precisely

Top Use Cases for IBM i Data in Splunk: IT Operations Analytics

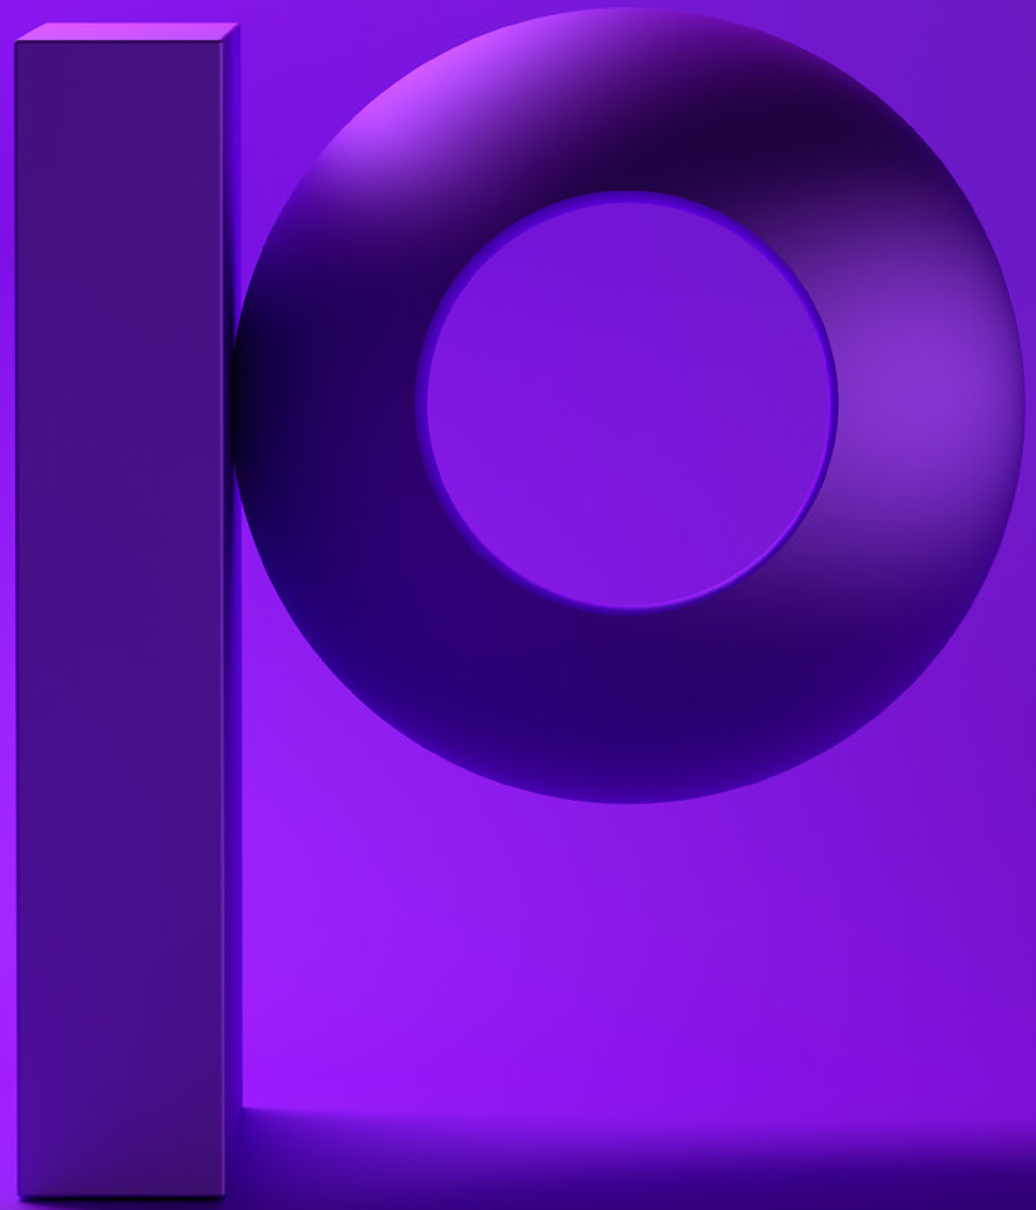

Introduction

This is Part 3 in a 3-part series on the use cases for using IBM i Data in Splunk.

Today's computing environments are a complex arrangement of many hardware components and several software layers, and it is vital that each of these parts functions to the best of its ability. In the case of customer-facing systems, the failure of one element can impact hundreds, thousands or even millions of users.

Ensuring the optimal performance and availability of IT systems and applications, while also controlling IT costs and maximizing the use of critical resources, has become a significant challenge for IT professionals.

For decades, organizations have collected different types of data and monitored systems to support better operations and address issues.

With today's powerful IT Operations Analytics (ITOA) platforms, we are able to unlock the value that has been hidden in the detailed logs that are generated by enterprise systems. However, as with most technology solutions, the devil is in the details.

1. STRUCTURE. Machine data usually comes in files of semi-structured, unformatted data. Of course, each system or "machine" has its own way of logging data, which makes this task even more challenging.

2. SEQUENCE. Machine data is mostly sequential. This means, to get the most meaningful insights, you must look at the entire chain of events.

3. VOLUME. Machine data volumes can be massive. With hundreds of servers and other types of systems, dozens of applications and logs recording every step of a given event or transaction, volumes can easily reach terabytes of data per day.

4. TIMING. We all know data loses value over time, but when it comes to operational intelligence, the value diminishes exponentially. That's why both real-time data and a researchable historical record for analysis is critical for machine data.

These factors alone can make an operational intelligence initiative far more complicated than a traditional business intelligence project. As a result, operational intelligence requires an unconventional approach involving ITOA powered by multiple, simultaneous streams of machine data, correlated together and possessing a searchable, continuous machine record.

Breaking Down (Most) Silos with Splunk

Unless a DIY approach is for you, an end-to-end solution like Splunk can spare you the pain of setting this up from scratch – so you can spend your time analyzing the insights that will address your operational issues and security mandate and drive your business.

Many organizations, across all industries, are turning to Splunk solutions to meet their ITOA challenges. Splunk integrates log and other machine data from across an IT infrastructure, eliminating silos of legacy tools with partial views that need to be cobbled together before it's possible to understand and respond to issues. With Splunk, users can examine their data in depth and in real time – all in one place – so they can predict, prevent and fix problems fast, while also reducing costs through data center optimization.

Splunk automates the collection, indexing and alerting of machine data that's critical to your operations; unfortunately, not all systems are as easy to integrate into Splunk as others. In particular, IBM i systems have been left out of most Splunk environments because the machine data generated by these systems is unique and requires specialized skills to work with.

Where to Find IBM i Insights for ITOA

The next page contains the primary sources of information logged on IBM i, and the valuable insights they contain – if you can access and make use of them.

System Audit Journal

This contains a variety of information focused on supporting security information and event management. Security events include things such as changes to system values, object authorities, profiles, authorization lists, access attempts, power user activity, transfers of objects to production libraries, actions on spooled files, adopted authorities, exit points, access of sensitive objects, and more. Essentially anything that is happening within the system environment that might impact security is contained within the System Audit Journal.

Operator Messages - QSYSOPR Message Queue

QSYSOPR messages may be used to indicate an action required by an operator or simply to provide the system operator notification of changes to the environment. Messages alerting the operator to a condition that needs attention may include loading a next tape volume, servicing a printer, saving a version of QHST, replying to a message to cancel a job, or allowing a job to exceed its spooled printer output limit. Operator notifications, for example, include messages such as jobs being held or released by a user.

System and Application Messages - QSYSMSG Message Queue

This optional queue contains messages primarily related to end user connections that are failing along with the reason for the failed connection. The system sends certain messages to QSYSMSG, to QSYSOPR, or to both QSYSMSG and QSYSOPR, depending on the system reference code (SRC) sent with the message and whether the SRC is being logged with critical message handling.

Accounting Journal

This contains information related to job and printer activity. There is an entry for each job completion/termination which contains details related to the job's execution, including resource usage such as processor time. Information about printer and spool activity is also recorded in the accounting journal.

QHST History Log

The history (QHST) log consists of a message queue and a physical file known as a log-version. Messages sent to the log message queue are written by the system to the current log-version physical file. QHST contains a high-level trace of system activities such as system, subsystem, job information, device status, and system operator messages. Records stored here can be complex in nature and require some re-formatting for usability in reporting.

Collection Services and Logs for Performance Data

The IBM i platform can be configured to collect an abundant amount of performance information. When performance data collection services are in effect, the IBM i operating system logs relevant performance metrics to a Management Collection object. The metrics are extracted from the object and stored into Db2 tables either in real time or at a later point. These Db2 tables may optionally be journaled as well. Performance information can be extracted from either the Db2 tables or the journals and analyzed to understand the performance of the operating systems, system components, and executing applications.

Including IBM i in Splunk for a True 360° View

Spunk offers a powerful solution for organizations needing to comply with industry and government regulations. However, Splunk does not natively collect essential security and compliance data from the IBM i platform, leaving a significant blind spot and vulnerability. That's where Precisely Ironstream comes in.

Ironstream seamlessly feeds IBM i logs to Splunk, ensuring that critical machine data for the entire IT landscape is available in a single tool. Splunk turns mountains of incomprehensible data into visual insights.

Together, Ironstream and Splunk help organizations achieve satisfactory security and compliance audits, and provide security event tracking, real-time monitoring of security events, automated reporting, and complete visibility into the health and security of all systems in the enterprise.

precisely

Precisely is the global leader in data integrity, providing accuracy and consistency in data for 12,000 customers in more than 100 countries, including 90 percent of the Fortune 100. Precisely's data integration, data quality, location intelligence, and data enrichment products power better business decisions to create better outcomes. Learn more at www.precisely.com.

www.precisely.com