



---

## Highlights

- Cognitive systems provide intelligent predictive maintenance—the prevention of natural failures becomes the new normal
  - SparkCognition uses machine-learning algorithms and IBM Watson to exceed the accuracy of traditional methods and reduce false alarms
  - IBM Power Systems are designed for big data with innovation that optimizes the performance of SparkCognition
  - IBM FlashSystem storage drives up efficiencies and performance while driving down complexity and cost
- 

# Cognitive advantage at the speed of power

*IBM and SparkCognition drive industrial analytics*

To those in the industrial sector, the Internet of Things is much more than a buzzworthy label for work like turbine monitoring, aircraft systems maintenance and transportation logistics.

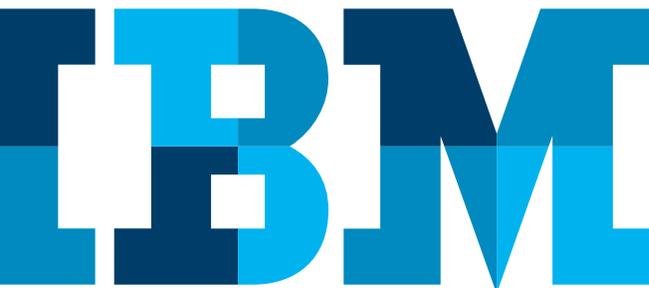
The industrial Internet with sensors embedded everywhere, used in conjunction with cognitive systems and applications, brings an intelligence and speed to machine maintenance that hasn't existed before.

Statistical modelling of event probabilities, the status quo, provides one level of responsiveness—indicating when a failure might take place. Yet conditions are always changing—weather is dynamic, temperatures fluctuate, stresses vary. Models can try to account for different scenarios, but they lack up-to-the-moment data from real-world events on which to fine tune their predictions.

Today, machine-learning algorithms and advanced systems can scour and analyze every parts manual, repair history, change log, diagram, outage report, manufacturer alert—essentially every piece of available documentation—absorbing and interpreting changing variables. They are now capable of providing actionable predictive intelligence informed by an ability to quickly understand meaning hidden in data.

Harnessing this thinking ability through cognitive computing can help leaders in the industrial sector transform the core of their business—making unscheduled maintenance during peak demand a thing of the past in the cognitive era.

Industrial sector organizations need the speed and innovation that SparkCognition® applications together with IBM® systems can deliver to anticipate and prevent natural failures.



## Failure is predictable

Among its security and reliability solutions, SparkCognition adds intelligence to the industrial Internet with SparkPredict™ and its cognitive machine prognostic capabilities that protect the health of machine assets, minimize maintenance costs and optimize schedules.

SparkPredict's state-of-the-art algorithms can be applied to a variety of assets, including wind turbines, aircraft, and railroad equipment to identify failures associated with issues such as gearboxes, blades, generators, mechanical systems and more.

SparkCognition uses a twofold approach to extrapolate dynamic insights from data. Powered with proprietary artificial intelligence algorithms and machine learning techniques, SparkPredict builds accurate and real-time health prognostics to deliver predictive maintenance insights.

Then, SparkPredict leverages IBM Watson™ and its immense corpus of knowledge as an in-context advisor to provide contextual diagnosis and remediation advice to fix the identified problem.

Since its triumph on the television quiz show *Jeopardy!* IBM has advanced Watson's capabilities and made it available via the cloud, while also opening the Watson platform to developers. SparkCognition has trained IBM Watson to understand complex machine maintenance questions. SparkPredict builds machine-scale pattern recognition

models to monitor mechanical systems, and together with Watson, presents answers and solutions based on the ever-expanding corpus of structured and unstructured data—in real time—for physical asset failure detection and prediction.

For example, when analyzing boiler feed pumps, SparkPredict and Watson can identify potential issues such as mechanical seal failures, rotor lock, bearing damage and water leaks. Warning and failure alerts are generated based on real-time operational behavior.

In the case of aircraft maintenance or turbine monitoring, technicians can ask natural language queries to find documents and other digital resources to address issues. SparkPredict can integrate with other systems such as diagnostic databases, maintenance records, and personnel records to help classify fault codes, recommend the right personnel and schedule maintenance in an optimal manner.

## The Power for intelligence

On-premises cognitive computing systems such as SparkPredict that must match the speed of IBM Watson in the cloud require the speed and I/O throughput necessary to quickly analyze and provide actions.

IBM Power Systems™ with POWER8® processors and IBM FlashSystem® storage are optimized for cognitive applications and the requirements to ingest and process a tremendous amount of data in an extremely short period of time. Running on Power Systems and IBM storage, SparkPredict and other SparkCognition applications can deliver computation, analytics and insight more efficiently than conventional infrastructure.

### SparkPredict and IBM Watson

**Improved Precision:** Sophisticated algorithms exceed accuracy of traditional methods and reduce false alarms

**Faster Results:** Automated model-building can shrink days or week-long model building processes to minutes

**Scalability:** Minimum human engagement enables large-scale deployments

**Adaptability:** A data-agnostic approach adapts to different asset sets to create the best prediction

**Continuous Improvement:** Self-learning and adaptive algorithms ensure that results continue to improve over time

*“In order to stay at the vanguard of the industry, we need systems that enable quicker queries over larger data sets, and IBM Power Systems is an ideal fit.”*

—Amir Husan, Founder & CEO, SparkCognition

Figure 1: SparkPredicts uses IBM Watson, built on IBM Power Systems, as an in-context advisor



Figure 2: IBM Power Systems and IBM FlashSystem storage create an optimum environment for analytics and insight.

### Processing power

Power Systems are designed for big data with innovation that optimizes analytics performance. They make it fast and easy to deploy services in the cloud, while also providing the ability to scale out or scale up with security built in.

IBM POWER8 processors have become the servers of choice for running today's most demanding, data-intensive workloads. POWER8 systems deliver computing power with 50 percent more cores, twice the simultaneous threads per core, and with more cache and I/O bandwidth versus the previous system generation—giving SparkPredict the optimum environment to deliver insight.

For analytics, users on IBM Power Systems receive two times more query results per core per hour than on x86 servers.<sup>1</sup> The huge advantage over x86 processors in multithreading and memory bandwidth make POWER8 a strong choice for high-end servers running memory-intensive workloads.

In addition, IBM Watson runs nearly two times faster with OpenPOWER acceleration innovations,<sup>2</sup> and OpenPOWER-based systems provide 45 percent better price-performance than x86-based servers.<sup>3</sup>

### Cognitive acceleration

Data has a value tied to its timeliness—the faster always-changing data is analyzed, the quicker actionable intelligence can be delivered; this is measured in microseconds, and every microsecond counts.

IBM FlashSystem storage systems provide the latency, scalable performance and operational efficiency needed to unlock insights from massive volumes of data. Flash memory also reduces energy consumption, requires less space than physical hard drives and provides superior performance—up to 50 times better performance than disk systems, four times greater capacity in less rack space, and the ability to scale to over two petabytes of flash capacity.

Where other systems have settled on one millisecond response times, IBM FlashSystem goes much further and faster with latency as low as 90 microseconds—which translates into increased industrial machine uptime, for example, because greater storage and server efficiency produces faster analytical insight and prediction.

Additionally, the combination of market-leading software defined storage capabilities and dynamic performance enable IBM FlashSystem to provide shared storage solutions for both highly specialized applications such as SparkPredict as well as any other workloads in the data center, lowering costs and complexity.

Optimizing throughput and enabling faster algorithms are important to meet service requirements. FlashSystem together with the POWER8 Coherent Accelerator Processor Interface (CAPI) technology enables accelerated computing. CAPI connects via Fibre Channel to FlashSystem arrays and allows programmers, as an example, to issue read and write commands from the application. As a result, this innovative use of flash technology can help eliminate 97 percent of code path length and save 20-30 cores per one million I/Os per second.

### Why SparkCognition and IBM

The amount of equipment-generated and sensor-generated data is massive; it is constant, 24/7/365 and knows no boundaries. For a human expert to keep up to date is impossible. However, cognitive applications, trained by human experts, with the power of advanced technology can turn data into accurate, actionable and timely insight.

IBM Power Systems with POWER8 and IBM FlashSystem storage are optimized for big data and cognitive computing; they create an infrastructure that can improve the performance of SparkPredict and Watson, while reducing complexity, power consumption and rack space.

When those systems are combined with the sophisticated cognitive algorithms of SparkPredict and the in-context advice of IBM Watson, the vast quantity of industrial sensor data becomes intelligent to improve operations, velocity and asset reliability.

### For more information

To learn more about IBM Power Systems, IBM FlashSystem and SparkCognition, contact your IBM representative or IBM Business Partner, or visit: [ibm.com/power](http://ibm.com/power), [ibm.com/storage/flash](http://ibm.com/storage/flash) and [ibm.com/cognitive](http://ibm.com/cognitive).



1. 2.03X more query results is based on IBM internal testing of a sample analytic workload; current as of October 20, 2015. Performance improvement figures are cumulative of all queries in the workload. Individual results will vary depending on individual workloads, configurations and conditions. IBM Power System S822LC; 20 core s / 80 threads, POWER8; 3.5GHz, 768 GB memory, DB2 10.5 / Ubuntu 14.04. Competitive stack: Intel x86 2-socket server with 36 cores / 72 threads; 2xIntel E5-2699 v3; 2.3 GHz; 768 GB; DB2 10.5 / RHEL 7.2
2. Watson using OpenPOWER acceleration innovations is at nearly 2X faster when running retrieve and rank workloads: <http://www.ibm.com/press/us/en/pressrelease/48075.wss>
3. Results are based on IBM internal testing of single system running multiple virtual machines with Sysbench read only work load and are current as of October 18, 2015. Performance figures are based on running 24 M record scale factor per VM. Individual results will vary depending on individual workloads, configurations and conditions. IBM Power System S822LC; 20 cores / 160 threads, POWER8; 2.9 GHz, 256 GB memory MariaDB 10.0.19. Ubuntu 14.04.03, PowerKVM 3.1 Competitive stack: HP Proliant DL380 Gen9; 24 cores / 48 threads; Intel E5-2690 v3; 2.6 GHz; 128 GB memory, MariaDB 10.0.20. Ubuntu 14.04.03, KVM Each system was configured to run at similar per VM throughput levels and number of VMs were increased for each system until total system throughput showed maximum throughput levels. Competitive pricing was taken from available web-based pricing.



---

© Copyright IBM Corporation 2016

IBM Corporation  
Route 100  
Somers, NY 10589

Produced in the United States of America  
June 2016

IBM, the IBM logo, ibm.com, FlashSystem, MicroLatency, Power Systems, POWER8, and Watson are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NONINFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.



Please Recycle