

The Ultimate Guide to IBM i Machine Data Analytics



Introduction: What is IBM i Log Data?

IBM Power Systems running IBM i are trusted, reliable powerhouses tailored for the needs of companies that process a high volume of transactions, such as financial, healthcare, retail, and logistics organizations. Over its lifespan, the IBM i has been known as the iSeries, AS/400, System/38 and more – names you still hear from long-time users of the platform. There are several different data sources that are available within the IBM i environment that can be leveraged to provide insight into the operational health of the system and applications, as well as providing visibility into security and compliance issues. Most of these come in the form of journals, data base files, or message logs recorded by components of the IBM i OS. They provide information about security, operational issues, system resource utilization (CPU, Storage, I/O), and performance.

The data sources provide valuable insight into performance, availability, health, and security of the system and its applications. Organizations understand and see the value of incorporating IBM i machine data into their analytics processes. With all the data sources and volume of data available on the IBM i platform, each use case becomes an exercise in identifying the right set of information needed to meet the needs of the organization. Let's look at the different data sources and how they can be used for two primary use cases: IT Operational Analytics (ITOA) and to address Security and Information Event Management (SIEM challenges.





Security and Information Event Management (SIEM)

SIEM technology aggregates and provides real-time analysis of The IBM i operating system can produce an abundant amount of security alerts using event data produced by security devices, network information related to security. Typically, every security event – application logons, object accesses, and more may be recorded in some form. The infrastructure components, systems and applications. A primary function of SIEM is to analyze security event data in real-time for internal and external problem is the volume of security information and how to make sense of threat detection to prevent potential hacks and data theft. This typically it. Events related to potential security breaches, including logon failures includes user behavior analytics (UBA) – understanding user behavior and or unauthorized object access attempts, must be continuously monitored how it might impact security. SIEM technologies also collect, store, analyze, and analyzed to understand threats and breaches. Inevitably there will be and report on data needed for regulatory compliance to ensure that audit inadvertent attempts to access an object by an unauthorized user, or a logon failure due to an incorrect password being entered – but separating requirements are met as dictated. this "noise" from real threats requires deeper analysis of patterns over time.

Understanding potential security threats and preventing security breaches has become paramount to just about every IT organization. The image and financial damage caused by a security breach can cripple an organization. Compliance initiatives further increase the burden put on most organizations to get a better handle on security.



Forwarding IBM i security events into an analytics platform enables indepth analysis of access patterns, authorization failures, user behavior, and other security-related information to detect potential security threats and prevent breaches.





Real-Time Alerting

User Activity Monitoring

Dashboards

Reporting

File Integrity Monitoring

System & Device Log Monitoring

Log Retention

precisely.com | 877 700 0970



0

IT Operational Analytics (ITOA)

According to Gartner Group, IT Operational Analytics (ITOA) is a market for solutions that bring advanced analytical techniques to IT operations management use cases and data. ITOA solutions collect, store, analyze, and visualize IT operations data from other applications and IT operations management (ITOM) tools, enabling IT operations teams to perform faster root cause analysis, triage, and problem resolution.

The IT infrastructure of most large organizations is comprised of different systems and platforms each with unique toolsets requiring domain specific skills to extract and leverage value from those tools. Forwarding IBM i machine data and logs into an analytics platform opens a variety of opportunities to enhance IT operational intelligence and gain better organizational efficiency, without having to rely solely on platform specific tools. For example, the following ITOA use cases can be supported:

- Monitoring job activity
- Monitoring physical and virtual processor utilization
- Monitoring network activity
- Monitoring database performance
- Monitoring memory and disk utilization

Free Space by Unit



Example of IT Operational Analytics for disk space utilization and performance.



0

IBM i Data Sources

These are the primary sources of information logged on IBM i:

- System Audit Journal
- Operator Messages QSYSOPR Message
- Oueue System and Application Messages QSYSMSG Message Queue
- Accounting Journal
- OHST History Log
- Collection Services and Logs for Performance Data

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 arewritten 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 reformatting 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 systems 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.



IBM i Data Elements for Analytics

Listed below is a sample of the data elements that can be captured from IBM i machine data and logs. With this information available for analytics, organizations can address many ITOA and SIEM use cases.

Security Information

- User profiles
- System values
- Object attributes and authorities
- Authorization lists
- Job descriptions
- Commands
- Active jobs
- Spool files
- Changes to values, authorities, profiles, authorization lists
- Access attempts for authentication or object access
- Sensitive object access

Job Information

- CPU used Disk I/O detail
- Page faults
- I/O pending faults
- Socket sends/receives
- Stream file, directory and Symlink reads
- Stream file writes
- Seize/Wait time
- Communication puts/gets

Physical Processor Information

- Time used
- Owning partition

CPU Information Per Virtual CPU

- Time used
- Number of CPUs active



ρ

Virtual Processor Information

- Status, time active, time used
- Configured/Uncapped available time
- Instruction count

Memory Pool Information

- Database & non-database faults
- Job transitions size
- Disk I/O stats
- Pages aged and stolen

Disk Information

- Reads/Writes
- Disk Capacity and Space Availability
- Disk Busy
- Disk Response Times

TCP Communications

- Detailed stats at the datagram level
- Network fragmentation information





Gain Valuable Insight from IBM i Log Data

Analytics platforms make it simple to collect, analyze, and act upon the Information from key IBM i data sources, logs and message queues can untapped value of the big data generated by technology infrastructures, be normalized and forwarded to analytics platforms. security systems, and business applications – providing the insights to drive operational performance and business results. They typically collect and This enables analytics systems to provide total visibility into the IBM i index log and machine data from any source and provide powerful search, environment and the applications it supports. There is no need for special analysis, and visualization capabilities to empower users of all types. The knowledge and expertise to correlate IBM i data with that coming from other platforms. The automatic collection and forwarding of information problem is that most of these analytics platforms have no good connectors from IBM i data sources demystifies the complexity of those data sources or mechanisms to get to machine data produced on IBM i platforms. and enables your organization to easily gain operational and security insights into the IBM i environment and address operational and security issues across your entire IT infrastructure.

Precisely Ironstream[™] software is the industry's most comprehensive solution for automatic forwarding of machine and log data to analytics platforms. In the past, the IT shops needed significant expertise and resources to extract data from the different data sources on an IBM i system. With Precisely's log forwarding solutions, it is easy to feed data to Splunk and other analytics platforms for visibility into activity on all systems — including the IBM i — from one integrated user interface.



Ю



About Precisely

As a global leader in data integrity, Precisely ensures that your data is accurate, consistent, and contextual. Our portfolio, including the Precisely Data Integrity Suite, helps integrate your data, improve data quality, govern data usage, geocode and analyze location data, and enrich it with complementary datasets for confident business decisions. Over 12,000 organizations in more than 100 countries, including 93 of the Fortune 100, trust Precisely software, data, and strategy services to power AI, automation, and analytics initiatives. Learn more at www.precisely.com

www.precisely.com