

**Hewlett Packard
Enterprise**

Unleash the Value of Your Trapped Data with HPE Shadowbase

Paul J. Holenstein
Executive Vice President
Shadowbase Products Group
Gravic, Inc.

May 2019



Introduction



Paul J. Hostenstein
Executive Vice President
Shadowbase Products Group
Gravic, Inc.



Disclaimer

This presentation contains forward-looking statements regarding future operations, product development, product capabilities and availability dates. This information is subject to substantial uncertainties and is subject to change at any time without prior notification. Statements contained in this presentation concerning these matters only reflect Gravic, Inc.'s predictions and/or expectations as of the date of this presentation and actual results and future plans of Gravic, Inc. may differ significantly as a result of, among other things, changes in product strategy resulting from technological, internal corporate, market and other changes. This is not a commitment to deliver any material, code or functionality and should not be relied upon in making purchasing decisions.

Specifications are subject to change without notice and delivery dates/timeframes are not guaranteed...purchasing decisions should not be made based on this material without verifying the desired features are available on the platforms and environments desired.

All trademarks mentioned in this presentation are the property of their respective owners.

Agenda

The value of data

A customer's trapped data

Rearchitecting operations

Further comments

HPE Shadowbase product suite overview

Summary & for more information





The value of data

Data is...

The new currency

- The more current data is, the more valuable and actionable it becomes
- Data managed by online transaction processing systems is among the *most valuable* and the *most sought-after*
- Conversely, the older data is, the less valuable it is



Data is...

Trapped

- Trapped data is data stuck in a single application, system or database
- Competitive opportunities are missed because its full value is not being exploited
- Rapidly creating additional value-add solutions is not possible because it is not easily accessible
- We call this a “data silo”



Data is...

An opportunity and a challenge

Opportunity

- Immediate access to real-time data creates new business opportunities
- Leveraging existing information reduces both time and cost to market
- Examples include:
 - Decision support systems
 - Fraud detection
 - Shop floor control
 - Automated alerts and notifications to drive (trigger) new services



Challenge

- However, immediate access to real-time data impacts the online system's throughput and response time
- Worse, as the demand for this data increases, the load on the online system continues to increase
- Therefore, the challenge is to make the real-time data immediately available in a usable format to the interested parties with minimal impact to the online environment



This presentation discusses such a challenge, and the clever way it was resolved using *online data replication*




A customer's trapped data

Trapped data

Offload querying/reporting from the host

Situation

- A large European steel tubing manufacturer runs its shop floor control applications on an HPE NonStop system
- The manufacturer periodically runs reports on Linux servers that remotely access the NonStop Enscribe database to gather state information
- The reporting application uses a nonstandard interface into the NonStop database
- These reports need to heavily transform the non-relational Enscribe database information into a usable format, and require transmission of the data across the network for each report run



Trapped data

Offload querying/reporting from the host

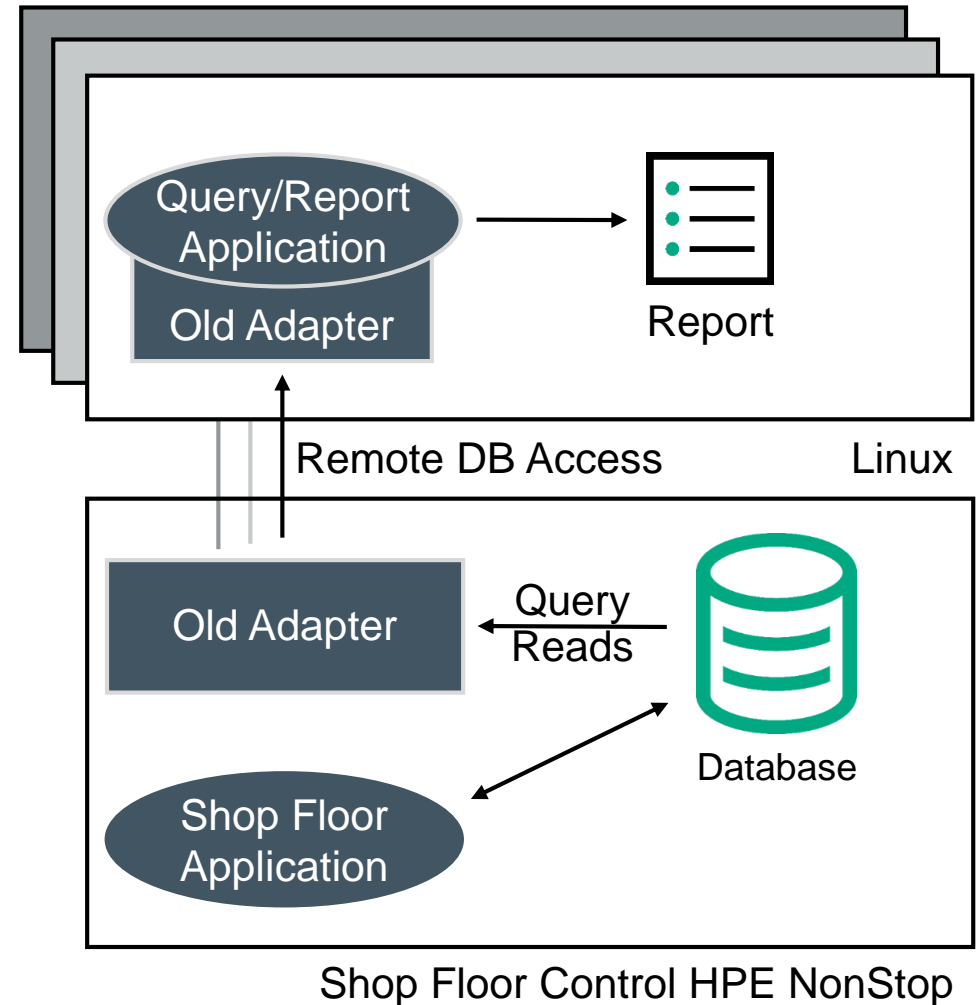
Old architecture

- The query/report application polls the old adapter to extract the data for the report, and returns it to produce the report

Challenges

- Each time the report is run, it reaches across the network, extracts the production data, transforms it, and sends it to the Linux reporting application
- The raw Enscribe data is nonnormalized, full of arrays and redefines, and includes data types that must be cleansed and transformed into a usable reporting format
 - Each report run could dramatically impact the production environment
- As the query/reporting workload grows, the increased workload threatens the shop floor application
- This leads to an architecture that did not scale and is susceptible to network failures and timeouts

These challenges required a new architecture to address these issues and meet new requirements





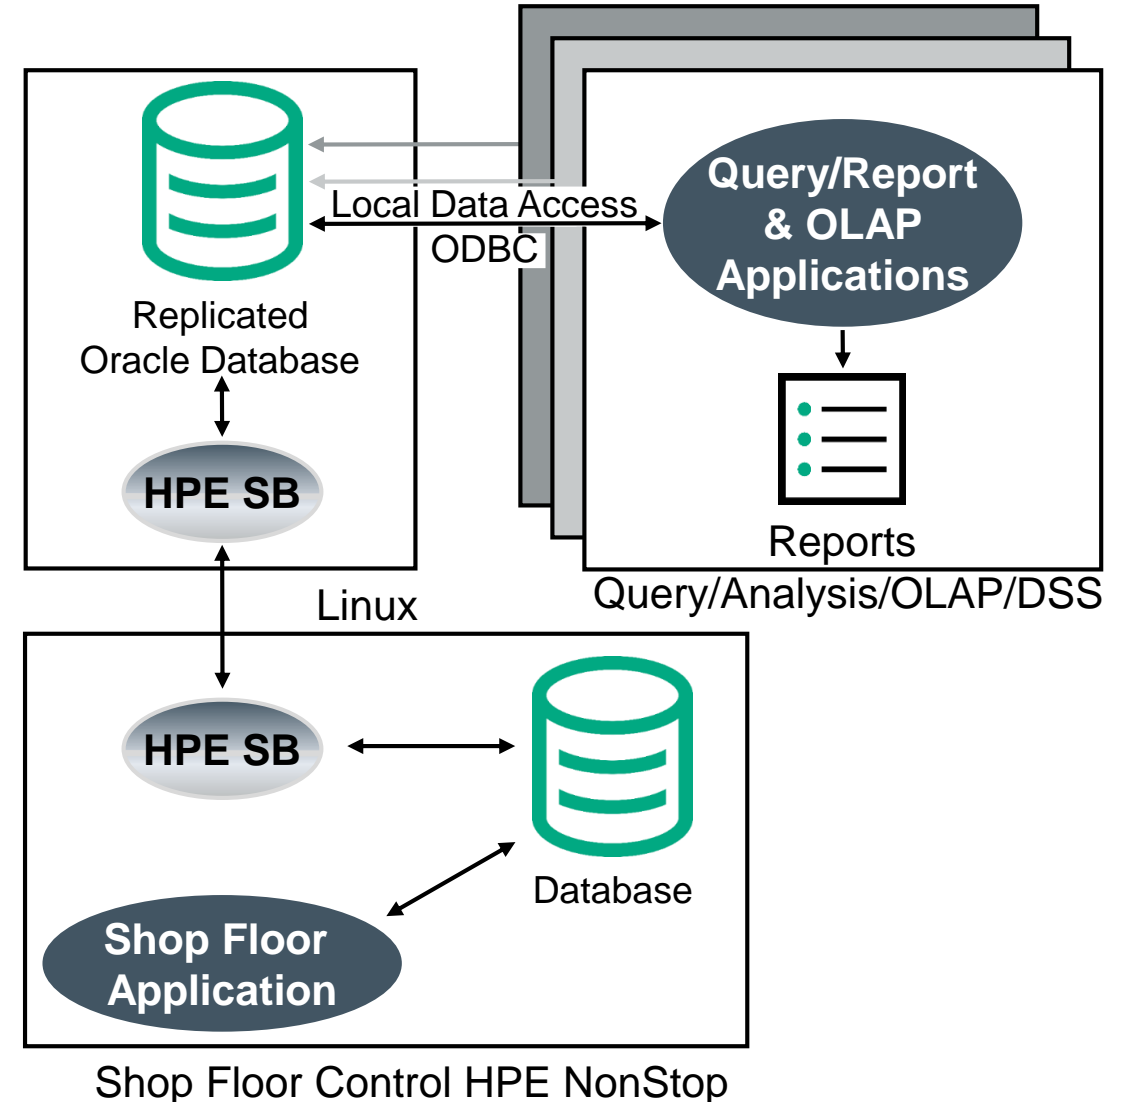
Rearchitecting operations

Rearchitecting operations

Offload querying/reporting from the host

New architecture

- The manufacturer selected HPE Shadowbase to copy production data to a commodity-based reporting environment
- The new architecture efficiently cleanses, normalizes, and replicates the data across the network only once (when it changes), rather than each time a query/report runs
- The customer selected standard ODBC connectivity for access to the query/reporting database for generating reports
- The customer also implemented new OLAP tools for analytical processing, and configured Shadowbase to bi-directionally replicate the OLAP results back to the NonStop for improved shop floor control

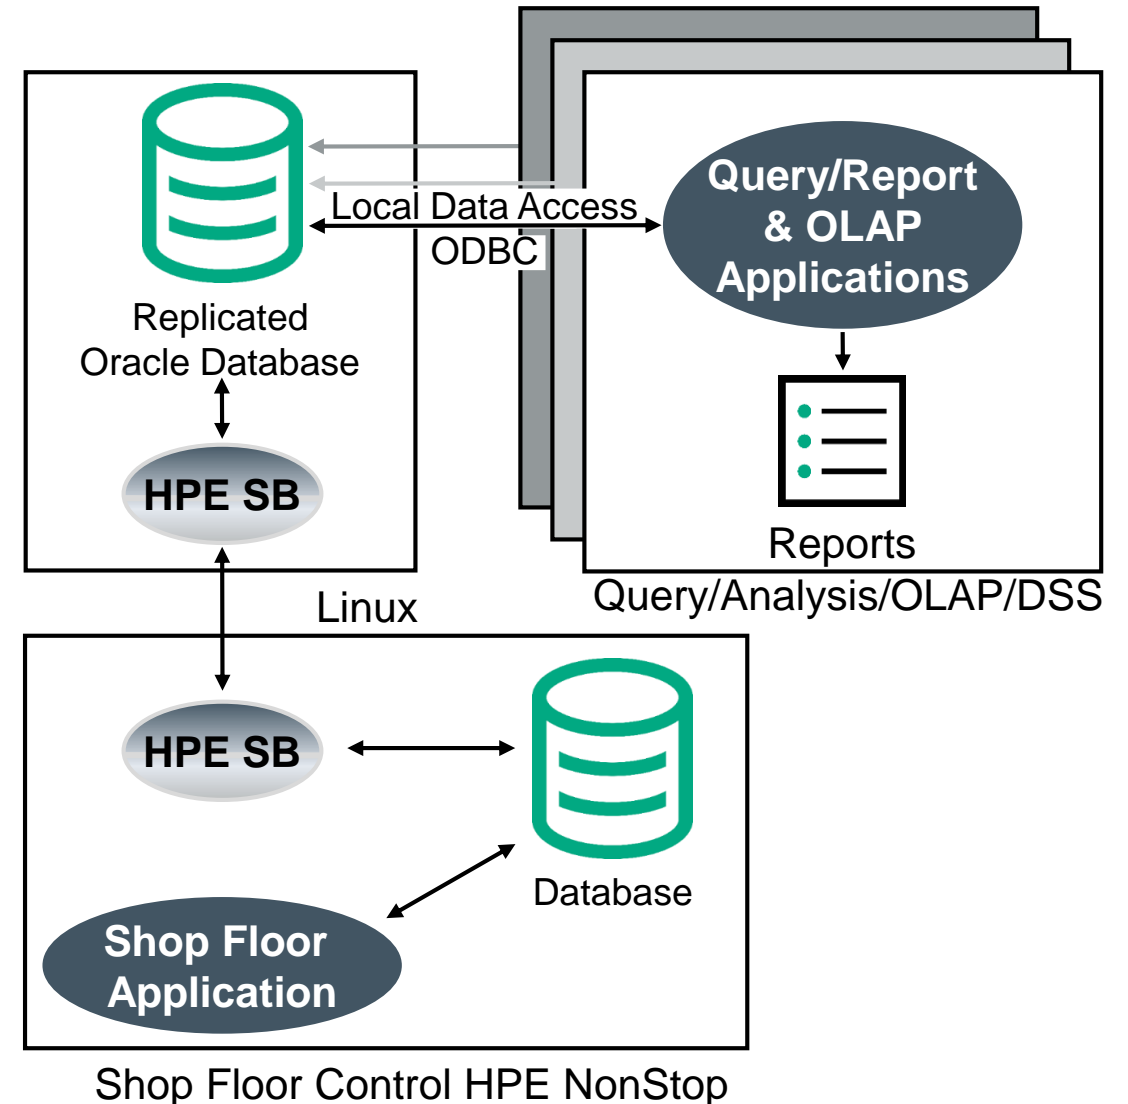


Rearchitecting operations

Offload querying/reporting from the host

Key benefits

- Enables a scalable on-demand reporting infrastructure without impacting the NonStop host environment
- Transforms the nonnormalized host data into an SQL relational format for improved reporting/query processing
- Standardizes access to the local replicated database via ODBC, opening up the data to DSS tools
- Eliminates network failures or timeouts from impacting the query/reporting and analytical subsystem
- Facilitates sharing the reporting and OLAP results back to the NonStop application, which optimizes shop floor control, and was completely impossible with the original solution





Further comments

Using HPE Shadowbase for Data Integration

Further comments

Further comments

- Customer originally chose Linux/Oracle as the reporting subsystem environment based on corporate policy
- Offloading the reporting from the operational system to another platform is a standard way to spread the load and isolate dynamic (and intensive) read/processing activity to another environment
- However, Shadowbase could have replicated the data from Enscribe into a local SQL relational target database (e.g., SQL/MX)
 - As an alternative, the customer could have chosen to keep all of the data resident on the host platform, although would have to create the reporting database on different CPUs/disks than the operational system was using
 - The Nonstop is a mixed-workload environment (MWE) that can handle both the operational application as well as a reporting subsystem (hybrid IT)
- Alternatively, use HPE Shadowbase to replicate the operational system data to a disaster recovery backup system, and run reporting/OLAP there to maximize capacity utilization
 - Shadowbase will keep the target databases “fresh,” important for business continuity RPO as well as data integration processing using current state data



Read the case study:
[Large Manufacturer Uses Data Replication to Offload Online Host Processing and Enable New Business Applications](#)



What is HPE Shadowbase?

What is HPE Shadowbase?

Video: <https://www.shadowbasesoftware.com/>



What is HPE Shadowbase?

Business continuity

Capabilities

- Extreme availability
- Zero downtime migrations



Eliminate

- Unplanned downtime
 - Active/Passive Disaster Recovery
 - Sizzling-Hot-Takeover (SZT)
 - Active/Active (Hot-Hot) Continuous Availability
- Planned downtime
 - Zero Downtime Migrations (ZDM) for upgrades, new versions, and platform refresh

What is HPE Shadowbase?

Data integration

Capabilities

- Integrate data from one database into another for competitive advantage



Master data management

- Homogeneous and heterogeneous environments
- Data transformation, scrubbing, filtering, and cleansing
- Extension of replication capabilities with embedded application logic

Popular uses

- Data streaming and change data capture
- Data warehouse feeds
- Offline and online data loading/synchronization

What is HPE Shadowbase?

Application integration

Capabilities

- Integrate one application with another for new solutions



Build

- Event-driven architectures
 - Process events as they occur; no more polling for needed data
- Real-time architectures
 - Process events when they occur; no more working with “stale” data
- Integrated applications without application code changes
 - Integrate at the data-layer, avoiding costly adapters, middleware, and code changes

Popular uses

- Decision support systems
- Real-time business intelligence
- Real-time fraud detection or sales analysis

What is HPE Shadowbase?

Utilities

Capabilities

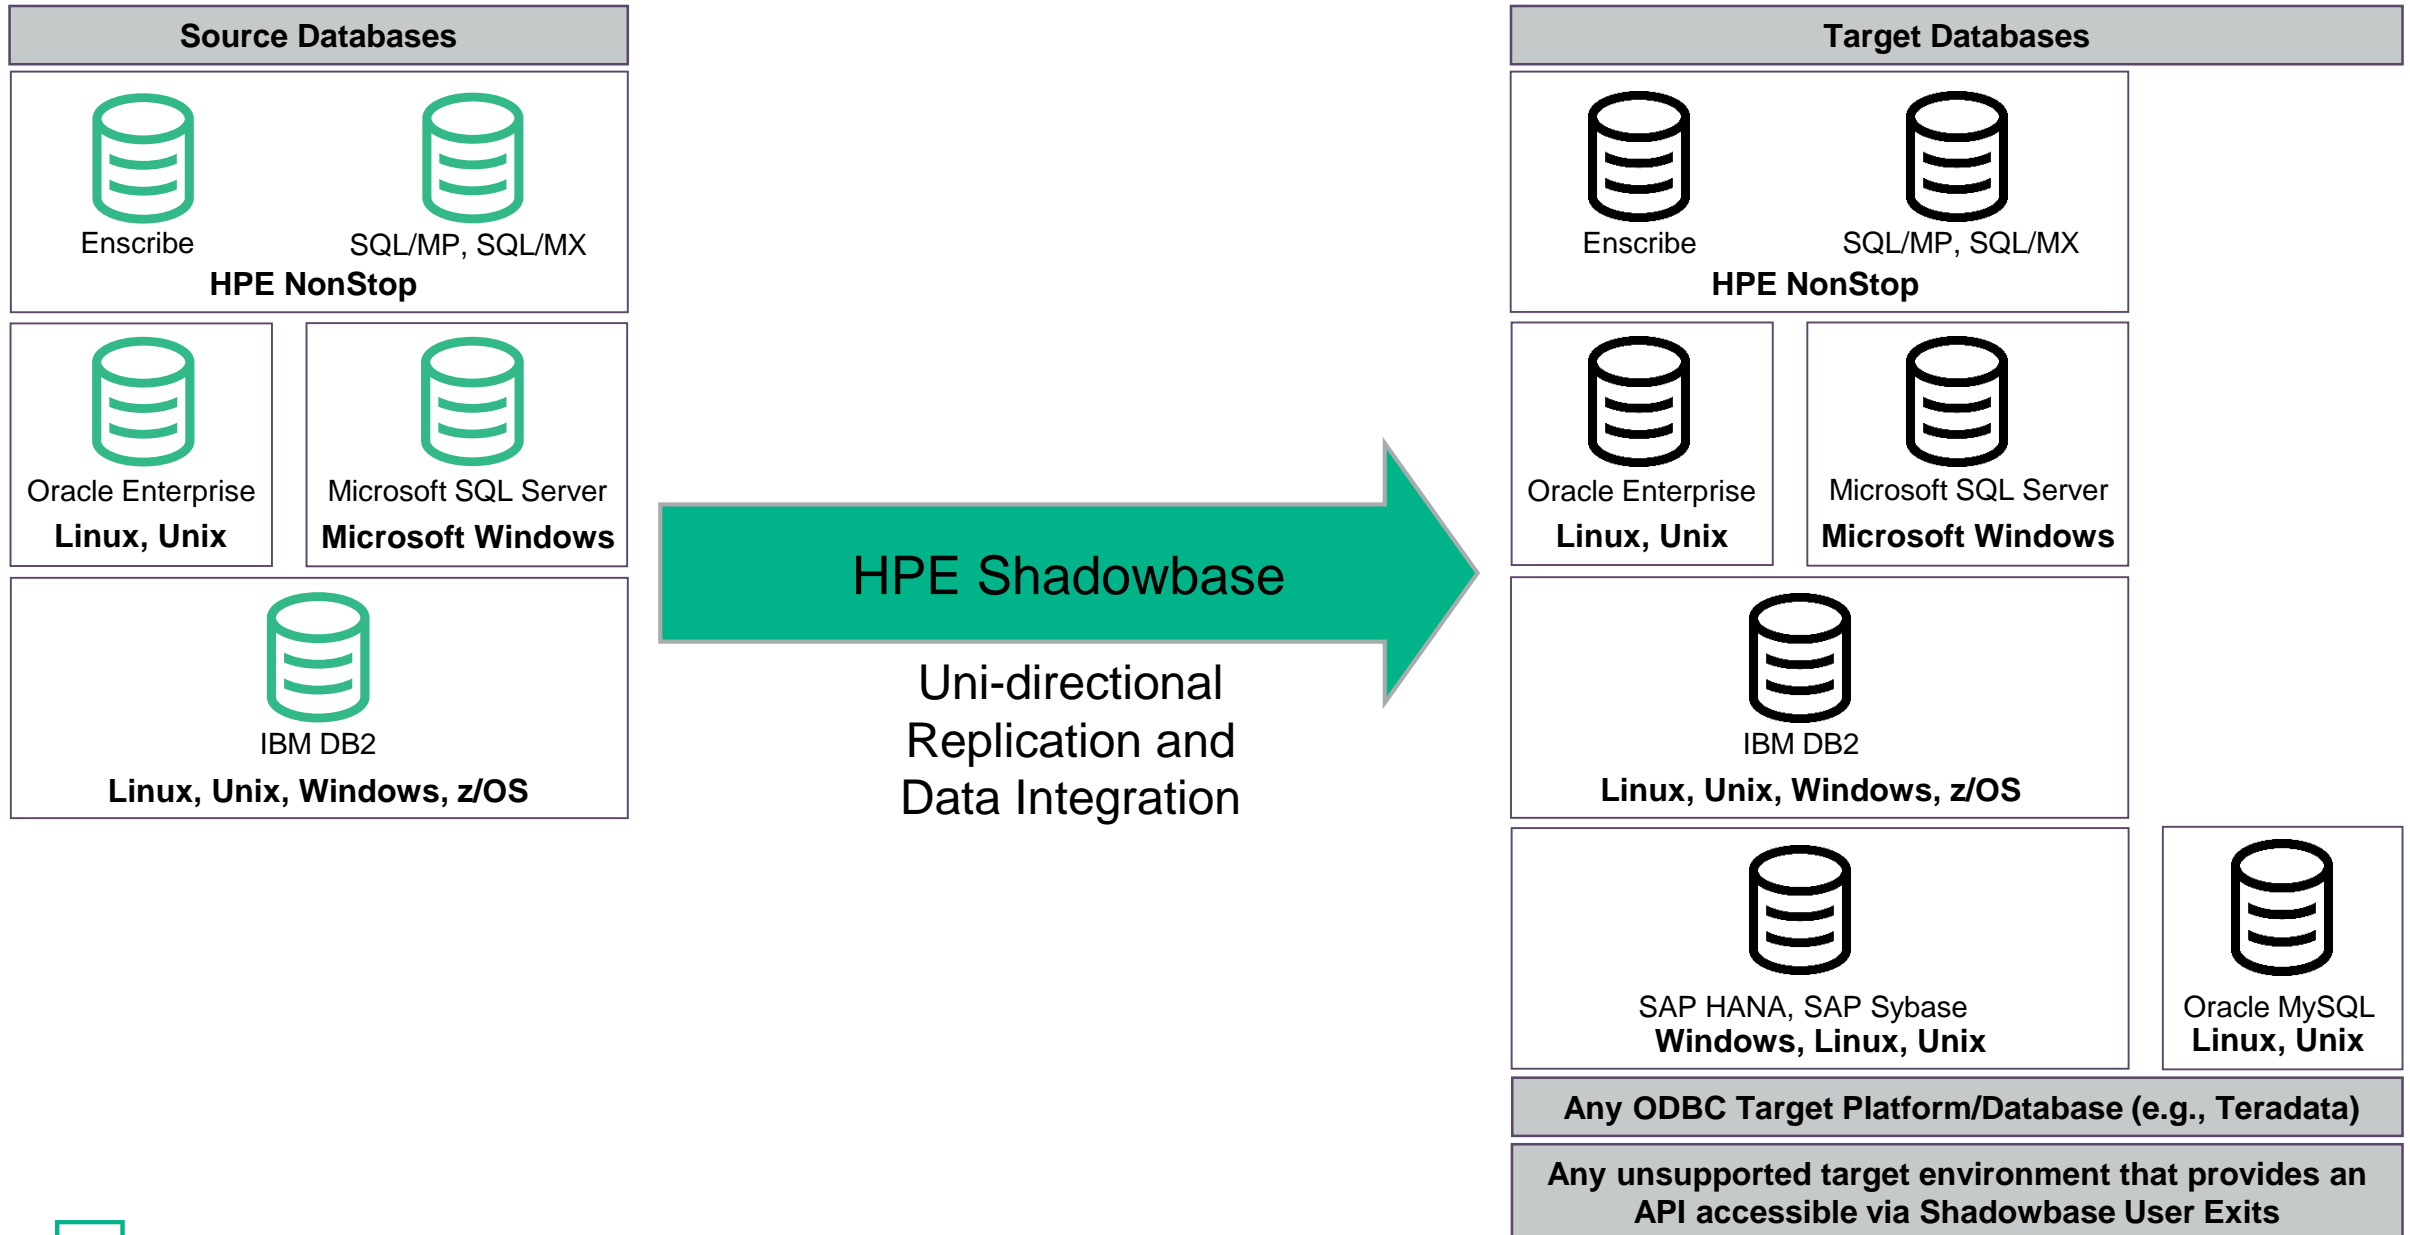
- Audit compliance, data validation, data transformation, ETL, online data recovery



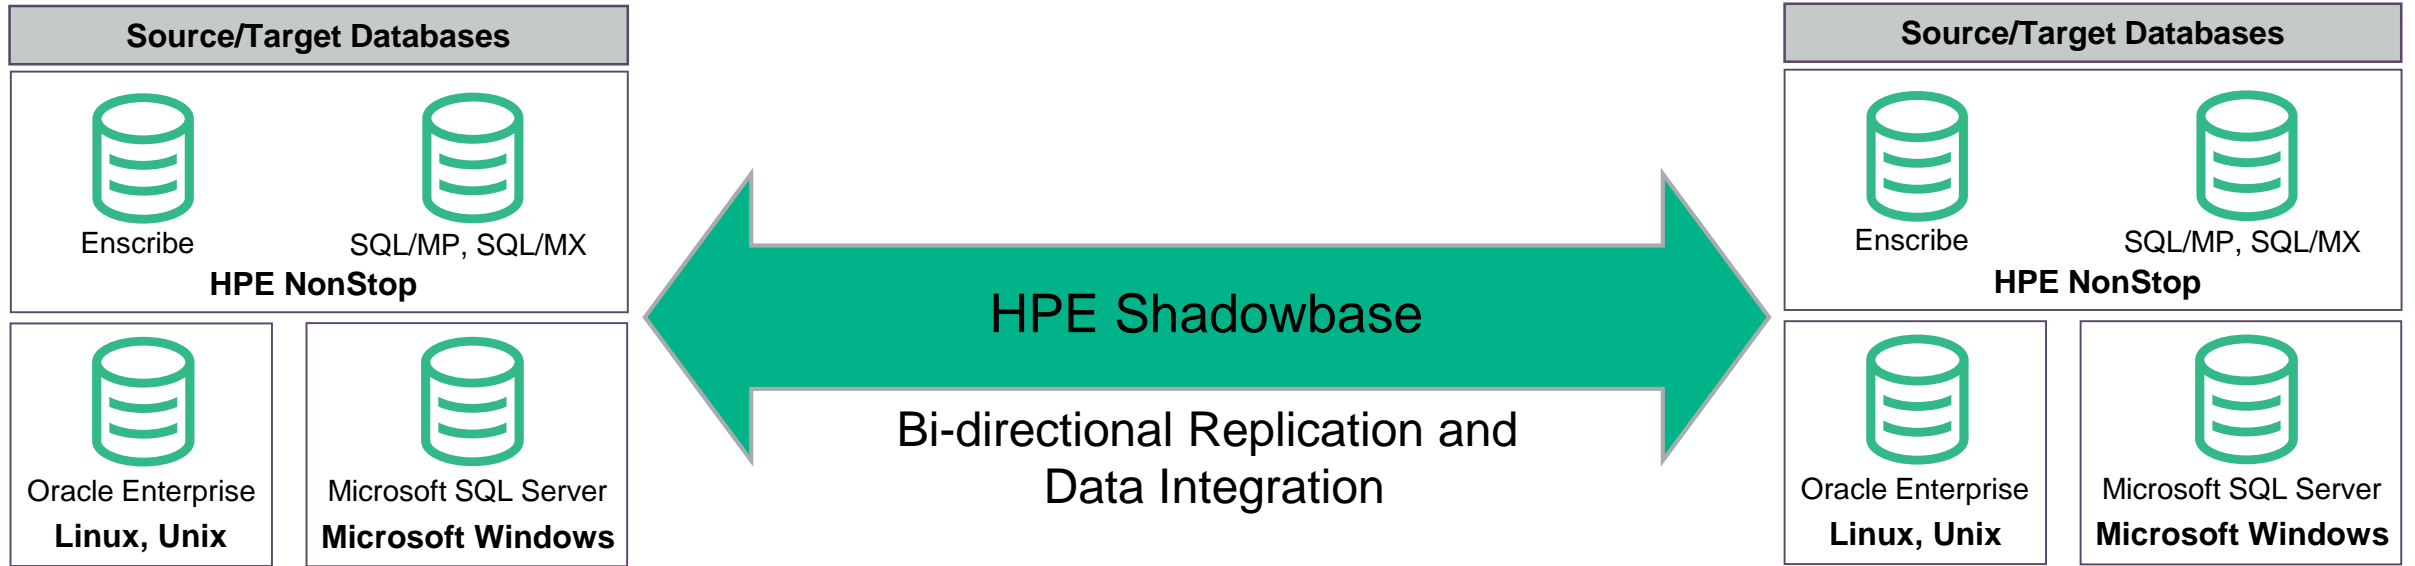
Manage

- Review what data was changed, and when
- Validate two data sets
- Convert and map data between environments
- Extract, transform, and load (ETL support)
- Restore corrupted databases (while online)

Homogeneous & Heterogeneous Uni-directional Data Replication



Homogeneous & Heterogeneous Bi-directional Data Replication





Summary & for more information

Why choose Shadowbase?

Proven technology

- Shadowbase is deployed at hundreds of sites, including many of the most-demanding HPE NonStop sites

Flexible solutions for your business challenges

- Business continuity, data integration & synchronization, data warehouse feeds, application integration, real-time business intelligence



Global sales organization

- Global reseller presence from HPE Sales (or HPE’s regional reseller)

Global 24x7 support & professional services organizations

- Global support presence from the HPE GNCS (or HPE’s regional reseller)
- Global Professional Services from HPE Pointnext and HPE SDI (or HPE’s regional reseller)

Affordable and committed to the HPE NonStop platform

- Improves TCO via overall cost advantage and features
 - “One product, many solutions”
- We are partnering and investing with HPE in many innovative enhancements



For more information – marcom 1 of 3

Preparing for your HPE Shadowbase experience



Business Continuity

Disaster Recovery (DR), Sizzling-Hot-Takeover, and Active/Active application availability



Data Integration

Replicate data from one DB to another while transforming, scrubbing, or filtering



Application Integration

Share data among applications
(access to core data)

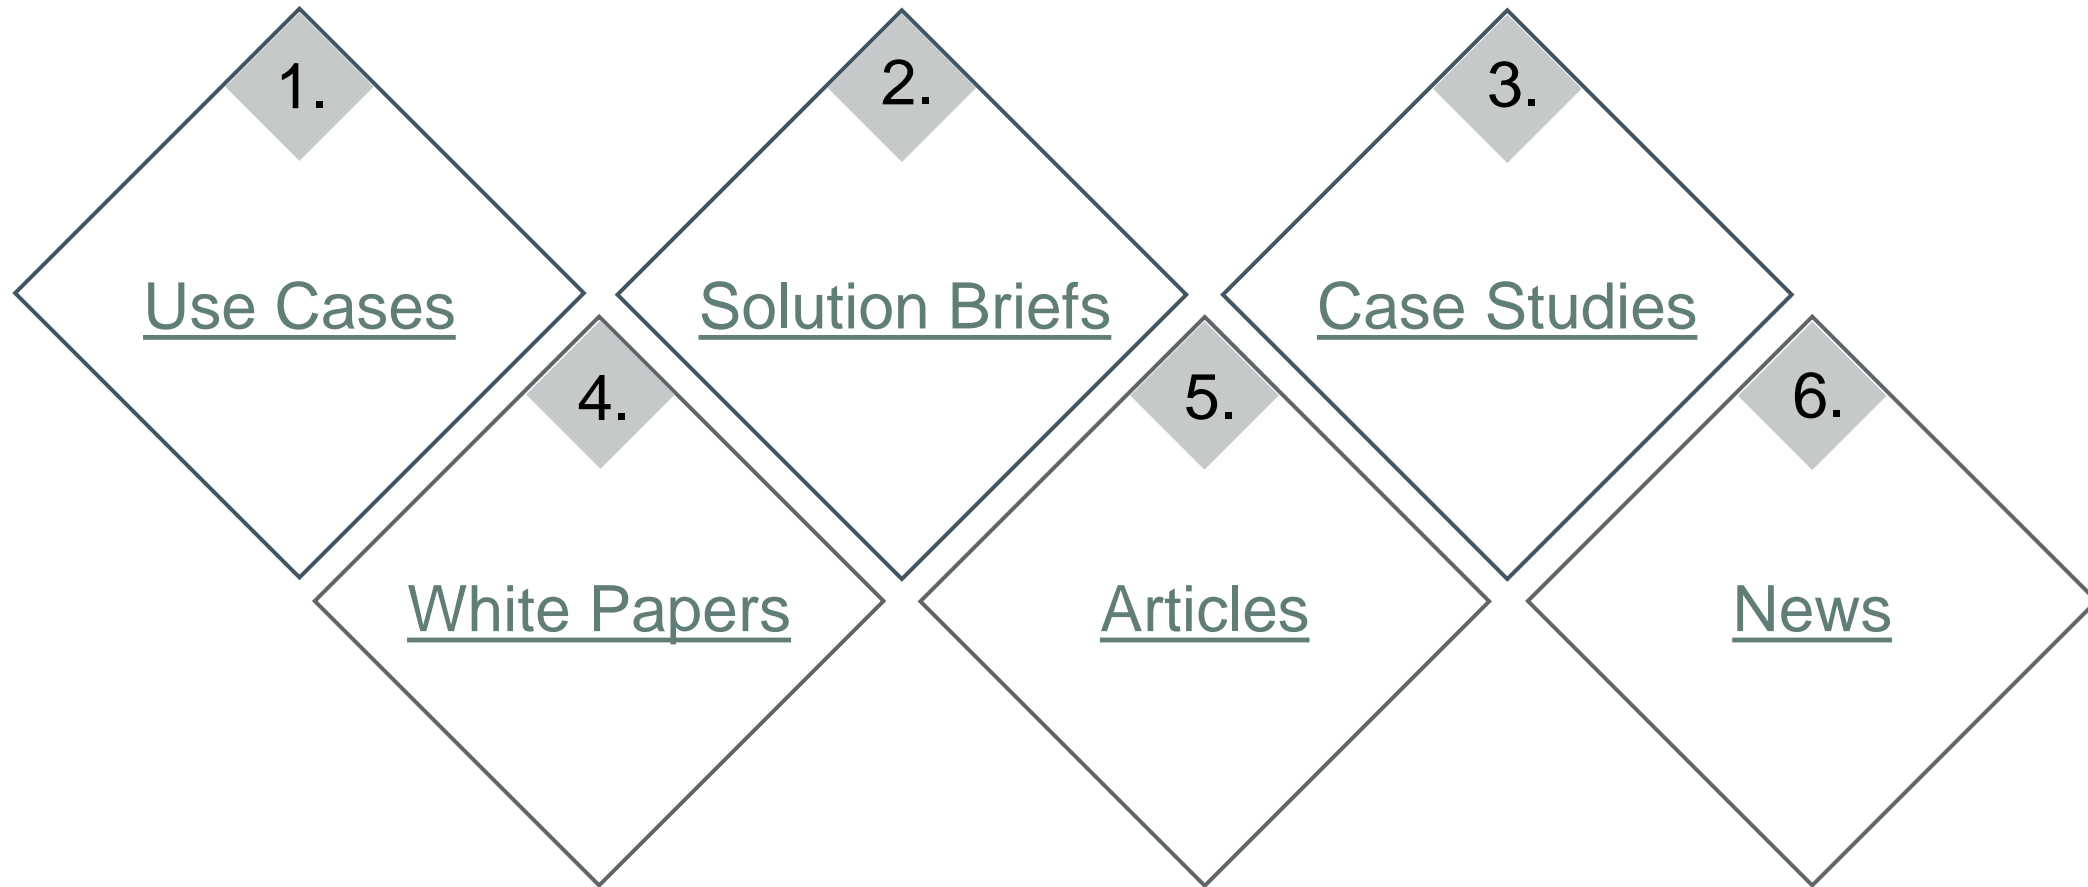


Utilities

Manage, monitor, and correct data


For more information – marcom 2 of 3

Preparing for your HPE Shadowbase experience



For more information – books 3 of 3

Breaking the Availability Barrier Book Series




Volume 1

Survivable Systems for
Enterprise Computing



Volume 2

Achieving Century Uptimes
with Active/Active



Volume 3

Active/Active Systems in
Practice



Hewlett Packard
Enterprise

Thank you

SBProductManagement@gravic.com