

HPE Structured Data Manager

Manage explosive application
information growth



HPE SDM at work at HPE

Hewlett Packard Enterprise, one of the world's largest IT companies, was experiencing similar challenges to those described in this data sheet. To address these issues, HPE's IT group established an application information archiving practice based upon HPE Structured Data Manager. Implemented in 2012, HPE IT has achieved the following benefits:

- Query time reduced by 89%
- Storage reduced by 48%
- Full database backup reduced by 37%

HPE Structured Data Manager (HPE SDM) reduces the total cost of ownership of application infrastructure, increases business productivity, raises information value, and mitigates the risks associated with tightening compliance requirements.

Organizations of all sizes are struggling to cope with application information growth. There is also increased complexity from tightening regulations, new database applications, and legal discovery requirements. As these challenges intensify, so does the amount of customer data to manage and the length of time you're required to manage it. This daunting scenario poses challenges to line of business owners, IT, and compliance and legal professionals alike.

Some of the problems organizations face managing enterprise applications include the following:

1. **Increasing data volumes** that negatively affect database and application performance, availability, and manageability. Data overgrowth also compromises critical business processes such as financial close, order processing, payroll, and management reporting.
2. **Retaining and managing large volumes of historical data** online in production databases on high-availability servers and storage puts unnecessary pressure on already strained budgets and overburdened IT staff.
3. **Failing to effectively address regulatory compliance demands** that can lead to stiff financial penalties and or sanctions.

HPE Structured Data Manager application archiving solution helps you address these challenges by effectively migrating eligible transactions off of your Oracle, Microsoft® SQL Server, IBM DB2, Sybase, or other JDBC-accessible databases while seamlessly preserving data access and integrity to support long-term retention and access requirements.

Comprehensive information management solutions from HPE

HPE Structured Data Manager provides unmatched flexibility to harness the power of a holistic approach to information management with the following built-in integrations:

- **HPE Records Manager:** Captures, manages, helps secure, and disposes of information according to corporate and compliance policies
- **HPE Consolidated Archive:** Central repository of structured and unstructured data with advanced search, classification, discovery, investigation, and surveillance functionality
- **HPE Vertica:** Provides high performance analytics so you can gain better business insights from your information
- **HPE Cloud Services and Amazon S3:** Leverage as storage to continue to lower cost and increase ROI for your cloud investment

Products/options description

HPE Records Manager	On-premise
HPE Consolidated Archive	On-premise and Cloud
HPE Cloud Storage and compute	Cloud
Amazon S3	Cloud
HPE Vertica	On-premise and Cloud
Hadoop	On-premise and Cloud

Database archive target location

Oracle database	On-premise
Microsoft SQL database	On-premise
IBM DB2 database	On-premise

Control mission-critical databases

HPE Structured Data Manager helps to control the growth of mission-critical databases by automating the migration or retirement of data while preserving its business value and meeting the desired access requirements. Data can be relocated to a separate, online database for fast, transparent access, or to standards-based XML, CSV, or JSON documents for long-term retention based on retention rules and policies that align with your business. You can perform out-of-the-box database archiving of structured data to any Hadoop-powered file system, regardless of the configuration: cloud, on-premise, or hybrid support. SDM also includes an integrated set of components that facilitate design, deployment, and on-going management of archiving processes throughout the lifecycle of applications and data.

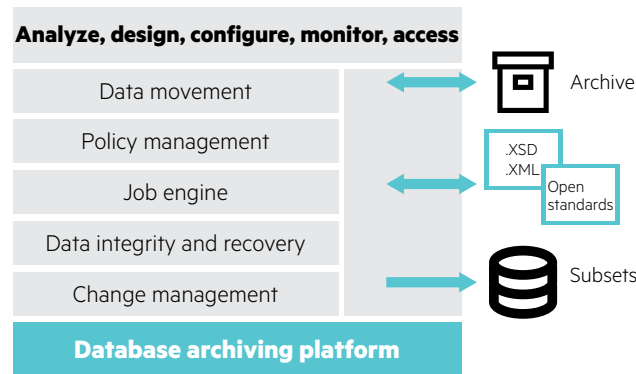


Figure 1: HPE SDM Platform Components

Adjust for application complexity

HPE SDM offers capabilities that address different levels of application complexity, data volumes, and archive access requirements. The components include:

- **Designer:** Provides a visual interface to model data and create business-aligned data migration rules with ease.
- **Data movement:** Makes sure data relocation is performed to meet volume requirements while retaining application integrity.
- **Archive access:** Provides a full range of access capabilities to meet requirements for business operations, regulatory compliance, and legal discovery. Simplified access via a Web-based console does not require third-party tools.
- **Job engine:** Automates application archiving tasks with built-in recovery and restart.
- **Management console:** Provides system configuration, job monitoring, job launching, and complete audit trail capabilities. Whether you are running applications on Oracle, Microsoft SQL Server, Sybase, DB2, or open standards JDBC environments, HPE Structured Data Manager offers the rich set of application archiving capabilities required to control and manage database growth.

Scale to Big Data

Today, information is at the core of the changing enterprise. The dramatic increase in volume, velocity, variety, and vulnerability of information is transforming the core of businesses and governments. HPE is helping information-complex companies make sense of Big Data with a broad and flexible portfolio based on open source standards and the availability of a comprehensive partner ecosystem. HPE Structured Data Manager offers full support for Big Data, allowing enterprises to be more agile and proactive with the ability to gain insights provided not only by HPE Structured Data Manager but the rest of the HPE information archiving portfolio, including HPE Vertica.

Rapid support for third-party, custom and in-house applications

The SDM Designer, a standard component included with the HPE Structured Data Manager license, provides a rich graphical interface for modeling application transactions and creating archive policies that extend archiving support to your third-party and custom applications. The Designer makes it possible for your business analysts and domain experts to focus on business-specific application archiving rules, instead of highly-specialized coding techniques, which simplifies the development and testing process. Designer is integrated with the HPE Structured Data Manager software run-time environment for easy deployment and maintenance.

Benefits of SDM Designer

- Enables fast development of archive modules for third-party and in-house applications, through visual modeling of data tables and creation of archive policies and rules
- Provides multi-level rule creation for intuitive and self-documenting design capabilities
- Allows preview reporting against production data sets to iteratively test archive rules and policies prior to server deployment
- Enables sharing and collaboration of archive projects using transportable design files
- Provides easy integration of custom scripts and external programs into automated archive operations using Business Flows

With the SDM Designer, you get a better return on investment by extending database archiving to virtually any application or database. You can also reduce your need for manual scripting and SQL coding, while relying on your in-house development team or preferred database VAR or integrator to build archiving modules quickly.

Faster data access and lower costs

Database-to-database archiving gives you the option to relocate data from production databases to a secondary, online archive database, allowing you to archive data in its original format. As a result, you gain the benefit of reducing the size of your primary databases, which has a two-fold effect:

- First, it dramatically accelerates end user access to production data while still providing transparent access to the archived data.
- Secondly, because it doesn't require additional database design skills or storage technologies, you can take advantage of in-house database administrator skills, which means no additional headcount.

HPE SDM helps you to capture greater storage cost savings by archiving inactive structured data to cost-effective private or public cloud storage while enabling users to directly search and/or query the archived data in the cloud.

HPE Structured Data Manager enables you to:

- Archive data from production databases directly to a secondary database with commit consistency and transaction integrity
- Give business users access to archived data using their native application interface, the WebConsole, or enterprise reporting tools in read-only mode
- Provide simultaneous access to production and archived data using combined reporting—a single query can run without change from standard application screens or reports
- Deliver scalable data movement capabilities to support the broadest spectrum of application transaction volumes
- Provide integrity of production and archive data times, through native database support and transaction consistency models
- Support native Oracle partitioning and partition swapping for fast data movement while archiving related non-partitioned data together as complete sets




Figure 2: Reduce data growth through regular archiving

Benefits of database-to-database archiving

- Boosts performance and availability of database applications
- Streamlines database backup, recovery, and cloning operations
- Reduces infrastructure costs by deferring server upgrades and reducing storage requirements
- Avoids any end-user retraining or application modification
- Deploys without costly, error-prone scripting and SQL coding

Benefits of database-to-file archiving

- Enables critical data to survive longer than the originating applications or databases
- Assists in meeting industry-specific compliance requirements for long-term data retention
- Deploys without costly, error-prone scripting and SQL codings
- Accelerates application retirement processes to shorten time-to-cost savings

Long-term compliance and standards-based data access

Database-to-file archiving gives users the option to archive data from production or archive databases to an industry-standard XML, JSON or CSV format. This type of automatic migration of data to open standards based documents helps you to:

- Achieve database and application independence
- Gain long-term viability of your archive
- Comply with industry-specific regulations that require data to survive longer than the originating applications
- Encapsulate relevant reference data with the transaction to make sure that the archive can stand alone; for example, customer-specific data elements such as name, address, and contact information are archived along with sales orders transactions

With HPE SDM, you can:

- Archive data from primary or archive databases to vendor-neutral, industry-standard XML/CSV/JSON files for long-term retention that is independent of database version, originating application, operating system, and hardware.
- Combine database-to-database and database-to-file archiving, through multistage policies that migrate production database data to secondary online archive databases. Then, as the data ages further, to XML, CSV, or JSON files.
- Run enterprise reporting tools and standard SQL queries unchanged, directly against the file-based archive.
- Store archived data in standard ASCII files on file systems, or on WORM or Content Addressable Storage systems. This function lets you choose your application data directly in the WebConsole for retiring. All you need to do is select the data to retire and then let the Retirement job run automatically.

Scalability and security for the largest, complex database applications

Easy to configure, HPE Structured Data Manager helps you meet business requirements and unique characteristics of large-scale database environments, all the while handling massive data volumes with integrity. Production and archived data are kept in a consistent state because data is either archived as complete business transactions or not archived at all. With retention policies being enforced for the first time that means the initial archiving runs can involve a huge amount of data. Standard archive jobs can be run for ongoing operations. Encrypted passwords and detailed audit trails provide security for archiving operations support.

Features

- Supports chained relationships (multiple recursive relationships) to address business cases with complex transactional dependencies in addition to table level policies and rules.
- Automates recovery from failures during archive runs by archiving complete transactions and built-in job recovery capabilities, no matter how many tables are involved.
- Uses detailed audit trails to capture archive cycle information, messages, and statistics.
- Stores any password required for operations in encrypted format to meet compliance with security standards.
- Restores entire archive runs or individual transactions back into production. Safety checks are performed to make sure the data is inserted correctly without jeopardizing production database integrity.

Benefits

- Keeps both production and archive data highly available
- Scales to meet needs of largest and complex enterprise database applications
- Provides long-term data retention solution for production databases
- Facilitates business reporting and eDiscovery of structured data

Robust run-time environment and native database integration

HPE Structured Data Manager is certified for multiple open system databases and is designed to leverage vendor-specific database capabilities, to help you simplify operations, increase application archiving performance, and reduce maintenance. This native API and deep integration approach makes sure that once the database archiving modules are designed and implemented, that you can run them on a regular basis—monthly or quarterly—to fully maintain your databases and monitor your archiving jobs just like any other data center task. Each archiving operation is launched as a single job; in the event of an error, the job engine handles recovery tasks and re-launches the job, as required.

HPE SDM features

- Supports major versions of Linux®, and Windows®
- Supports major versions of Oracle, SQL Server, Sybase, and DB2 along with JDBC, through API-level integration for improved archiving performance
- Provides configurable multi-threaded parallel job execution, to align archive operations with available server resources
- Deploys in flexible configurations, e.g., you can run archiving software on the database server or separate server, as desired
- Supports multiple languages and countries through multi-byte character and object support
- Provides job launching and monitoring capabilities so you can view jobs and drill down into exceptions, as required
- Leverages existing investments in data center tools, e.g., job schedulers, through supported command line instructions on supported operating systems
- Allows online archiving operations so users can remain online in the production applications without jeopardizing integrity
- Provides Eligibility Analytics to investigate data eligibility issues and dependencies, with transaction-level drill down to diagnose and resolve complex chaining situations

HPE SDM benefits

- Provides fast installation, set up, and rapid return on investment
- Enables database archiving deployments across enterprise applications, with lesser impact on database administration and operations staff
- Leverages existing investments in IT operations management products
- Supports global and international deployments

Ready-to-run modules for Oracle Applications

HPE Structured Data Manager out-of-the-box modules for Oracle E-Business Suite and PeopleSoft Enterprise offer comprehensive support for Financials, including Sub-Ledger Accounting, Project Management, Supply Chain, Manufacturing, CRM, Human Resources, and Payroll applications.

Comprehensive services, support, and training

Hewlett Packard Enterprise offers a full set of implementation services and multiple support options to help you realize the full potential of your HPE Information Management solutions, enable information availability, mitigate business risk, and achieve a superior return on your investment. For more information about HPE service and support, visit [hpe.com/services](https://www.hpe.com/services).

Available standard products and options

Choose from the following products and options to customize HPE Structured Data Manager to fit your information infrastructure.

PRODUCTS/OPTIONS	DESCRIPTION
Base System	<ul style="list-style-type: none"> • Graphical-based Installer and Deployment Assistant: Provides ease of installation and deployment • Platform support: Offers certified support for Windows Server®, UNIX®, and Linux as archive server and/or database server • Designer: Graphical design environment features rich modeling capabilities, multi-level eligibility rules and policy building, preview reporting, project sharing, and offline cache mode • Run-time environment: Delivers powerful job execution capabilities, including native database and platform support, job engine automation, built-in recovery/restart, and extensibility framework • Standard Data Movement: Facilitates database-to-database and database-to-file data movement • Business flows: Allows multiple jobs or activities to be consolidated into a single logical run-time operation, including the ability to call custom jobs/scripts/external programs • Web-based Management console: Enables configuration, monitoring, administration, and job launch capabilities • Support for transactional and batch reload capabilities • Advanced Data Movement Option: Provides support for high-volume and complex data environments by leveraging specialized data movement techniques, including support for table partitioning. Also, provides automated support for source table reorganization during archiving operations • Rapid Application Retirement: User friendly GUI to rapidly retire multiple applications with minimal manual processing
Database Connectors	<ul style="list-style-type: none"> • Native support for Oracle database environments, including, 9i, 10g, 11g, and 12c. • Native support for Microsoft SQL Server 2005, 2008, and 2012 environments • Native support for Sybase environments, including 12.5 and 15 • Native support for DB2 environments, including 9.7 • Mainframe Generic support for open standards JDBC environments
Archive Query Server Cache Option	<ul style="list-style-type: none"> • Direct SQL access to XML, JSON, or CSV archive files from any ODBC/JDBC/.NET client
Archive Access Transparency Option	<ul style="list-style-type: none"> • Support for native application access in query-only mode. Also supports combined reporting for access to current and archive data in a single report or query • Archive access from WebConsole provides remote access to archive stores, and reduce the need for third-party tools or source application
Developers License	<ul style="list-style-type: none"> • Unrestricted capabilities to design, test, and deploy archive modules in development environments



Sign up for updates

★ Rate this document



Learn more at
hpe.com/software/sdm

© Copyright 2015–2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Oracle is a registered trademark of Oracle and/or its affiliates. UNIX is a registered trademark of The Open Group. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

4AA5-8166ENW, March 2016, Rev. 1