XPLG

XpoLog Data Sheet

New frontier in big log data analysis and application intelligence

Technical white paper
May 2015

XpoLog is a data analysis and management platform for Applications IT data. Business applications rely on a dynamic heterogeneous applications infrastructure, such as cloud and virtualization.

The XpoLog data analysis and management platform helps collect, manage, analyze, and search any log in the environment.

XpoLog collects, indexes, and correlates data in a searchable repository from which it can generate graphs, reports, alerts, dashboards, and visualizations. The technology analyzes the environment automatically, and helps uncover hidden values in the environment to increase application quality and availability.

XpoLog delivers applications and operational intelligence to application owners, operations, IT administrators, and other stakeholders across the IT life cycle.

Visit our online knowledge base – **http://wiki.xpolog.com**
For more information, a product demo, or any other assistance, please don't hesitate to contact us at **support@xpolog.com**

Installation and Scalability

Installation

The software installs in less than five minutes on any JAVA supported OS, and can also be deployed as a web application on common J2EE application servers. After installation is completed, XpoLog can be accessed immediately by users using any common browser from anywhere in the organization.

Scalability

XpoLog scales easily to support high volumes of daily data.
XpoLog's recommended deployment is a cluster of several nodes that works correspondently as a cluster. It is possible and recommended to define certain node(s) to function as processor(s), and others as UI node(s). The processor(s) are responsible for the entire backend processing, while the UI node(s) are dedicated to providing high quality service to the users. As part of this mechanism, XpoLog automatically checks that all nodes are up and functioning as intended. In case of an outage, the mechanism alerts the system administrator immediately and other node(s) in the cluster get temporary responsibility to keep the mechanism active and updated. See more details at:
http://wiki.xpolog.com/display/XPOL/XpoLog+Cluster+Installation. Either way, sources are not affected by this mechanism. When the failed node(s) recover, they are immediately brought to their original structure.

Data Collection and Indexing

Data Collection

The data collection mechanism collects and stores log data with high priority and high credibility. XpoLog detects dynamic changing log data and manages references to the sources to ensure complete data collection. As part of the XpoLog cluster, it contains a self-monitoring mechanism to avoid any outages and informs system administrators on issues that should be addressed. The XpoLog cluster also works in failover mode, meaning that if a node in the cluster fails, another node automatically identifies it, alerts on it, and continues data processing to avoid data loss.

Data Storage and Retention

The XpoLog solution manages a set of collection policies. There can be a global collection policy for all sources, or different policies per different groups of sources. The collection policy determines the storage location (requires DAILY_VOLUME X RETENTION X 40%) and the collection frequency which is configurable from seconds upwards. XpoLog can interact with multiple storage devices to manage the data.

The mechanism collects data incrementally (only new records on each cycle) and also alerts system administrators in case of a failure (connectivity, permission, and more) or any other issue that system administrators should be aware of.

There is also an option to tune the data collection by applying filter rules to drop undesired records during the collection process. Based on the defined criteria, irrelevant data is ignored and is not collected to the XpoLog repository.

Supported Log Sources

XpoLog is completely nonintrusive and has a very low footprint (<2% from a single CPU core) on the sources. It has no impact on the sources in order to ensure that the applications' and systems' normal flow of work is not affected. It is also possible to determine the number of processes and connections that XpoLog established simultaneously to a source, in order to have full control and visibility on the XpoLog activity against the sources.

XpoLog supports Windows and UNIX/Linux platforms. Logs are collected from Windows platforms over the network using UNC paths (**Error! Hyperlink reference not valid.**...) and from UNIX/Linux platforms over SSH. The above methods are completely agentless and do not require changes in the existing infrastructure, only providing XpoLog read permissions to the required sources. In cases with specific requirements, XpoLog also supports an agent-based architecture to enable connectivity to the sources using an agent. In addition, there are XpoLog support database tables (using JDBC driver), Syslog, Windows Event Logs, and more.

There is no limitation on the source files' sizes; XpoLog manages the data in a very efficient way so that source files can be 1 KB or 1 TB, without any difference in search time per source.

Configuration Management

Log Types

XpoLog supports any textual format files; the collection and indexing mechanism supports dynamic rotated files (log.1, log.2, …, log.n). The system automatically recognizes multiple types of logs such as Windows Events, Access logs, IIS, Log4J, Apache, Syslog, J2EE application servers, Database tables, and many more.
In addition, it is possible to customize the parsing on specific log types and manage a set of user-defined templates in the system to be applied on a wider environment. While the data is being digested, XpoLog normalizes the data to a structured model to enable deep analysis and monitoring on all log fields.

Data Parsing and Normalization

XpoLog contains an advanced and highly flexible normalization engine – any log format (structured or unstructured) is supported and can be normalized by the system into a structured model. This flexibility is very important for users to gain more value from the logs – specific searches, deep statistics, and more accurate monitoring on any of the log fields and parameters. XpoLog supports multiline messages easily. Furthermore, it is possible to use the parsing engine to extract specific parameters/phrases from the multiline message / XML code of the message and extract them into virtual columns to be presented as regular log fields. (There is no limitation of record size in XpoLog. It can be configured to support any log record size.)

Data Management / Tagging

XpoLog offers several tagging options on logs, which can be used to manage large and heterogeneous environments. The basic configuration structure is managed in a hierarchical manner, i.e. in a user-defined Folders and Logs structure (similar to directories/files on File Systems). In addition, logs are automatically tagged to their source servers. Another available tagging is ‘Application’, which is a virtual entity that users can create to tag multiple folders/logs under a specific group. This approach has several advantages: easy navigation between sources, focused searches and analysis (log/folder/source server/application), and flexible implementation of security constraints (what a user can see and do).

Search Console

Search Options

XpoLog Search provides a wide variety of features for log analysis and troubleshooting processes. Rich search syntax (AND, OR, NOT, CONTAINS, = != < >, regular expressions, wild cards, and more) is available to isolate matching events from multiple sources on selected timeframes; complex syntax is available to generate deeper analysis such as aggregations, statistics, trends, correlations, min/max/average calculation, and more.

Search Augmentation – Integrated Knowledge Layers

One of the most interesting parts of the Search is 'Augmented Search', where XpoLog automatically detects errors in the search results and presents these as suggestions (tagged to low/medium/high severities) on the console side by side to the search results. This search augmentation boosts troubleshooting and exposes the hidden issues in the logs so that users can find out faster the source of their problems. For example, a user can run a search on some log source(s) and as part of the result, regardless of the search itself, XpoLog presents problems that were already detected and tagged by the system as suggestions. This makes the entire troubleshooting process deeper, faster, and more proactive.

Ad-hoc Visualization

Use of XpoLog Search regular queries, complex querying, and visualization options enables deep and comprehensive analysis of the data, starting from aggregations over time of log fields, and going through many functions that can be used to generate different calculations such as average, minimum, maximum, time bucketing, data normalization and many more. As part of the complex querying functionality of XpoLog, it is possible to create correlation rules based on multiple parameters that are common to a flow. Using the correlations/transactions functionality, it is possible to measure the number of transactions over time, quickly isolate a transaction's log event, measure times, and monitor the transactions level.
The data can be stored in XpoLog for as long as needed to enable both live analysis and monitoring, as well as historical analysis.

Data Visualization (Dashboards and Reports)

XpoLog contains an advanced dashboards and reports engine. It is possible to create multiple custom reports on the system and expose them to users using the interface or a scheduled export (PDF files that are scheduled to be sent).

The XpoLog reporting mechanism generates reports and stores the reports results for as long as defined. The data may be removed from the system; however, the reports results remain available for different purposes such as historical analysis or compliance regulations.

Logs Monitoring

XpoLog contains a comprehensive monitoring mechanism that can be scheduled to proactively look for predefined rules or automatically detected errors (Analytics engine) and trigger alerts accordingly. There are several alerting capabilities that can be easily implemented for different criteria – email, JMS, SNMP, and custom Scripting executions.

Analytics – Proactive Problems Discovery

Analytics is a proactive engine that scans all logs data as it is being added to XpoLog and identifies, without any pre-knowledge, critical problems in logs and servers.
Analytics presents a unified analysis report of all logs, based on several criteria:

1. **Problems Analysis –** General applications behavior presented over time. Analytics maps problems over time on a log or application level. Log problems are listed with drilldown links to the log events. A user can select a preferred view type (per log, per application, or per server) and quickly navigate between different sources over any desired timeframe.
2. **Risk Level Analysis** – Each problem indication that Analytics presents is tagged to a severity. Severity on automatically detected problems is assigned by XpoLog (and can be easily modified by users) or users, as predefined problems are added to the system.

The Analytics console is integrated with XpoLog Search in Augmented Search knowledge layers. Problems which the Analytics engine detects automatically, appear on top of the search results in the XpoLog Search console.

Security

Users Authentication and Authorization

XpoLog is integrated with LDAP, Active Directory, and additional SSO systems for authentication and authorization of users/groups who sign into the system.
Users are authenticated against the organizational authentication system and relevant permissions are applied on the user/group as required (what to do / what to see).

Data Masking and Compartmentalization

It is also possible to mask sensitive information from selected users/groups in order to ensure that only allowed data is available and presented. Masked values are not available for viewing and searching by the restricted users. The masking is defined on the parsing level, where fields or part of fields in the log can be masked.

API

XpoLog APIs' external applications can query XpoLog remotely and receive matching events for their own internal processing.
The reporting and report exporting are supported as well. It is possible to integrate external systems with XpoLog dashboards and reports, in order to consolidate everything in a centralized location.

System Health and Administration

XpoLog contains a system status console, which monitors all aspects of the system, starting with system resources (memory, storage, processing times, etc.) and going through detailed measurements of sources' storage capacity, average daily capacity, and more. This is done at the log source level, folders level, servers level, and applications level to provide a deep visualization on the sources' logging behavior to system administrators. XpoLog alerts system administrators on suspected outages and problems, so that action can be taken immediately, if needed.

Tel: +972 3 634 3884
Kfar Truman 1, 73150
P.O.B 174, Israel
Email: info@xplg.com
WWW.XPLG.COM