



Easily Modernize and Expose Valuable Enscribe Data – Eliminate Data Silos

A Gravic, Inc. Article



The Manufacturer’s Trapped Data

A large European steel tube manufacturer ran its online shop floor operations on an HPE NonStop Server. To exploit the currency and value of this online, trapped, and siloed data, the manufacturer periodically generated reports on Linux Servers using a customized application and connectivity tool that remotely queries the online NonStop Enscribe database and returns the results (Figure 1).

With this original architecture, every time a Linux query/report was run, processing on the NonStop Server was required, and as query activity increases, this workload started to significantly impact online shop floor processing. This impact was compounded due to the high volume of data transformation and cleansing required for converting the Enscribe data into a usable format for the reports. Periodically, the company needed to suspend the execution of reports due to these production impacts. In addition, the remote connectivity architecture was not very robust or scalable, and was also susceptible to network failures, timeouts, and slowdowns when operating at full capacity. Since the data adapter used was nonstandard, the manufacturer could not access the data using standard ODBC and OLAP¹ tools to take advantage of new analytical techniques (such as DSS²). Furthermore, the company had a new requirement to share the OLAP analysis with the online NonStop applications in order to optimize shop floor control, which was completely impossible with the original solution. Therefore, a new architecture was required to address these issues and meet the new requirements.

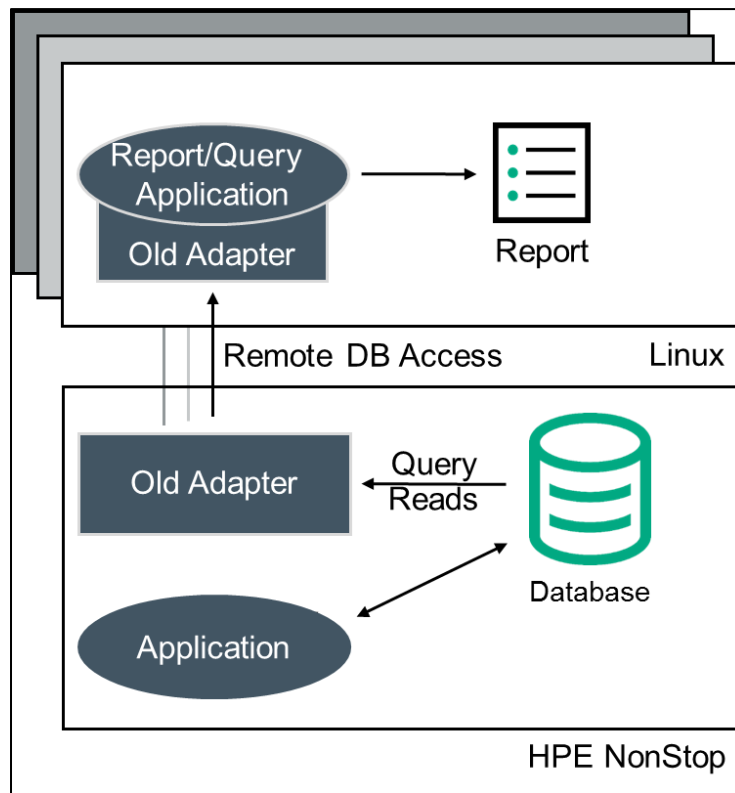


Figure 1 – Original Remote Query Architecture

¹Online Analytical Processing (OLAP) is the technology behind many Business Intelligence (BI) applications. OLAP is a powerful technology for data discovery, including capabilities for limitless report viewing, complex analytical calculations, and predictive “what if” scenario (budget, forecast) planning.

²A Decision Support System (DSS) is a computerized information system used to support decision-making in an organization or a business. A DSS lets users sift through and analyze massive reams of data and compile information that can be used to solve problems and make better decisions.

Rearchitecting Operations

The manufacturer completely rearchitected its Linux-based querying/reporting application. Rather than remotely querying the NonStop Enscribe data each time a report is run, the data is replicated by HPE Shadowbase software in real-time from the Enscribe database to a copy of that database hosted on a Linux system. This architectural change only requires the data to be sent across the network once (when it is changed), instead of every time a query/report is run.

The raw Enscribe data is non-normalized, and full of arrays, redefines, and data types that do not have a matching relational data construct (e.g., SQL data type). As part of the replication process, the non-relational Enscribe data is transformed and cleansed by HPE Shadowbase software into a relational format, and written to an Oracle RAC database. This architectural change only requires the data to be cleansed and/or normalized once when the data is changed, rather than every time a query/report is run.

Since the data is now local to the Linux Servers and presented in standard relational format, it is possible to use standardized SQL data query and analysis tools (Figure 2). The benefit of achieving this level of integration cannot be overstated; the customer now has the full power provided by the analytical tools market to manipulate its *raw information* to achieve *data intelligence*.

In addition, HPE Shadowbase software allows the manufacturer to reverse replicate the OLAP results and share those results back with the NonStop applications to better optimize the shop floor manufacturing process. This integration dramatically improves shop floor control operations and avoids having to architect yet another interface.

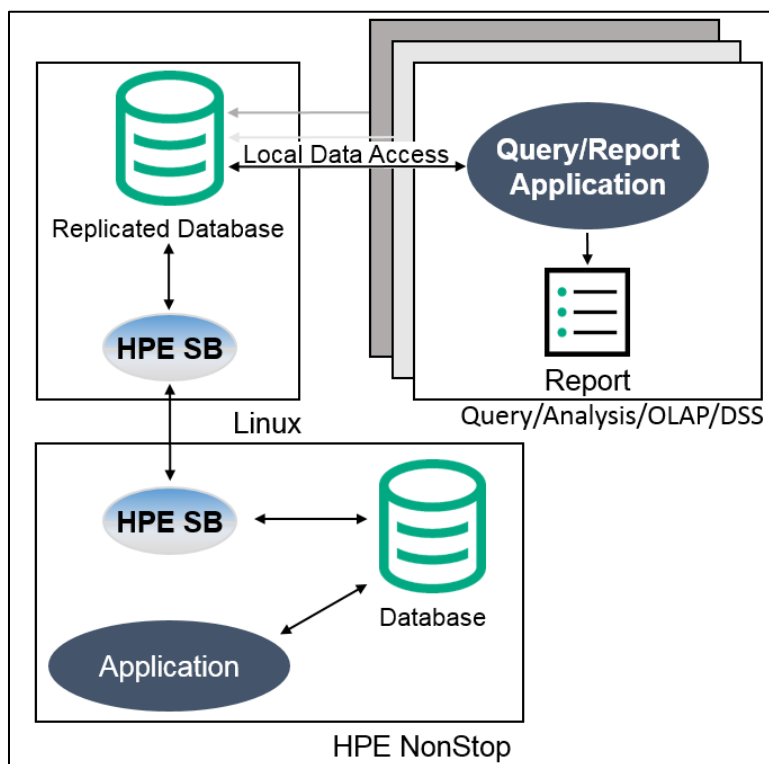


Figure 2 – The Manufacturer's New Architecture Featuring Local Database Access

Key Project Benefits

This new architecture completely resolves the issues with the old mechanism by decoupling the data access and transformation/cleansing process from the querying/reporting process, and by presenting the data locally on the Linux system in standard relational format:

- **Querying/reporting and data transformation/cleansing impact on the production NonStop system is eliminated.** Online data replication has much lower overhead than executing queries and

performing data transformation/cleansing directly on the NonStop Enscribe database every time a query is executed. This architecture dramatically reduces the workload on the NonStop Server, avoids throughput and response time issues with the shop floor applications, and eliminates the periodic need to suspend running reports.

- **Querying/reporting performance is dramatically improved.** It is much faster to run queries and reports against an already transformed/cleansed local relational database using standard ODBC access, than it is to remotely access the NonStop Enscribe database and spontaneously perform the necessary data transformation/cleansing every time a query or report is executed.
- **Querying/reporting application availability is improved.** A local Linux copy of the database is used so that an outage of a network or a NonStop Server will not impact the querying/reporting applications.
- **The shop floor application's scalability is improved.** A local database access is inherently more scalable than executing remote queries across a network, especially when removing the need for data transformation/cleansing on every query. Competition for resources on the NonStop Server is also greatly reduced by offloading workload to the Linux systems. In addition, the original architecture typically only allowed a single report to run at any one time (to minimize the impact on the NonStop production system). Now, with the data hosted on a Linux Oracle RAC database and the NonStop system out of the query loop, multiple reports and advanced analytics can be executed in parallel.
- **Analytics and data value are improved.** Transforming and cleansing the Enscribe data into relational format and storing it in a local relational database enables the use of standard ODBC database access. It also enables additional off-the-shelf querying/reporting tools to extract even more value from the data (including OLAP and DSS). The ability to share the OLAP analysis with the NonStop production applications using Shadowbase bi-directional data integration enables new uses for the data (in this case, to better optimize the shop floor manufacturing process).

Summary

This case study only scratches the surface of data replication's potential. It can:

- **Integrate disparate applications**, such as an online production shop floor application with a reporting application.
- **Eliminate data silos** by exposing valuable data and enables the creation of new business processes.
- **Distribute data wherever it is needed** with low overhead while it transforms and cleanses that data.
- **Distribute data where it can easily be consumed by other applications** (typically without requiring any application changes).
- And much, much more.³

Data replication is not only for providing backups for business continuity, but also for moving data in real-time as required in order to leverage its value wherever needed. Think about where you have valuable, yet isolated, data and consider how you could use data replication to unlock this value for competitive advantage and to build new solutions for your business.

³For more information, please see the Gravic white papers, [HPE Shadowbase Streams for Data Integration](#) and [HPE Shadowbase Streams for Application Integration](#).

Two Options for Eliminating Trapped, Siloed Enscribe Data: Off-platform vs On-platform Data Integration

This case study describes the solution chosen by the customer known as “off-platform” data integration, where Shadowbase software is used to transform and replicate the trapped and siloed source data (HPE NonStop Enscribe) into the desired target data structure (SQL), and apply it into a remote (off-platform) target platform and database (Linux/Oracle).

An alternate approach, known as “on-platform” data integration, can also be considered (see Figure 3).

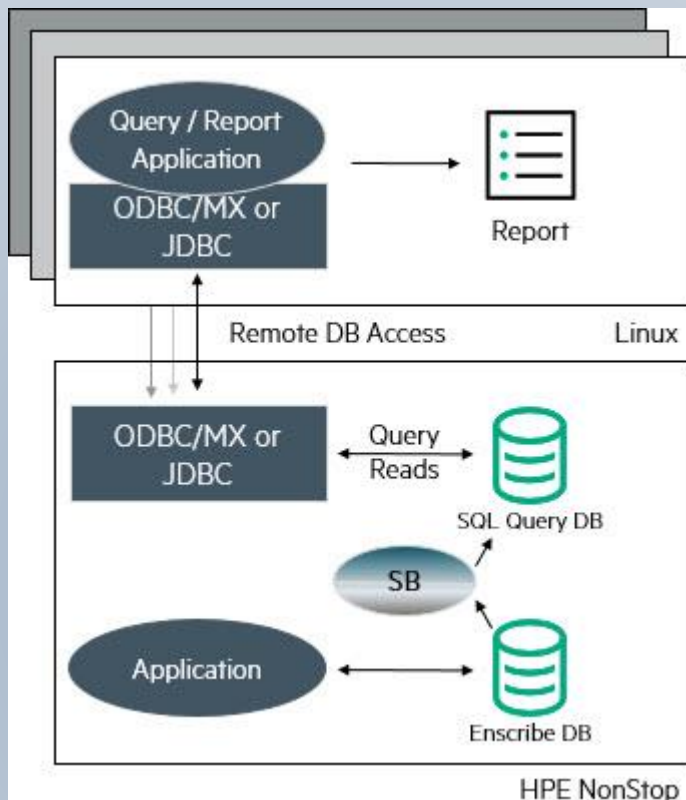


Figure 3 – On-platform Data Integration

In this approach, Shadowbase software is used to transform and replicate the source data into the desired target structure and apply it into a “local” (on-platform) target database (either NonStop SQL/MP or SQL/MX). Remote applications can then use standard ODBC or JDBC query tools to remotely access the data in the SQL/MP or SQL/MX database (e.g., via ODBC/MX).

The on-platform architecture is very similar to the original architecture employed by the customer – only the remote query application client library would have to change – which can be very advantageous if the customer cannot easily accommodate or manage off-platform systems (e.g., Linux, Windows, Unix) and databases (e.g., Oracle, SQL Server, Db2®). Additionally, keeping the database on the NonStop Server provides all of its inherent fault-tolerant availability benefits, including its unequalled NonStop SQL database query capabilities (e.g., provided by SQL/MX). Of course, when adding additional processing to the existing system, a sizing and performance analysis should also be performed to ensure that the system capacity is adequate (otherwise, additional CPU or disk may be needed).

In this case study, the customer chose to use Shadowbase off-platform data integration capabilities in order to take advantage of additional tools, and to integrate with additional utilities available for the Oracle database (e.g., OLAP and DSS).

International Partner Information

Global

Hewlett Packard Enterprise

6280 America Center Drive
San Jose, CA 95002
USA

Tel: +1.800.607.3567

www.hpe.com

Japan

High Availability Systems Co. Ltd

MS Shibaura Bldg.
4-13-23 Shibaura
Minato-ku, Tokyo 108-0023
Japan

Tel: +81 3 5730 8870

Fax: +81 3 5730 8629

www.ha-sys.co.jp

HPE Shadowbase Contact Information

17 General Warren Blvd.
Malvern, PA 19355-1245
USA

Tel: +1.610.647.6250

Fax: +1.610.647.7958

www.shadowbasesoftware.com

Email Sales: shadowbase@gravic.com

Email Support: sbsupport@gravic.com



Hewlett Packard Enterprise Business Partner Information

Hewlett Packard Enterprise directly sells and supports Shadowbase Solutions under the name **HPE Shadowbase**. For more information, please contact your local HPE account team or [visit our website](#).

Copyright and Trademark Information

This document is Copyright © 2019, 2020 by Gravic, Inc. Gravic, Shadowbase and Total Replication Solutions are registered trademarks of Gravic, Inc. All other brand and product names are the trademarks or registered trademarks of their respective owners. Specifications subject to change without notice.