own experiences, discussing the next course of action on a difficult case. The exception is that with cognitive systems, one of the parties is a computer with a vast knowledge base but with less intuitive sense and creative ability than its human counterpart.

Commissioning a cognitive system

Since cognitive systems are a relatively new development, we have only begun to explore the possibilities that it opens up. Medical research and healthcare have seen significant success in developing cognitive systems to aid the treatment of cancer and in the organization of healthcare. Both examples enabled professionals to operate more effectively due to the additional relevant information and insight surfaced by the cognitive system during a consultation.

Imagine the value that cognitive systems could bring to your organization, to have a computer system that is able to synthesize the best knowledge about your industry. It can weigh evidence in order to act as an advisor to your knowledge workers and other professionals as they make decisions and apply those decisions to the markets and people that your organization serves.

Cognitive systems apply best in domains with the following characteristics:

- A shortage of expertise that is impacting your organization's growth in delivering new services and higher customer service.
- Vast amounts of data and much of the data is in text form that currently can only be consumed by people, if at all.
- The decision making in the domain is not an exact science, calling for the judgment and experience of subject matter experts, who are interpreting evidence from various sources.

- Decision making within the domain calls for an adaptive, learning mode of operation where the people using the system are able to provide feedback and additional insight in order to improve the system's performance over time.

Manufacturing example

Consider an organization that designs and sells a range of complex technical products to individuals and businesses such as, an automobile manufacturer or an IT manufacturer (such as IBM). This organization is contemplating a cognitive system in order to improve their design and maintenance of their product line:

- The design engineers working on the next generation of products need to understand key information about their products such as:
  - How well the existing products were received by their customers?
  - How competitive the products are in terms of maintenance effort and costs, both to the manufacturer and to their customers?
  - What the latest developments in new technology are that could be incorporated into new product features?
- The engineers maintaining the products that customers have bought need to be able to diagnose the causes of problems in the product and decide on effective remediation.
- The customers who are looking to buy a product need to understand the options, cost of ownership, and operational aspects of owning the product. After they have purchased a product, they might need help in using it and choosing enhancements.

This scenario is well suited for cognitive systems. It requires information from various resources, including hand-written notes from engineers, text-based specifications of parts and products (potentially from multiple suppliers), feedback reports, and maintenance best practices, along with the enterprise data that is used to manage the organization's business. Analyzing and understanding each type of data takes various analytics techniques. And the people who will be working with the cognitive system are solving challenging problems that need to match information about the current situation with the large body of knowledge that the organization has collected.

It is important to understand that a cognitive system is not able to generate new knowledge. Its power lies in its ability to put existing knowledge to best use. When a new cognitive system is commissioned, it must be loaded with a knowledge base covering the domain that it is to operate in and be trained in how to interpret the information within it. This training comes from the domain experts that could be the principle users of the cognitive system going forward. However, cognitive systems need some form of ongoing interaction with these subject matter experts to keep its knowledge base current with recent developments. Ultimately, the quality of its recommendations is heavily influenced by the quality of the interactions that it has with subject matter experts, even if they are not the only users of the cognitive system.

Building an ecosystem for a cognitive system

A cognitive system is surrounded by an ecosystem that continuously feeds new knowledge and the expertise of professionals into its knowledge base as part of every day business. Many organizations might not have all of the expertise to operate the entire ecosystem alone. Instead, they could become a major partner in a consortium supporting the domain.

Figure 1 on page 4 illustrates the ecosystem around a cognitive system.

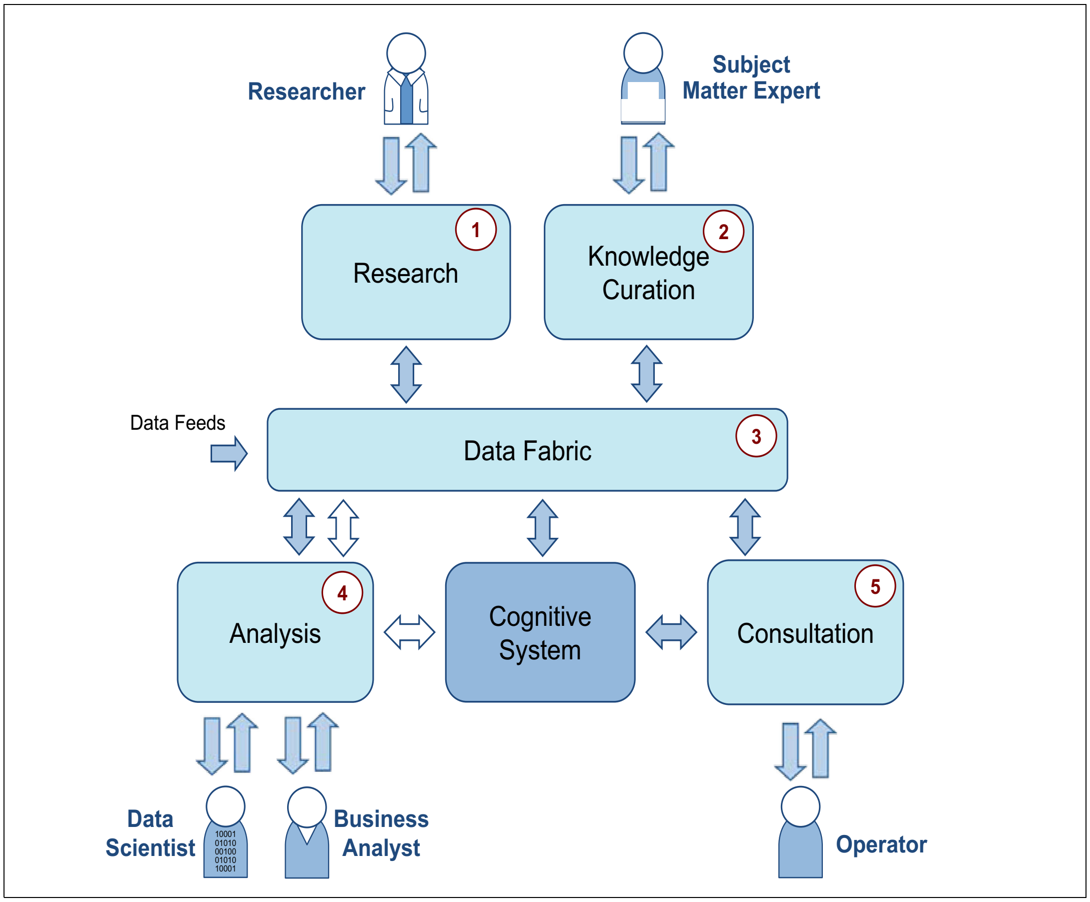

Figure 1 A cognitive ecosystem

The different subsystems that make up the cognitive system's ecosystem (Figure 1) are as follows:

- Research (Figure 1, item 1) is generating new knowledge in terms of documents and experimental results.
- Knowledge curation (Figure 1, item 2) is capturing the problem solving processes and techniques from subject matter experts. It also enables subject matter experts to provide feedback and correct the decisions and recommendations that are made by the cognitive system.
- Data fabric (Figure 1, item 3) is collecting, integrating, and transforming information from many sources and supplying it to the cognitive system and the systems around it. Each system takes the subset of this information that is relevant for its operation.
- Analysis (Figure 1, item 4) is developing new analytics to improve the processing inside the cognitive system and to create new information for its knowledge base.
- Consultation (Figure 1, item 5) describes one or more systems that are interacting with the cognitive system to make a decision about a specific situation.

The sections that follow describe the subsystems in more detail.

Data fabric subsystem

The data fabric subsystem is the information heart of the ecosystem, preparing and pumping the right data to the other subsystems when they need the data. Data passed to these subsystems includes:

- Domain entities describe the people, places, and assets details and business transactions that relate to the domain. For a manufacturer, this includes:
  - Customers, employees, suppliers, and business partners
  - Products and parts definitions
  - Sales transactions and maintenance visits
  - Invoices and other financial data

  This data is typically located in the enterprise systems and is structured or semi-structured data. This data is used directly by the consultation subsystem and anonymized versions of it can be added to the knowledge base of the cognitive system and the data stores used for research and analysis.

- Domain activities record the activity around the domain entities.

  For a manufacturer, these activities would include business transactions, such as sales of products, or comments and complaints from customers, onboarding suppliers, operations around an industry conference, and analyst reports.

- Domain codes are sets of values that appear in tables or hierarchies.

  This data is used to classify the domain entities and domain activities. A manufacturer may use codes to classify the products, parts, and assemblies, perhaps in terms of their features or success in the market. Locations can be classified by country codes and state codes. Customers can be classified in terms of their buying history or demographics. Classifications using codes are amenable to analysis, making these codes a valuable part of a domain's knowledge base.

- Domain documentation refers to the published specifications and related documentation that is published for the domain.

  This documentation may originate from the organization or from trusted external parties. For a manufacturer, this documentation could include specifications for the products and parts for the goods and services that are sold by the manufacturer, research papers on the materials used, published performance papers, and peer-reviewed publications. This documentation is believed to be good quality, but may contain gaps and inconsistencies in the use of terminology and background assumptions. Any of the subsystems may use this documentation as reference material.

- Domain notes include the notes and feedback that are recorded by subject matter experts and users of goods and services in the domain.

  For a manufacturer, the notes could be about engineer designing, testing, or maintaining products. These notes might be hand written and they are often incomplete since the individual that is writing them is often very busy. These notes are best supplemented with data captured automatically by domain monitoring to provide context to the observations made in the notes.

- Domain monitoring provides data from sensors and other instrumentation that records environmental conditions and the status of objects from the physical world.

  For a manufacturer, data could pertain to:

  - The operating environment of a product (such as, location, temperature, load, length of operation, and detected faults)
  - The tests that are being performed with results

– The products' current configuration as its parts are replaced during its lifetime

Domain monitoring data is often verbose and must be analyzed and summarized to extract useful information from it. This data must be correlated with the domain entities and activities.

- Domain insight describes information that is derived by harvesting and analyzing information about the domain.

  This insight is created outside of the cognitive system, possibility for another application and is managed in the data fabric.

Capturing data for the cognitive ecosystem is an ongoing process. Most of this data requires preprocessing to turn it into a useful form for the cognitive ecosystem to use. This preprocessing could be as simple as reformatting the information into a convenient and consistent structure, or it might involve complex analytics to parse and interpret the raw data being received.

After the preprocessing is complete, the data is stored in an appropriate form for the other subsystems to use. Figure 2 shows the details of the new data capture processing.

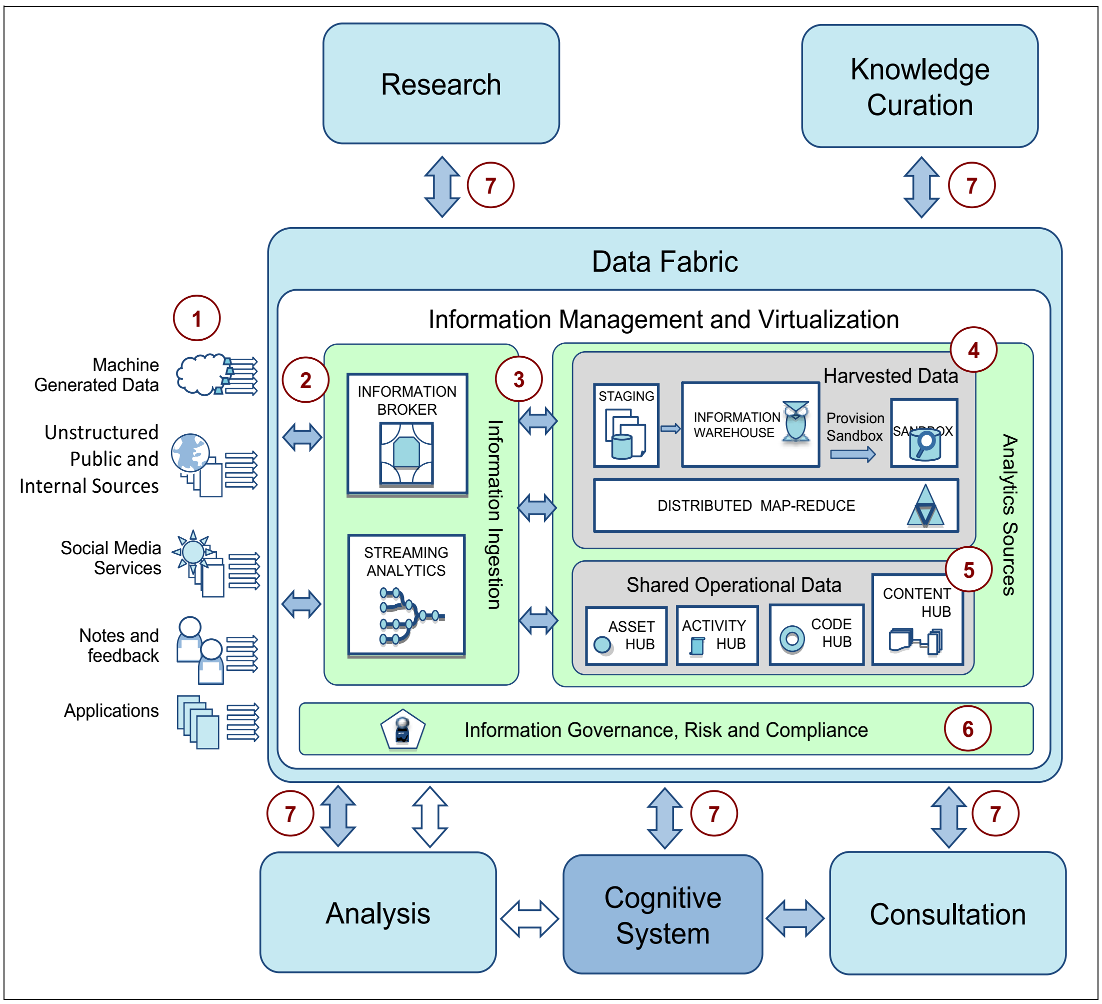

Figure 2 Capturing and managing data for the cognitive ecosystem

The details of Figure 2 on page 6 are as follows:

- New data (Figure 2 on page 6, item 1) can come in through many channels, and typically each type of data has different characteristics that affect how it can be processed. These characteristics include the structure of the data, how fast it is arriving, what level of quality it has, how much data is involved, and how the useful part of the data is separated from the data that is not relevant.
- Provisioning engines (Figure 2 on page 6, item 2) are used to receive the data, analyze, and assess the data before transforming it to the desired form.

  There are different types of provisioning engines, each specialized to manage particular types of data. Figure 2 on page 6, item 3 shows two examples of provisioning engines:

  - The streaming analytics provisioning engine is designed for data that is arriving too fast to store and must be processed in memory in order to extract out a small subset of relevant data that can be stored.
  - The information broker provisioning engine moves and transforms data between internal systems.
- Information is stored in various types of data stores to suit the way that the data must be stored, accessed, and managed. In general, these are divided into two groups:
  - Harvested data (Figure 2 on page 6, item 4) is data that is copied from many places and stored so that it can be used as a reference for analytics processing. After it is copied, it will be transformed, collated, and corrected to make it suitable for analytics. As this data is updated in the original systems, copies of the new values are fed to the harvested data so that a history of how these values are changing over time can be maintained. This data is useful in determining patterns of events, activities, and their resulting outcomes that can be incorporated into the analytics used within the cognitive ecosystem.
  - Shared operational data (Figure 2 on page 6, item 5) is data that is being used by applications and processes that are operating in the organization on a day-to-day basis. It typically contains the latest values with very recent history. Examples of this type of data include:
    - Asset hubs that contain information about the assets, people, products, accounts, and contracts that form the key domain entities of the organization.
    - Activity hubs summarize recent activity about these domain entities from many sources.
    - Code hubs act as reference sources for domain code translation tables and other lookup data that is used in the domain.
    - Content hubs are repositories for controlled documents and other media files such as, product images and videos.

  The shared operational data is one of the feeds into the harvested data.
- Information governance risk and compliance (Figure 2 on page 6, item 6) cover the people, processes, and technology that control the management of information in the data fabric. Information management is used to ensure that data is of appropriate quality and is managed and used according to the organization’s policies by all consumers of the data throughout its lifetime.
- There are the interfaces where the different subsystems consume data (Figure 2 on page 6, item 7). Each consuming subsystem (research, knowledge curation, analysis, cognitive system, and consultation) needs data in its own format. The data fabric provides information virtualization capabilities to ensure that the subsystems receive and update the information that they need without having to know how it is physically organized and structured.

The data feeds for the data fabric subsystem ensure that the cognitive ecosystem has the latest information to support its decisions, which includes the latest discoveries from research.

Research

Researchers all over the world are continuously pushing the boundaries of our knowledge. One of the benefits of cognitive systems is that they enable an organization to make use of new research.

Figure 3 shows more detail of the research process.

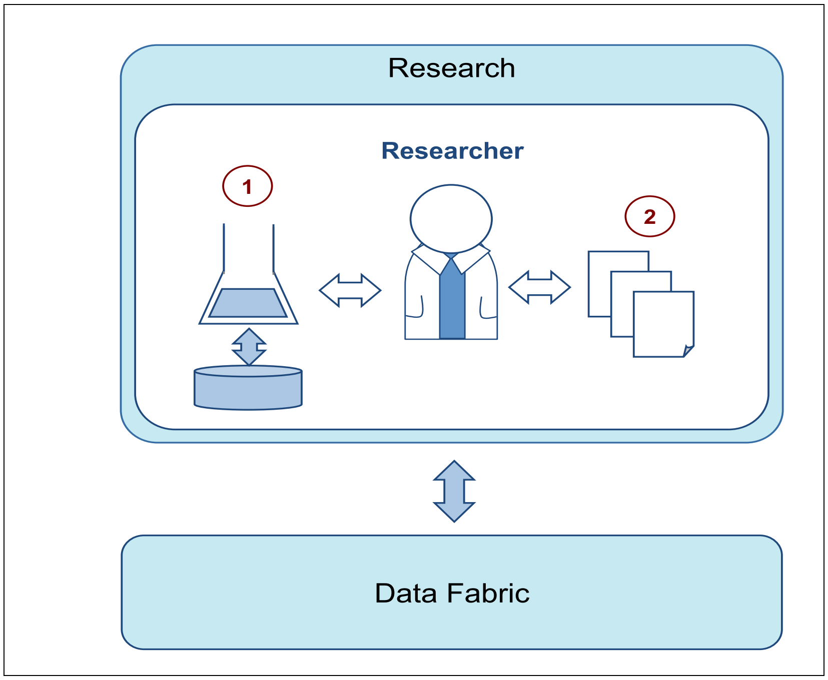

Figure 3 A researcher creating new information for the cognitive system

Most research is about performing experiments to test a theory or hypothesis. Research produces two types of data for the cognitive system:

- Data sets from experiments and test runs (Figure 3, item 1). Data that is used as either input to the experiment or produced by the experiment. This data can be referenced and analyzed by the cognitive ecosystem.
- Research papers and articles (Figure 3, item 2) that describe the research work and its results. These documents may refer to the data sets used during the research.

The results of relevant research are fed into the data fabric and prepared for processing by the information ingestion provisioning engines.

Analysis

Research is not the only source of information for the cognitive systems. Most organizations have access to a wide range of information, both from its internal systems and from public sources. However, researchers are not sure which information is relevant to the domain and how to use this information in decision making.

Figure 4 illustrates the process that is used by an organization to analyze and understand information sources in order to configure various decision models for use in the cognitive ecosystem.

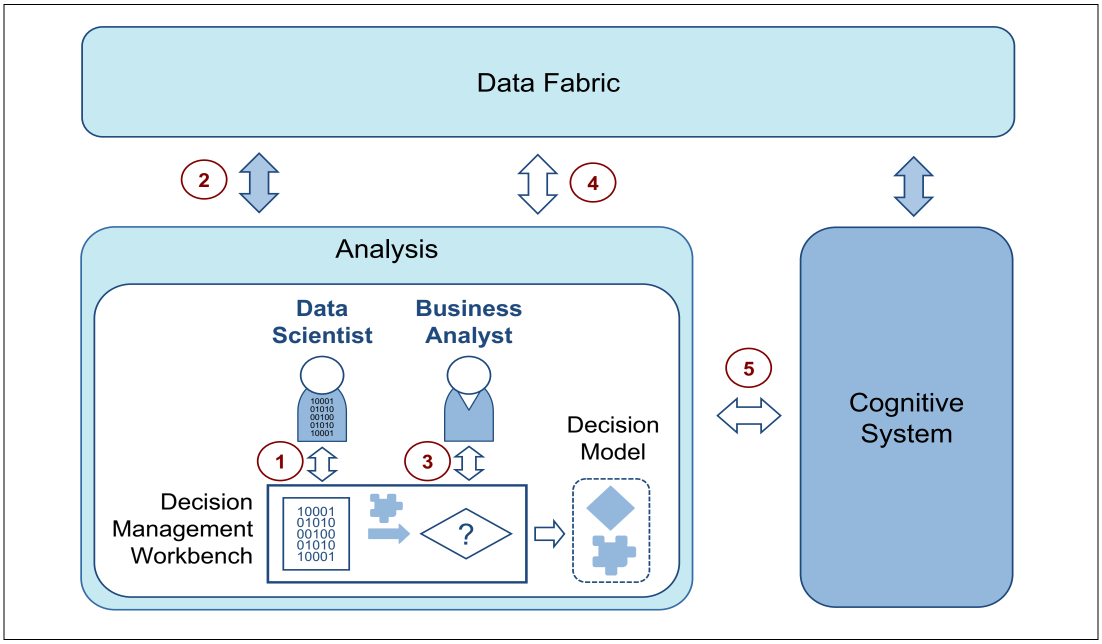

Figure 4 The development of decision models by data scientists and business analysts

In Figure 4, the process is as follows:

- A data scientist (item 1) performs analysis on a portion of the data that is extracted from the data fabric (item 2). The result of this work is one or more analytics models.
- A business analyst (item 3) selects the appropriate analytics models and combines them with business rules to create a decision model.
- The decision model can be deployed into the data fabric to improve the extraction, enrichment, and management of the knowledge base (item 4).
- Alternatively, the decision model can be deployed directly into the cognitive system (item 5) since it is pertinent to a decision-making process supported by the cognitive system.

The analysis process is an iterative and ongoing process that is responsible for both developing new analytics and refining the analytics already deployed in the cognitive ecosystem. The data fabric collects the results of executing the existing decision models for use by the data scientist when assessing and refining an analytics model.

Knowledge curation

Information that is stored in the data fabric might need input from subject matter experts to define how to interpret it, and to score the reliability of a particular source of information. Figure 5 on page 10 shows the activity that is involved in knowledge curation.

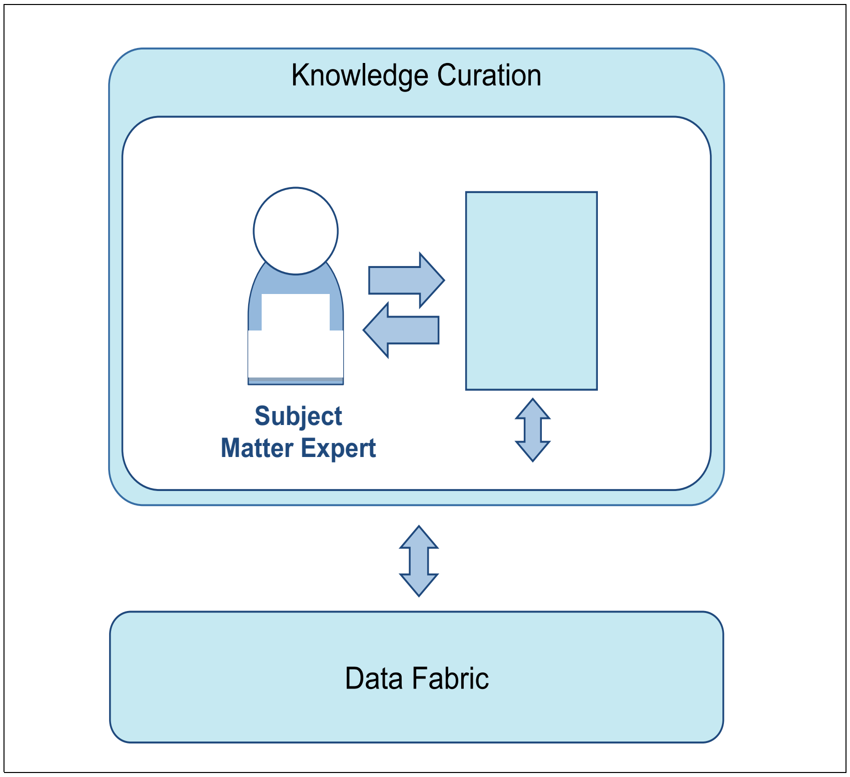

Figure 5 A subject matter expert validates and extends the cognitive system's knowledge base

In Figure 5, the subject matter expert uses a system to capture their expertise around assessing information sources and making decisions. The information that is supplied by the subject matter expert is stored in the data fabric.

Consultation

The purpose of the cognitive system is to support a person who must make a decision. Figure 6 illustrates this process.

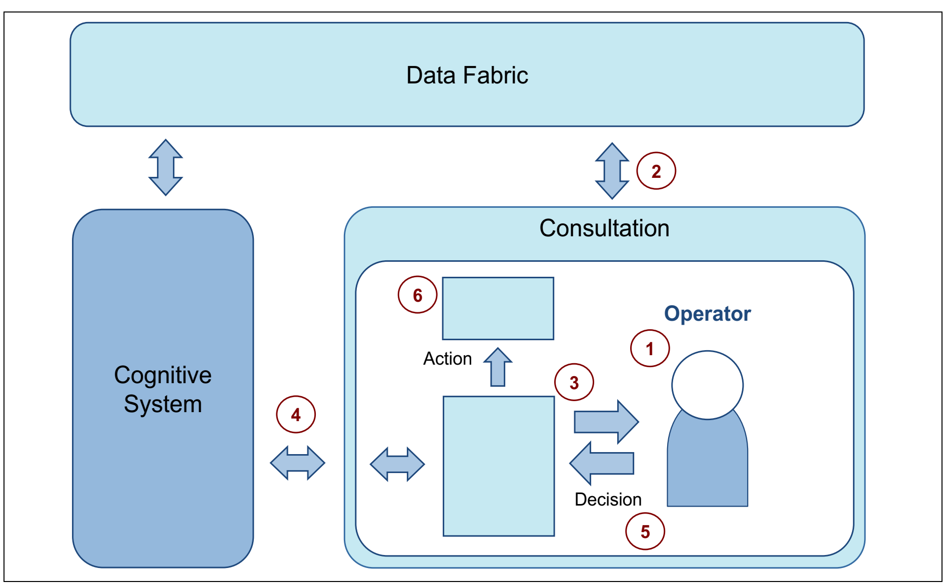

Figure 6 Consulting the cognitive system

The following activities depict consulting the cognitive system (Figure 6 on page 10):

- The operator (item 1) is the person making the decision. For a manufacturer, this could be an engineer who is trying to diagnose a difficult problem relating to a product.
- The operator is using an application that is designed to support their task (item 2).
- The application uses the data fabric to extract the latest information that is related to the specific decision (item 3), for example, details of the specific product that has failed.
- The application then starts to consult the cognitive system (item 4), feeding it relevant facts from the specific case. The cognitive system may return a number of possible causes and request more information. The application displays this information to the operator.
- The operator assesses the results so far and performs actions to acquire the additional data such as, performing a test on the product to create more diagnostic information.
- The application passes this to the cognitive system and the results are refined.
- The process repeats until the operator makes a decision (item 5) and the application triggers the appropriate action (item 6). For the manufacturer example, this could be to order a different part, print instructions on how to repair the product, or pass the problem onto a support team to resolve.

This ecosystem uses a common pattern when interfacing with different types of users that is repeated in the analysis both by the knowledge curation and the consultation subsystems. Each type of user has their own application to support their role and the cognitive system is called in the background by the application whenever it is needed. This approach avoids the need for an individual to switch applications when they need to make a decision, along with the added inconvenience of copying the data. Using the cognitive system becomes a natural part of the operator's every day work.

Developing a roadmap for cognitive systems

Figure 7 on page 12 summarizes the detail of the ecosystem around the cognitive system.

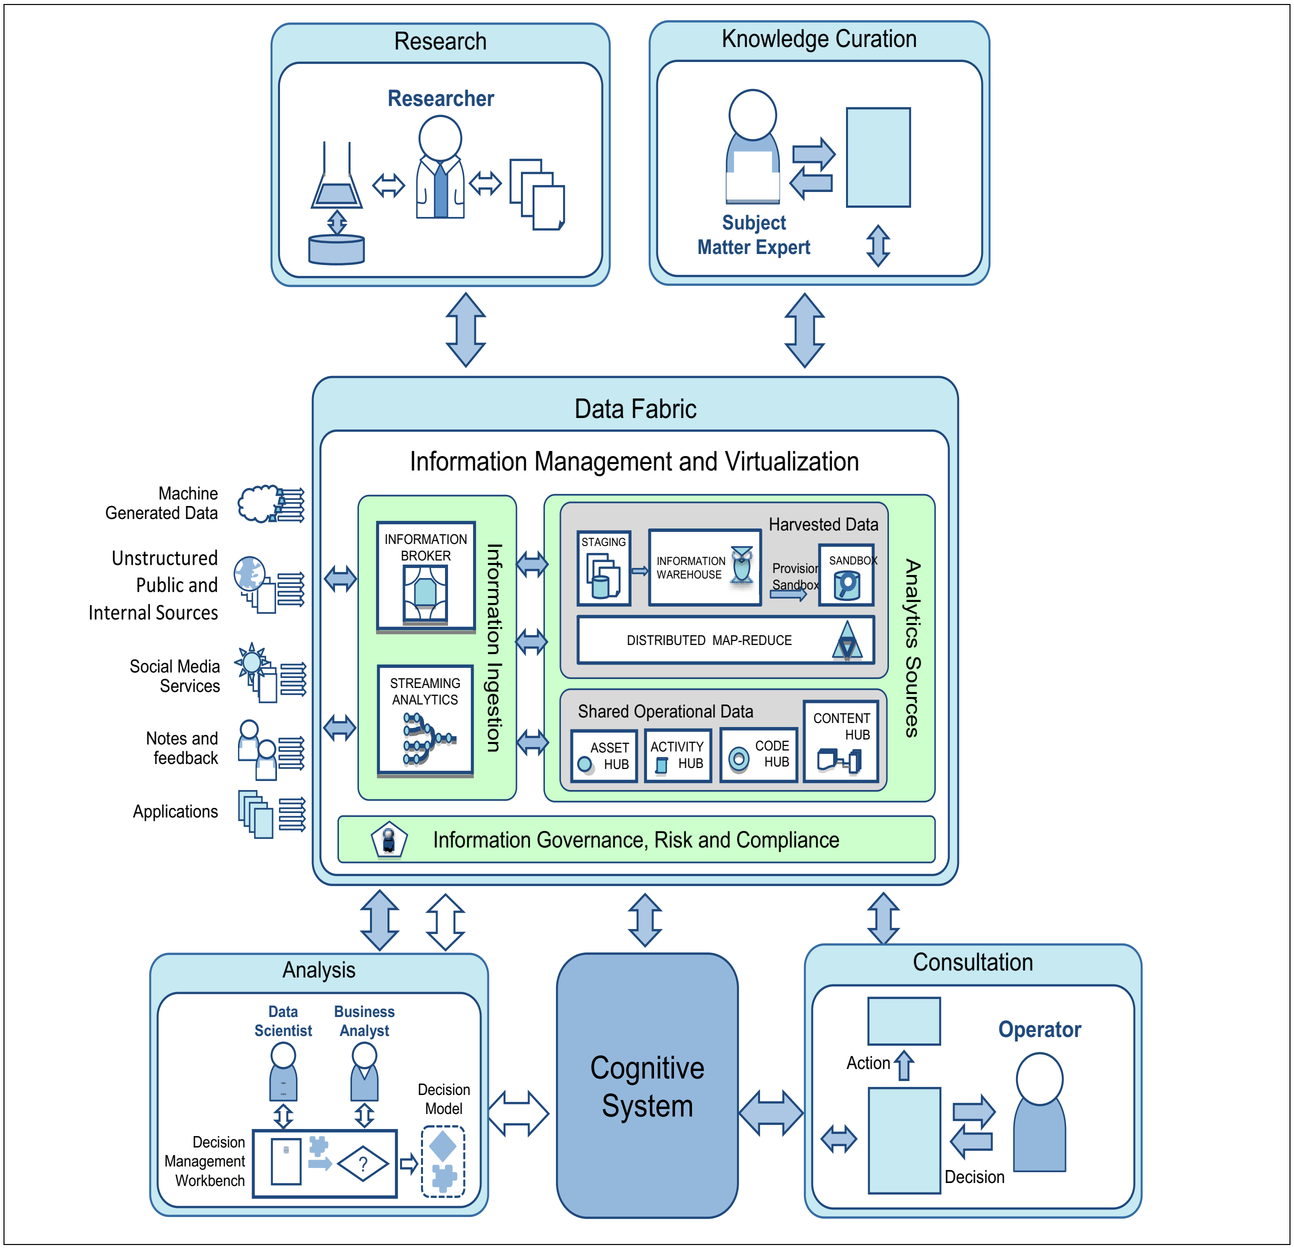

Figure 7 Detail of the cognitive ecosystem

Ideally, a cognitive ecosystem is developed incrementally. The organization's existing systems start as a baseline. These systems provide the initial information sources for the data fabric. Over time, a series of analytics solutions are deployed. Each solution delivers value to the business in its own right and contributes to the evolution of the cognitive ecosystem going forward.

Project by project, the cognitive ecosystem is expanded and enhanced by focusing on three aspects in each solution:

- Enhancing information management

  Enhancing information management aligns data to the needs of the cognitive system in the data fabric.

At each enhancement, more information sources are introduced into the data fabric. The level of governance over that data is enhanced and more data is made available via information virtualization to meet the needs of the new solution.

- Enhancing analytics algorithms

  Enhancing analytics algorithms builds out the ability to generate insight that is needed by the cognitive system going forward.

  At each enhancement, more decision models are built, covering a broader range of decisions and potentially using increasingly sophisticated analytics models.

- Enhancing processes and skills

  Enhancing processes and skills ensures that the organization manages information appropriately, consults the cognitive system as necessary, and acts on the insight that is created by the cognitive ecosystem.

  At each enhancement, more support is given to a broader range of users. Their applications are updated to initially connect to analytics services and then later to the cognitive system itself. The individuals that use the cognitive ecosystem are given additional training on how to perform a consultation and how downstream processes are affected by decisions made in the consultation.

Figure 8 shows the enhancement of information management, analytics, processes, and skills as each solution is rolled out.

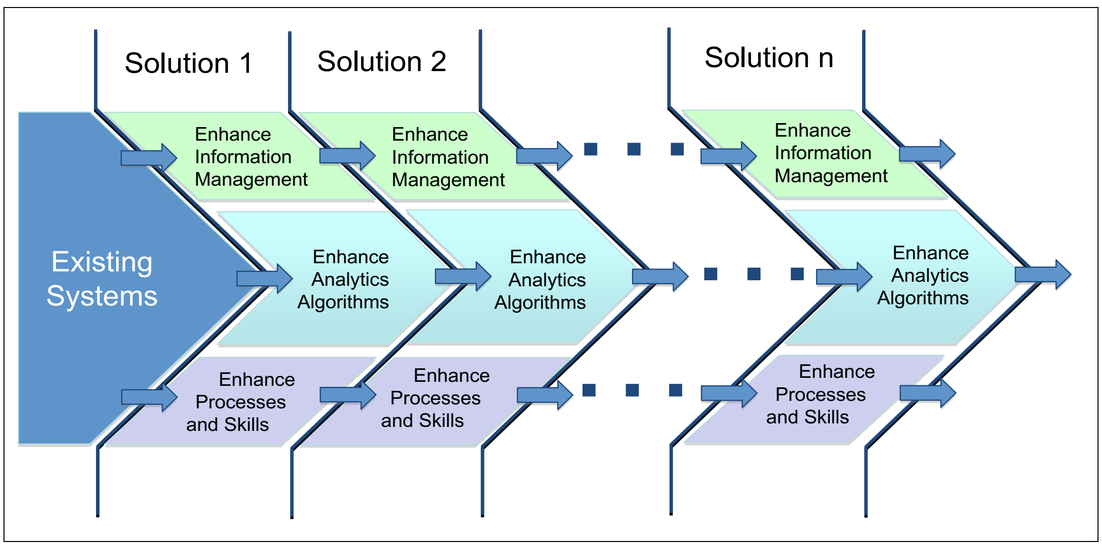

Figure 8 Developing a cognitive ecosystem

To understand how this approach works, it is worth returning to the manufacturer example. Initially, the manufacturing systems focus on running the business with some management reporting for financial and planning purposes. The manufacturer's team develops and deploys the following solutions to iteratively build out the ecosystem.

Each solution captures and manages more data, adds new analytics to create additional insight, and improves the processes and skills of individuals in the organization:

- Solution 1: Product and parts maintenance portal

  This solution offers product and parts documentation with a search engine for use by customers and engineers. Table 1 shows how solution 1 enhances the organization.

Table 1 Three aspects of solution 1

Aspect of solution	Outcome
Enhanced information management	Locates and organizes documentation from internal teams and parts suppliers into a coherent collection so that a search engine can operate over them.
Enhanced analytic algorithms	Develops analytics to measure which products have the most searches against them.
Enhanced processes and skills	Engineers are given education to make them aware of the new portal.

- Solution 2: Personal maintenance history for products

  This solution offers customers and engineers an online maintenance book for each product sold. This maintenance book records these items:

  – The purchased product's configuration when it was sold along with links to the appropriate documentation from solution 1.
  – All changes due to maintenance or upgrade made to the product.
  – Safety checks and inspections that are made on the product.
  – Recommendations and other communications from the manufacturers.
  – Suitable accessories and extensions.
  – Opportunities to feedback comments to the manufacturer.

  The owner can enter additional information such as the fuel/power or other consumable items that are used when the product is operated. They can display metrics on efficiency, cost of ownership, and depreciation. Table 2 shows how solution 2 enhances the organization.

Table 2 Three aspects of solution 2

Aspect of solution	Outcome
Enhanced information management	Develops a complete record of each product instance that is sold and what happens to it during its lifetime. This data provides invaluable information to the design engineers looking at future products and provides immediate value to customers and maintenance engineers.
Enhanced analytics algorithms	New analytics are developed to provide customers with metrics around the performance of their purchase.
Enhanced processes and skills	The design engineers are trained to use summaries of the online maintenance books in their background preparation for the design of new products.

- Solution 3: Diagnostics rules for product maintenance

  This solution offers maintenance engineers a simple expert system that is based on decision rules to help them diagnose common problems with the product. The expert system is able to advise on the diagnostics to run and aid the interpretation of this information. Maintenance engineers are able to submit their own diagnosis advice into the system, which will be incorporated into the knowledge base after an expert verifies it. The experts can also comment on and provide a star rating for the advice documented.

  Table 3 shows how solution 3 enhances the organization.

Table 3 Three aspects of solution 3

Aspects of solution	Outcome
Enhanced information management	Uses a moderated crowd-sourcing approach to build up a knowledge base of product problems and their diagnosis.
Enhanced analytics algorithms	Enables data scientists to begin to explore the knowledge base to uncover patterns in the diagnosis information that could point to frequent issues and common root causes.
Enhanced processes and skills	Informs design engineers that are keen to eliminate issues in the next generation of products, and managers wanting to understand the future liability of existing products in service.

- Solution 4: Predictive maintenance advisor

  This solution offers notifications to customers and nominated maintenance engineers when maintenance is required on a product. The advice in the notification is calculated with analytics using the age of the product, monitoring information collected from sensors located in the product, and analysis of maintenance information from similar products operating in similar conditions. The advice is also stored in the online maintenance book from solution 2. Table 4 shows how solution 4 enhances the organization.

Table 4 Three aspects of solution 4

Aspects of solution	Outcomes
Enhanced information management	Captures sensor data that records the operating environment and parameters of the product's operation.
Enhanced analytics algorithms	Combines sensor data and diagnostic information from solution 3 to develop predictive analytics decision models that are used to determine when parts within the product need replacing.
Enhanced processes and skills	Maintenance is reduced because parts are only replaced when required, based on evidence of actual usage, rather than on a pessimistic schedule.

- Solution 5: Product design advisor

  This solution offers consultation on design choices around parts, suppliers, and product configuration that is based on a combination of information from solutions 1 - 4. The product design advisor also offers consultation on new information from potential suppliers, new research into materials and mechanisms, along with analyst and competitive reports. This solution assists design engineers to validate a proposed design and advise them of the recent advances in the industry.

  This solution provides a true cognitive system since it requires natural language processing to interpret research documents, analyst reports, and other trade information. It also relies heavily on the ecosystem of information, analytics, processes, and skills that are built up by the earlier solutions.

  Table 5 shows how solution 5 enhances the organization.

Table 5 Three aspects of solution 5

Aspects of solution	Outcomes
Enhanced information management	Adds additional information about suppliers, advances in the state of the art, and opinions from analysts and other reporters.
Enhanced analytics algorithms	Develops natural language processing interpret research documents, analyst reports, and other trade information.
Enhanced processes and skills	Creates a partnership between the cognitive system and the design engineers.

With the last solution, design engineers, maintenance engineers, and customers are working within the cognitive ecosystem. The cognitive system continues to evolve as these people make use of the analytics and provide feedback, and additional information sources and decision models are added.

Figure 9 summaries the rollout of the solutions that constitute the manufacturer's pathway to cognitive systems.

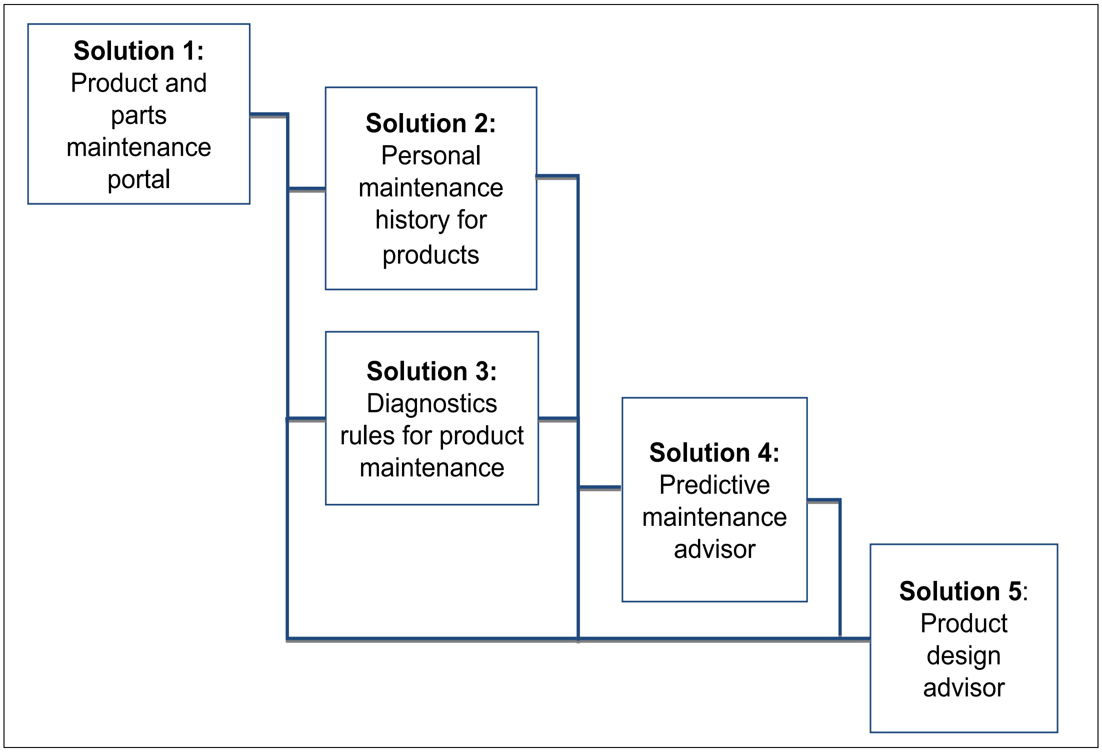

Figure 9 Summary of a possible manufacturer's solutions that lead to a cognitive ecosystem

Case studies

Table 6 provides examples of IBM case studies for the healthcare industry.

Table 6 Case studies for cognitive systems

Client	Solution case study information
Memorial Sloan-Kettering Cancer Center	IBM Watson™ helps fight cancer with evidence-based diagnosis and treatment suggestions: http://www.ibm.com/innovation/us/watson/pdf/MSK_Case_Study_IMC14794.pdf
WellPoint, Inc.	IBM Watson enables more effective healthcare preapproval decisions using evidence-based learning: http://www.ibm.com/innovation/us/watson/pdf/WellPoint_Case_Study_IMC14792.pdf

Summary

A cognitive system works with large amounts of information about a topic from many sources, which are stored in a variety of formats. It is able to extract facts from text documents and other material that is designed for human consumption. It organizes this knowledge and is able to engage in a natural language interaction with a professional/expert in the domain, as they jointly develop a recommendation on the next course of action to take in a complex situation.

Such a system needs an ecosystem around it that incorporates information feeds, analytic models, and subject matter experts to nurture the knowledge base behind the cognitive system. The cognitive system is embedded in the business operations.

IBM offers workshops, assets, products, and services engagements to help you build out a cognitive ecosystem. If this is of interest, contact your IBM representative or call us directly at: 1-800-966-9875.

Other resources for more information

For additional information about cognitive systems, access the following information:

- *The Era of Cognitive Systems - An Inside Look at IBM Watson and How it Works*, REDP-4955:

  http://www.redbooks.ibm.com/abstracts/redp4955.html?Open
- *Transforming the Way Organizations Think with Cognitive Systems*, REDP-4961:

  http://www.redbooks.ibm.com/abstracts/redp4961.html?Open
- *Artificial Intelligence: Learning Through Interactions and Big Data*, REDP-4974:

  http://www.redbooks.ibm.com/abstracts/redp4974.html?Open
- IBM Cognitive Systems:

  http://www-03.ibm.com/systems/power/solutions/cognitive-computing.html
- IBM Watson web pages:

  http://www-03.ibm.com/innovation/us/watson/index.shtml

Authors

This guide was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO).

Mandy Chessell is an IBM Distinguished Engineer and Master Inventor and is currently the Chief Architect of IBM InfoSphere® Solutions, working in the Chief Technology Officer (CTO) Office of SWG Information Management. She has expertise in designing information supply chains for information intensive solutions. Mandy joined IBM in 1987 and has held roles for developing new features for various IBM products such as IBM CICS®, IBM TxSeries, Encina, Component Broker, and IBM WebSphere® Application Server. She is a Fellow of the Royal Academy of Engineering (FREng), a Chartered Engineer (CEng), and a Fellow of the British Computer Society (FBCS).

Eric Jenney is the Program Director for the IBM Business Analytics and Optimization Architecture Board. He has been with IBM for 25 years, working on various products in these organizations: IBM Software Group, IBM Systems Technology Group, and IBM Sales and Distribution. Eric currently focuses on the IBM Analytics and Optimization solutions and works with associated Chief Technology Officers (CTOs) in order to develop consistent and coordinated offerings.

Scott Gerard is a Senior Technical Staff Member in the United States. He has 40 years of experience in software architecture, design, and implementation, and has worked at IBM for 34 years. His areas of expertise include software architectures, natural language processing, and multiagent systems. He has written extensively on various topics in software design and multiagent systems.

Bill Rapp is an IBM Distinguished Engineer in the United States and Chief Architect of Watson for Healthcare. He is leading the technical team that is adapting Watson for use by the healthcare industry. Prior to this, Bill was the Chief Technology Officer for the IBM Healthcare and Life Sciences Division. Bill has 35 years of experience in system and software architectures, 12 of those in healthcare-related technologies.

Thanks to the following people for their contributions to this project:

LindaMay Patterson
International Technical Support Organization, Rochester Center

Mike McRoberts
IBM SPSS® Chief Architect

Carolyn Calzavara
IBM Worldwide Marketing Programs, Watson Solutions

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new IBM Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document, REDP-5043-00, was created or updated on October 1, 2013.

Trademarks