IBM Predictive Maintenance and Quality (Version 2.0)
IBM Redbooks Solution Guide

The IBM® Predictive Maintenance and Quality solution (PMQ) uses information collected about products, processes, and assets to optimize maintenance schedules, production processes, and product quality. This integrated solution brings a variety of benefits to an enterprise:

- Predict the *best* time to perform maintenance on a monitored asset
- Uncover process deficiencies that can affect product quality
- Identify the root causes of asset failures

These capabilities are critical to high-level business objectives, such as capital planning, operations management, inventory optimization, and more. Enterprises quickly learn that just because the manufacturer recommends that an asset be serviced at certain times, the manner in which the asset is actually used, day-to-day, can require that it be serviced on a faster, or slower, schedule.

The PMQ solution (see Figure 1) builds upon the fact that today's production assets are often instrumented to collect and transmit vital data such as operating temperature, duration, error codes, and so on. Using innovative IBM integration and analytics products (IBM Integration Bus, SPSS®, Cognos® Business Intelligence, and so on), this data can be modeled to predict when an asset is most likely to fail or falter. This IBM Redbooks® Solution Guide describes the PMQ solution, which in Version 2.0 adds new features such as the Quality Early Warning System, to better detect and predict problems before they happen, and the Analytics Solution Foundation, which adds flexibility by enabling clients to adapt the solution components to their own use cases.
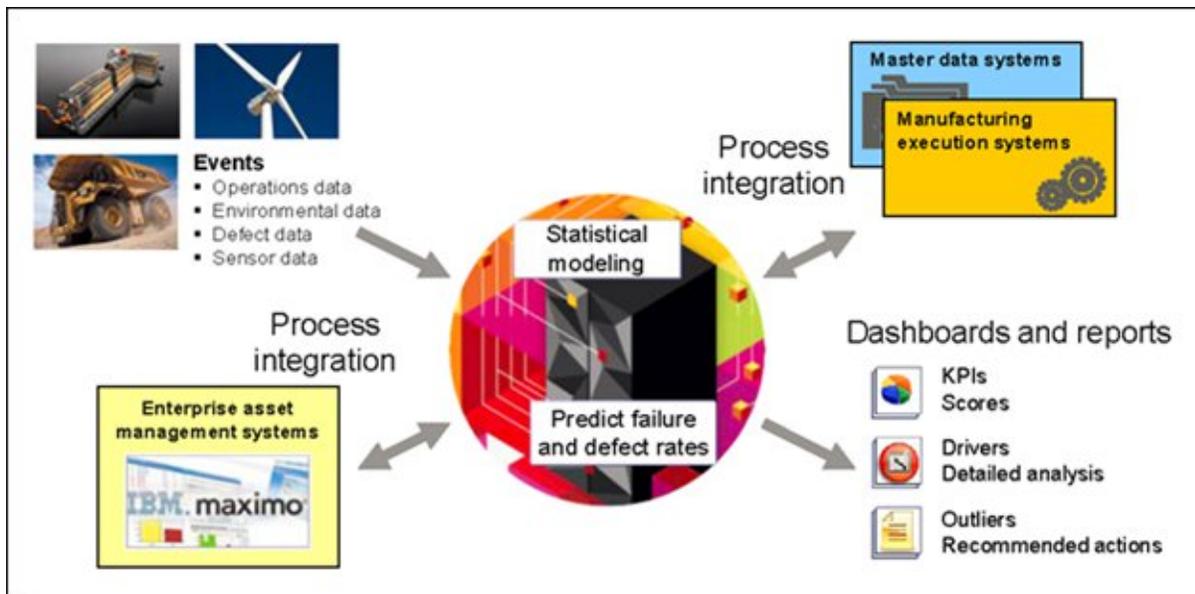

Figure 1. IBM Predictive Maintenance and Quality: Conceptual information flow

Did you know?

IBM consultants can work with you to assess your current asset maintenance and product quality issues and, using your existing historical data, demonstrate how the PMQ solution can help. PMQ includes integration with the IBM Maximo® Asset Management system and can be integrated with other asset management systems, such as SAP PM.

Business value

Predictive maintenance can be applied in any industry that relies on monitored production assets. These initiatives are increasingly seen in sectors, such as manufacturing, mining, drilling, farming, and security, covering everything from cars and trucks to wind turbines, cranes, and oil platforms.

The business value of predictive maintenance programs can be far-reaching. The benefits vary from sector to sector, but here are some of the top reasons enterprises are investing in these solutions:

- **Clearer correlation between processes and asset health**

  By enabling companies to relate the condition of monitored assets to the processes in which those assets are used, PMQ can help identify process changes that will extend the life of the assets. Managing asset information, such as operating conditions, failure episodes, critical runtime parameters, and maintenance records, is critical in any manufacturing operation.

- **Improved operational efficiency of individual assets**

  Every manufacturing asset needs to be kept and used under certain conditions for optimal performance and to increase its lifespan. Predictive maintenance helps companies identify the parameters that most affect performance and which, when optimized, will lead to superior operational efficiency.

- **Improved asset reliability and availability** *and lower costs*

  Predictive maintenance considers not only the manufacturer's recommended maintenance routines but also evaluates how an asset is truly being used, including its runtime performance characteristics, whether past repairs were conducted sooner or later than expected, and so on. The goal is to avoid costly downtime by ensuring the asset will always be available, and in optimal condition.

- **Higher-quality goods and services**

  Improving quality is an ongoing goal of any good business. By adopting predictive maintenance analysis, the enterprise can keep its manufacturing assets in optimal working condition and help ensure that every product it ships to market is of the highest quality.

- **Fewer product recalls**

  Products are recalled for many reasons and the result can be expensive for the enterprise and consumers. Predictive maintenance can reduce the probability of product recalls by recognizing manufacturing problems as soon as they are reported from monitored assets, and then issuing immediate recommendations and work orders to resolve the identified problem.

- **Greater customer satisfaction**

  High-quality goods and limited product recalls lead to greater overall customer satisfaction and more repeat business. And as has been shown, predictive maintenance can help in each of these areas by anticipating major problems *before* they occur.

Solution overview

The IBM Predictive Maintenance and Quality solution has these high-level logical components:

- Event sources: When monitored systems and devices submit their critical operating data, the events containing the data are typically routed through supervisory control and data acquisition (SCADA) systems or other, similar control systems that act as interfaces with the remote equipment, including distributed control systems (DCS). Event sources can also include production databases or any other, intermediate databases where event data is stored for periodic batch-loading, or external systems that send continuous, high-volume streaming event data.

- Integration bus: The integration bus connects the internal and external components of the solution. Its role is to transform the events that are received from monitored assets so they can be interpreted by other components, such as the analytics engine and scoring services, to obtain the predictive scores and recommended maintenance actions that stem from the received events.

- Analytic data store: The solution's analytic data store is used to store master data, device events, recommended maintenance activities, and so on.

- Analytics: The analytics layer looks for trends and potential maintenance issues using scoring models that are based on historical maintenance data and other inputs, such as key performance indicators (KPIs), for each monitored asset. The models produce scores that predict the future of a given asset and any maintenance recommendations regarding it.

- Enterprise information systems: Output from the PMQ solution can be sent to any downstream system that can utilize it. These include Enterprise Asset Management (EAM) systems, from which maintenance work orders can be generated; Master Data Management (MDM) systems, where the base information about the monitored devices must be kept up-to-date, and even Manufacturing Execution Systems (MES) that manage overall plant floor operations. In some cases, these systems also provide input to the PMQ solution.

- Reporting: The solution's business intelligence features include a customizable set of live dashboards and reports that provide different views of device events, predictive scores, and the maintenance recommendations that are generated based on those scores. Together, these views can provide valuable insights into the overall state of a complex manufacturing operation.

- Quality Early Warning System (QEWS): This component is new in Version 2.0 of the PMQ solution. It supports two important monitoring and analytical functions:

  o Manufacturing quality: QEWS can detect unfavorable changes in component quality based on data gathered from manufacturing processes and reports from customers and suppliers.

  o Warranty claims: QEWS can detect new warranty-related issues by monitoring KPIs such as the number of random failures in particular pieces of equipment and the wear-out rates of products once they are in use by customers. QEWS enables warranty analytics based on the actual date of sale of the product, the date of production of the end product, and the manufacturing date of the resource that is used in an end product.

- Analytics Solution Foundation: This new component provides additional capabilities in PMQ Version 2.0, including the ability to modify the data model, event processing behavior, and orchestration flows. This helps when you need to customize the solution to meet specific business needs.

- Installer: The installer enables automated installation and configuration of the Predictive Maintenance and Quality solution software.

These logical components of the solution are depicted in Figure 2.
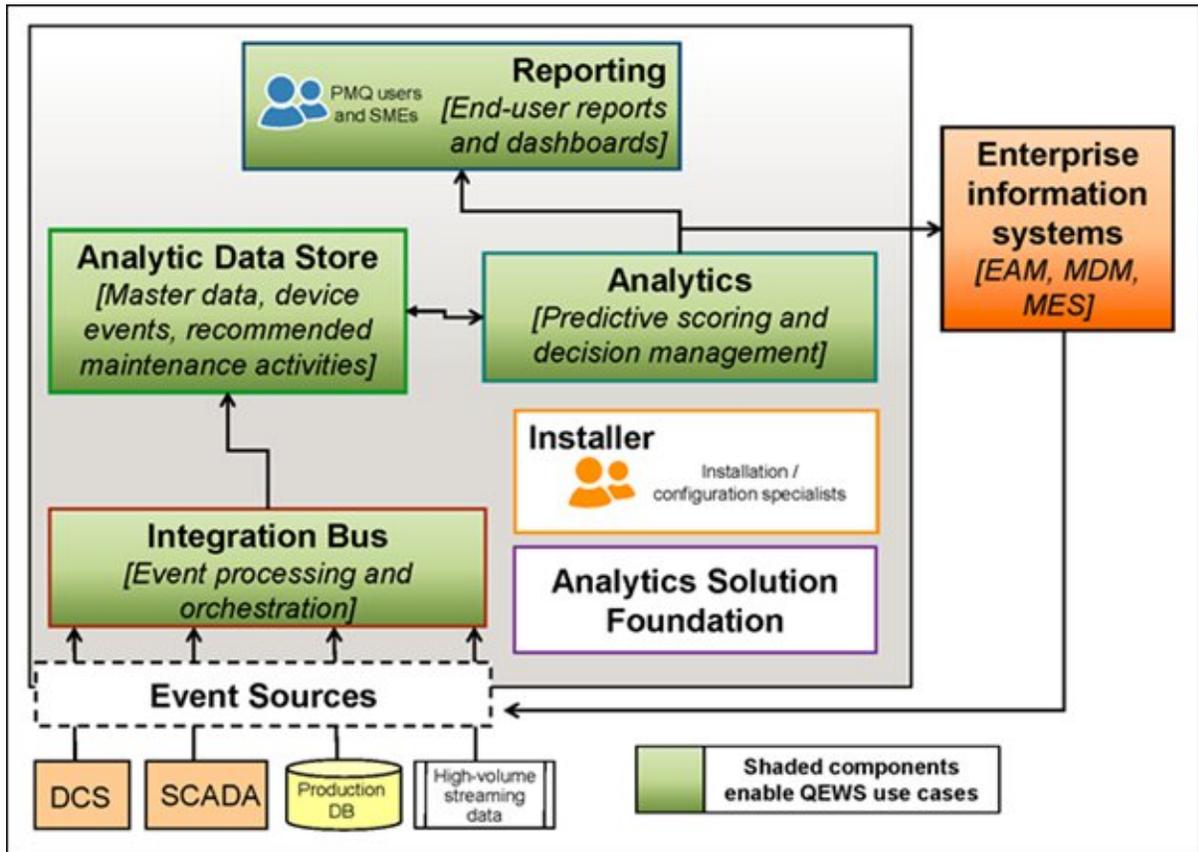

Figure 2. High-level logical components of the Predictive Maintenance and Quality solution

The physical components that make up the solution's logical components (shown in Figure 2) are described in more detail in Table 1.

Table 1: Physical components of the Predictive Maintenance and Quality solution (Part 1 of 2)

Component	Description
Master data management	This component is responsible for defining and managing the lifecycle of product master data: defining it, enriching it, and making it accessible to other processes within the solution. This component also manages any updates to the product master data over time.
Relational database server	The relational database server provides persistent storage for solution-related data, ensuring its integrity during access and storage operations. It is also where analytic data is stored, including a pre-built data schema for storing quality, machine, and production data, plus configuration information.
OPC adapter	This adapter converts OLE for Process Control (OPC) data into events that can be received by the integration bus (messages).

Table 1: Physical components of the Predictive Maintenance and Quality solution (Part 2 of 2)

Component	Description
Integration bus	IBM Integration Bus is responsible for classic enterprise service bus (ESB) functionality, including protocol conversion, message mediation, and transformation. It also includes various handlers to process inbound events, validating them against master data and invoking integration work flows, as required.
Integration work flows	Integration work flows are responsible for processing the inbound event data that is analyzed by the solution's predictive models. The flows run on Integration Bus and support functions, such as event aggregation, master data loading, and so on. Some solution integrators might create custom flows for integration with external systems.
Business intelligence	The business intelligence component is the platform for most of the visualization aspects of the solution and is provided by IBM Cognos Business Intelligence. It offers dashboards and reports through which users can learn the status of monitored assets and systems: • Site Overview: A summary of all monitored assets at all monitored sites, showing the KPIs that most affect asset maintenance. • Equipment: Numerous equipment-related reports are available. For example, users can view the range of recent measurements for a monitored device (Equipment Control), see when a device is falling outside of its predetermined operating limits (Equipment Outliers), and call up the recorded measurements for a device for a specific period in the past (Event Type History). • Product Quality: An overview of defects and inspection rates of a monitored asset or its components. • Master Data: A list of the master data for resources, materials, processes, production batches, and profile variances. • Material Usage: An overview of material usage by production batch. By correlating production batches with defects to material usage, a company can trace the impact of defective materials.
Statistical analytics and asset models	Statistical models are used for predictive analysis and help predict likely maintenance-related outcomes based on the current facts (readings from devices, applicable master data, updated KPI calculations, and so on). Users can create their own models specific to the enterprise or choose from among a set of useful, pre-built models, for example: • Health Score: Measures the likelihood that an asset or process will fail based on historical defect data, operational information, and environmental data • Lifespan Analysis: Estimates a device's remaining lifespan when functioning in a real-world scenario
Decision management	The decision management component provides a workbench for analyzing incidents, based on established business and maintenance rules, and determining any needed actions.
Installers	The PMQ solution includes installer components for both the software package itself and the content, such as configuration files, that must be loaded into the system.
Analytics Solution Foundation	The Analytics Solution Foundation is a new feature in Version 2.0 of the PMQ solution. It is essentially an additional solution layer dedicated to accommodating customer-specific needs and customizations, such as the use of database servers other than IBM DB2, custom master data models, custom event processing, and so on.

Solution architecture

From a pre-production perspective, the Predictive Maintenance and Quality solution starts with the loading of master data, such as a list of monitored devices and the types of measurements they produce, into the analytic data store. This can be achieved with CSV files, by using IBM InfoSphere Master Data Management (bundled with the solution), or through updates from external sources, such as IBM Maximo and similar Enterprise Asset Management (EAM) systems. Also loaded are the key performance indicators (KPIs) or thresholds that establish whether a certain asset is operating as intended or needs maintenance.

Events from field-level systems can arrive in batches or in real time using various protocols and formats, such as OLE for Process Control (OPC). OPC event sources are integrated using an external adapter, such as ILS deviceWise, or the IBM Integrated Information Core (IIC) OPC Adapter. After device events are received, the event processing component transforms them into the format required by the solution.

The analytic data store is implemented as a database using IBM DB2®. The information it contains is used for predictive scoring, in which a mathematical model developed in IBM SPSS Modeler puts a numerical value on the relative likelihood of device or component failure. These predictive scores are then passed to IBM SPSS Decision Management, which uses a predefined set of rules to determine appropriate actions to take in response to the various scores. For example, if immediate maintenance is needed, a work order can be created in the company's EAM system, such as IBM Maximo.

The information in the analytic data store is viewable in a variety of ways using IBM Cognos Business Intelligence, through which PMQ users can generate standard or customized dashboards and reports about the assets or processes that interest them. IBM Integration Bus is used to orchestrate event processing. Integration work flows are included for SPSS interaction, work order creation in Maximo, and integration with IBM Master Data Management, if it is used.

Installation scripts are included to support the installation of PMQ on a standard topology (additional deployment topologies are possible depending on the company's specific requirements). Version 2.0 of the PMQ solution includes the new Analytics Solution Foundation, which is described in detail elsewhere in this document.

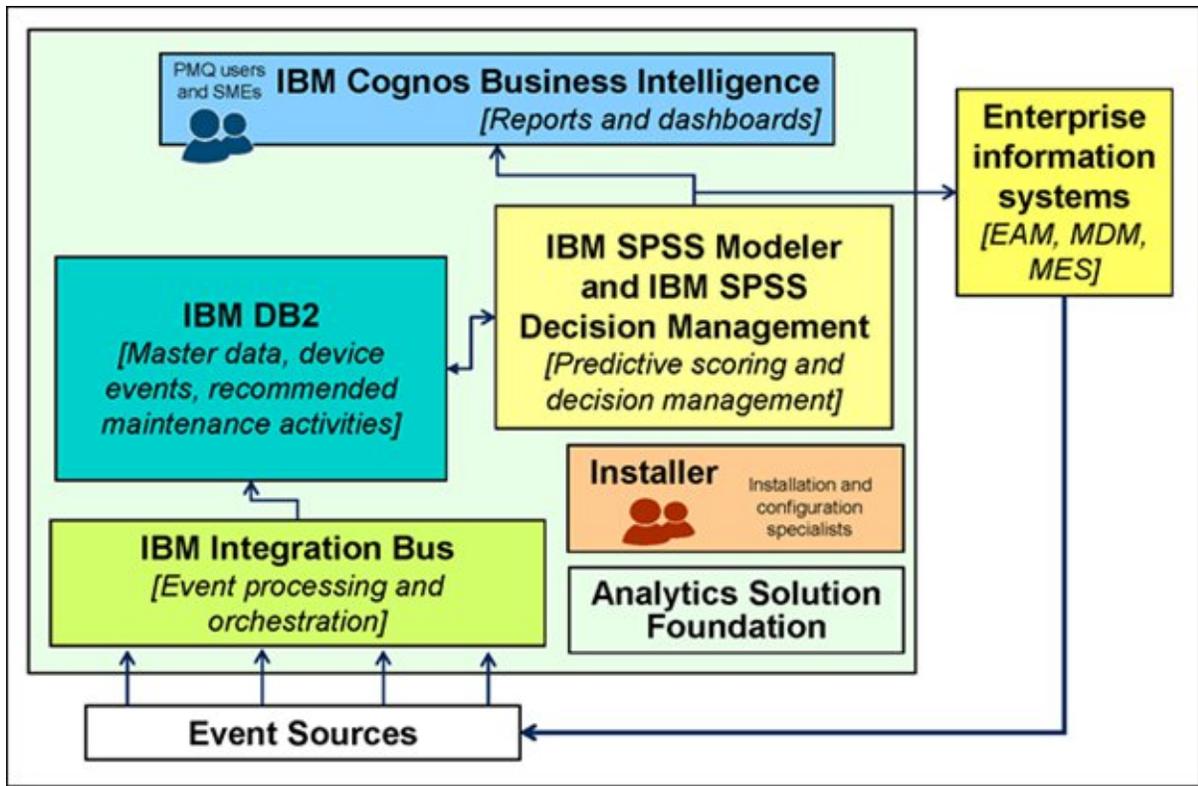The solution architecture is illustrated in Figure 3.

Figure 3. High-level architecture of the Predictive Maintenance and Quality solution

Usage scenarios

Four types of human actors are typically involved in using PMQ: the configuration specialist, runtime users, analysts, and report specialists. Each actor initiates particular use cases for which they are responsible, as noted here.

Load master data (configuration specialist)

Master data must be loaded into PMQ before events can be processed. The solution includes pre-built Integration Bus flows that load master data from files. Together, these flows are known as the flat file API.

Store and process events (runtime user)

The event processing flow removes events from the input channels and stores them in the solution's analytic data store. These events, which can contain many observations, are recorded in appropriate database tables corresponding to the resource (a resource can be either a device, such as a transformer, or an operator, such as the driver of a truck). If an event is received from a device that has not yet been loaded into the solution's master data, that event will not be processed and a message explaining the failure will be added to the event processing error log file. This use case is illustrated in Figure 4.

Execute analytical model and generate score (runtime user)

Scoring is a key part of predictive maintenance and involves the use of predictive models that use historical data to determine the probability of certain future outcomes. For example, a model can be created based on historical data regarding transformer temperature, current load, and occurrences of failure. In PMQ, predictive models are created in SPSS Modeler. The models are then deployed to SPSS Collaboration and Deployment Services, where they are available to be called as scoring web services.

Execute decision based on score (runtime user)

The solution's decision management feature allows it to recommend a specific action based on a particular predictive score. This analysis can be performed using a Integration Bus flow, but PMQ uses IBM SPSS Decision Management to control the decision-making process and the thresholds for taking action based on rules. With this component, rules can be authored, tested, optimized, and deployed to SPSS Collaboration and Deployment Services, where they become callable as decision management web services.

Generate recommendations (analyst)

SPSS executes pre-built analytical models, resulting in scores. KPIs are analyzed by the system on a continuous basis. In response to the scores and the current KPI values, the PMQ system generates recommendations using the pre-configured business rules.

Load events (real time) (runtime user)

Events from a device can be loaded into PMQ in real time using one of the connectivity options provided by Integration Bus, such as a web service or a message queue. A custom Integration Bus flow is required to transform, or map, the event information from the format used by the device-monitoring equipment to the format used by PMQ.

Load events (batch) (runtime user)

Certain devices lack the ability to supply events in real time so their events are instead loaded as batch files. The solution includes appropriate integration work flows for this purpose.

Create work order (runtime user)

IBM Maximo is an enterprise asset management system that supports the creation of work orders through a self-generated web service. PMQ includes an Integration Bus flow that can call the Maximo web service to create a work order. This flow is triggered when a recommendation for this type of action is received.

Extend solution to support custom master data or non -DB2 database servers (Configuration specialist)

The new Analytics Solution Foundation makes it possible to extend Version 2.0 of the PMQ solution to support custom master data tables and schema, or to use database servers other than IBM DB2, when dictated by specific business needs.

Detect random failures that could lead to warranty claims (Analyst)

The new Quality Early Warning System (QEWS) component in Version 2.0 of the PMQ solution makes it possible to detect higher-than-anticipated instances of random manufacturing failures, with the goal of identifying and fixing the problems before they result in increased warranty claims. Thresholds can be established based on product reliability goals (for example, the products in the field cannot exceed a specified failure rate) or financial liability goals (for example, the cost of satisfying predicted warranty claims must not exceed a specified total amount).

Detect wear-out scenarios involved in warranty claims (Analyst)

Again thanks to the new Quality Early Warning System component, PMQ 2.0 can detect higher-than-anticipated reports of product wear-out. Due to the serious nature of wear-out failures, the system can raise alerts whenever it detects evidence of unexpected wear-out rates, regardless of how many product units contributed to the detection.

Integration

The PMQ solution is often integrated with systems in these broadly defined categories (see Figure 4):

● Production data sources

These systems are not part of the PMQ software bundle but can be important in getting certain kinds of event data into the solution. For example, OPC data sources depend on this type of interface. OPC events are retrieved using third-party adapters, such as IBM IIC or ILS deviceWise.

● Enterprise asset management (EAM) systems

Depending on the need, PMQ can be integrated with EAM systems to both initiate work orders and retrieve updated master data related to the execution of those work orders. Support for one such system, IBM Maximo, is included with the solution in the form of software development kits (SDKs) to guide the integration. It is also possible to integrate with custom EAM systems, such as SAP PM.

● Master data systems

The solution also integrates with external master data systems, such as IBM InfoSphere® Master Data Management, which is included in the PMQ software bundle. Optionally, and depending on the enterprise's requirements, PMQ also can be used with other IBM products, such as IBM Business Process Manager or IBM Case Manager, using the protocols and interfacing options supported by those products.
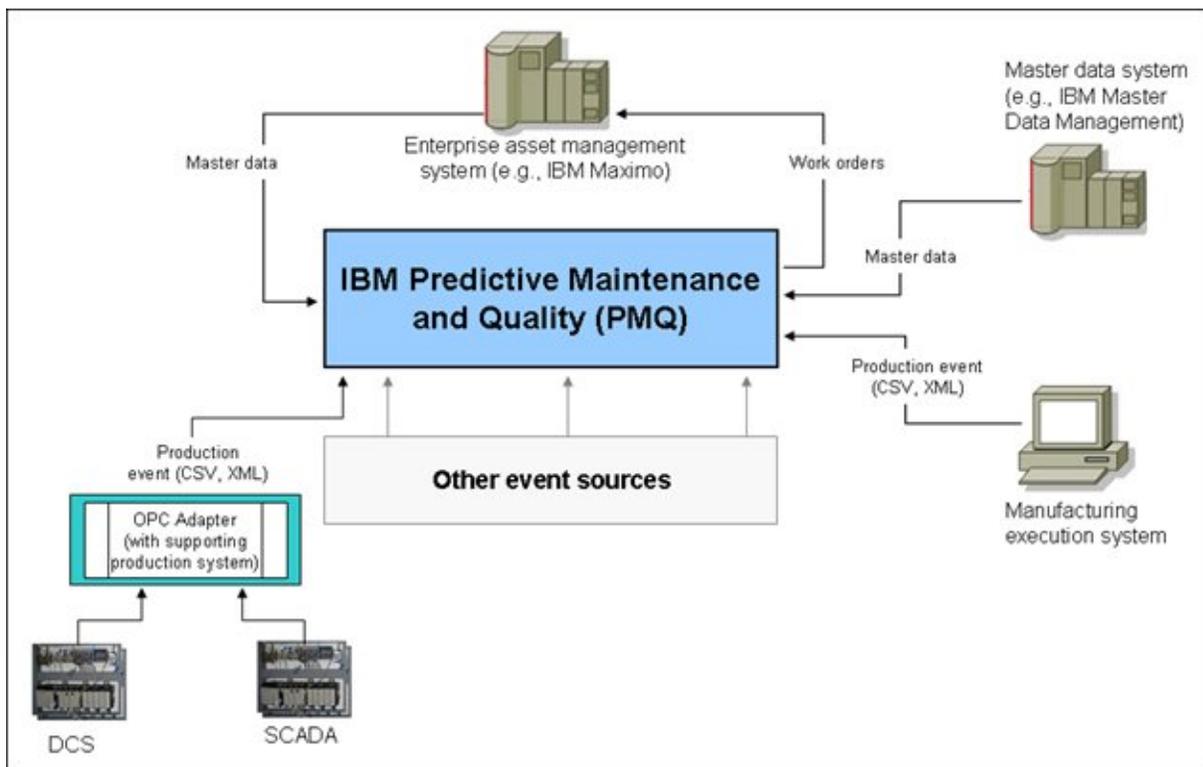

Figure 4. The nature of the PMQ solution means it integrates with multiple external systems

Supported platforms

PMQ is supported on Red Hat Enterprise Linux (RHEL) Version 6.3. From a hardware perspective, it can run on IBM System x® servers with quad core processors or other compatible 64-bit x86 servers.

Ordering information

The Predictive Maintenance and Quality solution is purchased through IBM Passport Advantage® using a variable pricing model. For more information, visit the PMQ website and contact an IBM representative:
http://www-03.ibm.com/software/products/en/predictive-maintenance-quality

Related information

For more information, see the following websites:

- IBM Predictive Maintenance and Quality website:
  http://www-03.ibm.com/software/products/en/predictive-maintenance-quality

- IBM Predictive Maintenance and Quality Technical Overview (Version 1.0):
  http://www.redbooks.ibm.com/redpieces/abstracts/redp5035.html

- IBM Predictive Maintenance and Quality Information Center:
  http://www-01.ibm.com/support/knowledgecenter/SSTNNL_2.0.0/

- IBM Offering Information page (to search on announcement letters, sales manuals, or both):
  http://www.ibm.com/common/ssi/index.wss?request_locale=en

  On this page, enter "IBM Predictive Maintenance and Quality", select the information type, and then click **Search**. On the next page, narrow your search results by geography and language.

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 give 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. 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.

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 measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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 was created or updated on November 4, 2014.

Send us your comments in one of the following ways:
- Use the online **Contact us** review form found at:
  ibm.com/redbooks
- Send your comments in an e-mail to:
  redbooks@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at http://www.ibm.com/redbooks/abstracts/tips1130.html .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at
http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Cognos®
DB2®
IBM®
InfoSphere®
Maximo®
Passport Advantage®
Redbooks®
SPSS®
System x®
WebSphere®

The following terms are trademarks of other companies:

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other company, product, or service names may be trademarks or service marks of others.