

DATA

Distributeur agréé de HVR en France.

HVR





- Introduction
- HVR Overview
- Technology
- Use Cases
- Case Studies
- Contact



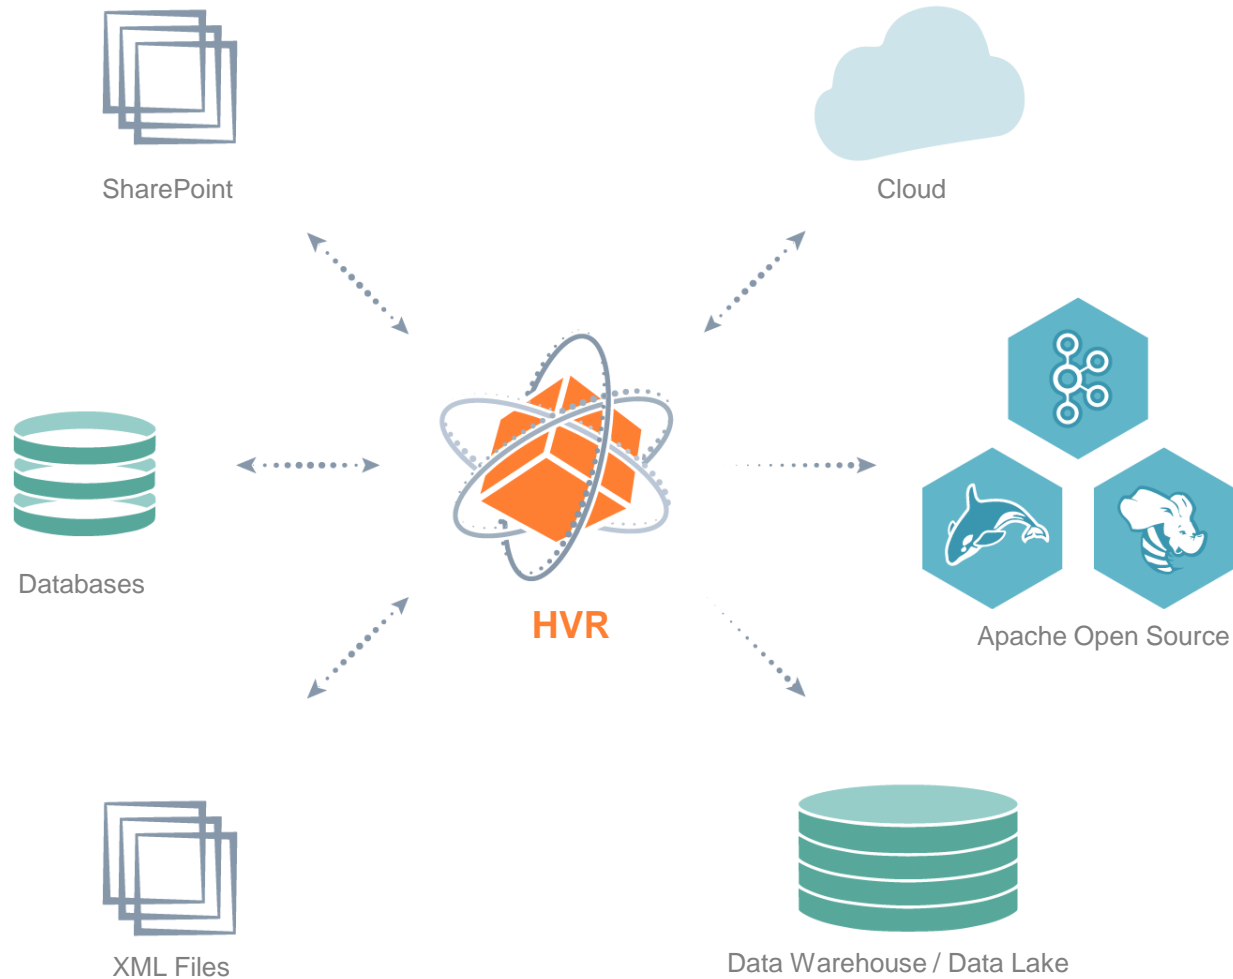


HVR OVERVIEW

Introduction



Fast Data Movement for Real-Time Updates



 **> 150** Customers

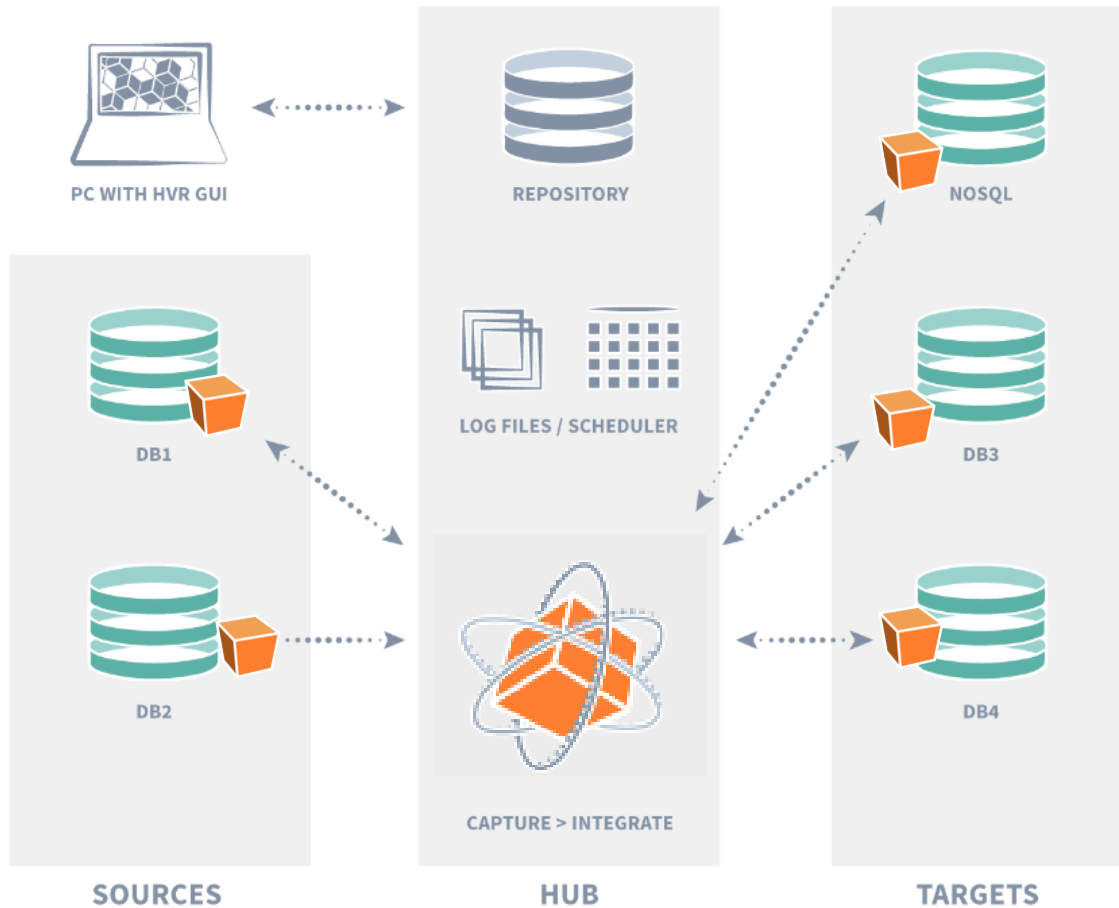
 **> 1000** Deployments

 **> 30** Countries



TECHNOLOGY

Overview



➤ Solution


- Fast – distributed, efficient
- Optimized network communication

➤ “All in One Box”

- Table Creation and Initial Data Load
- Change Data Capture
- Data Validation and Repair
- Automatic monitoring and reporting



Peer-to-Peer



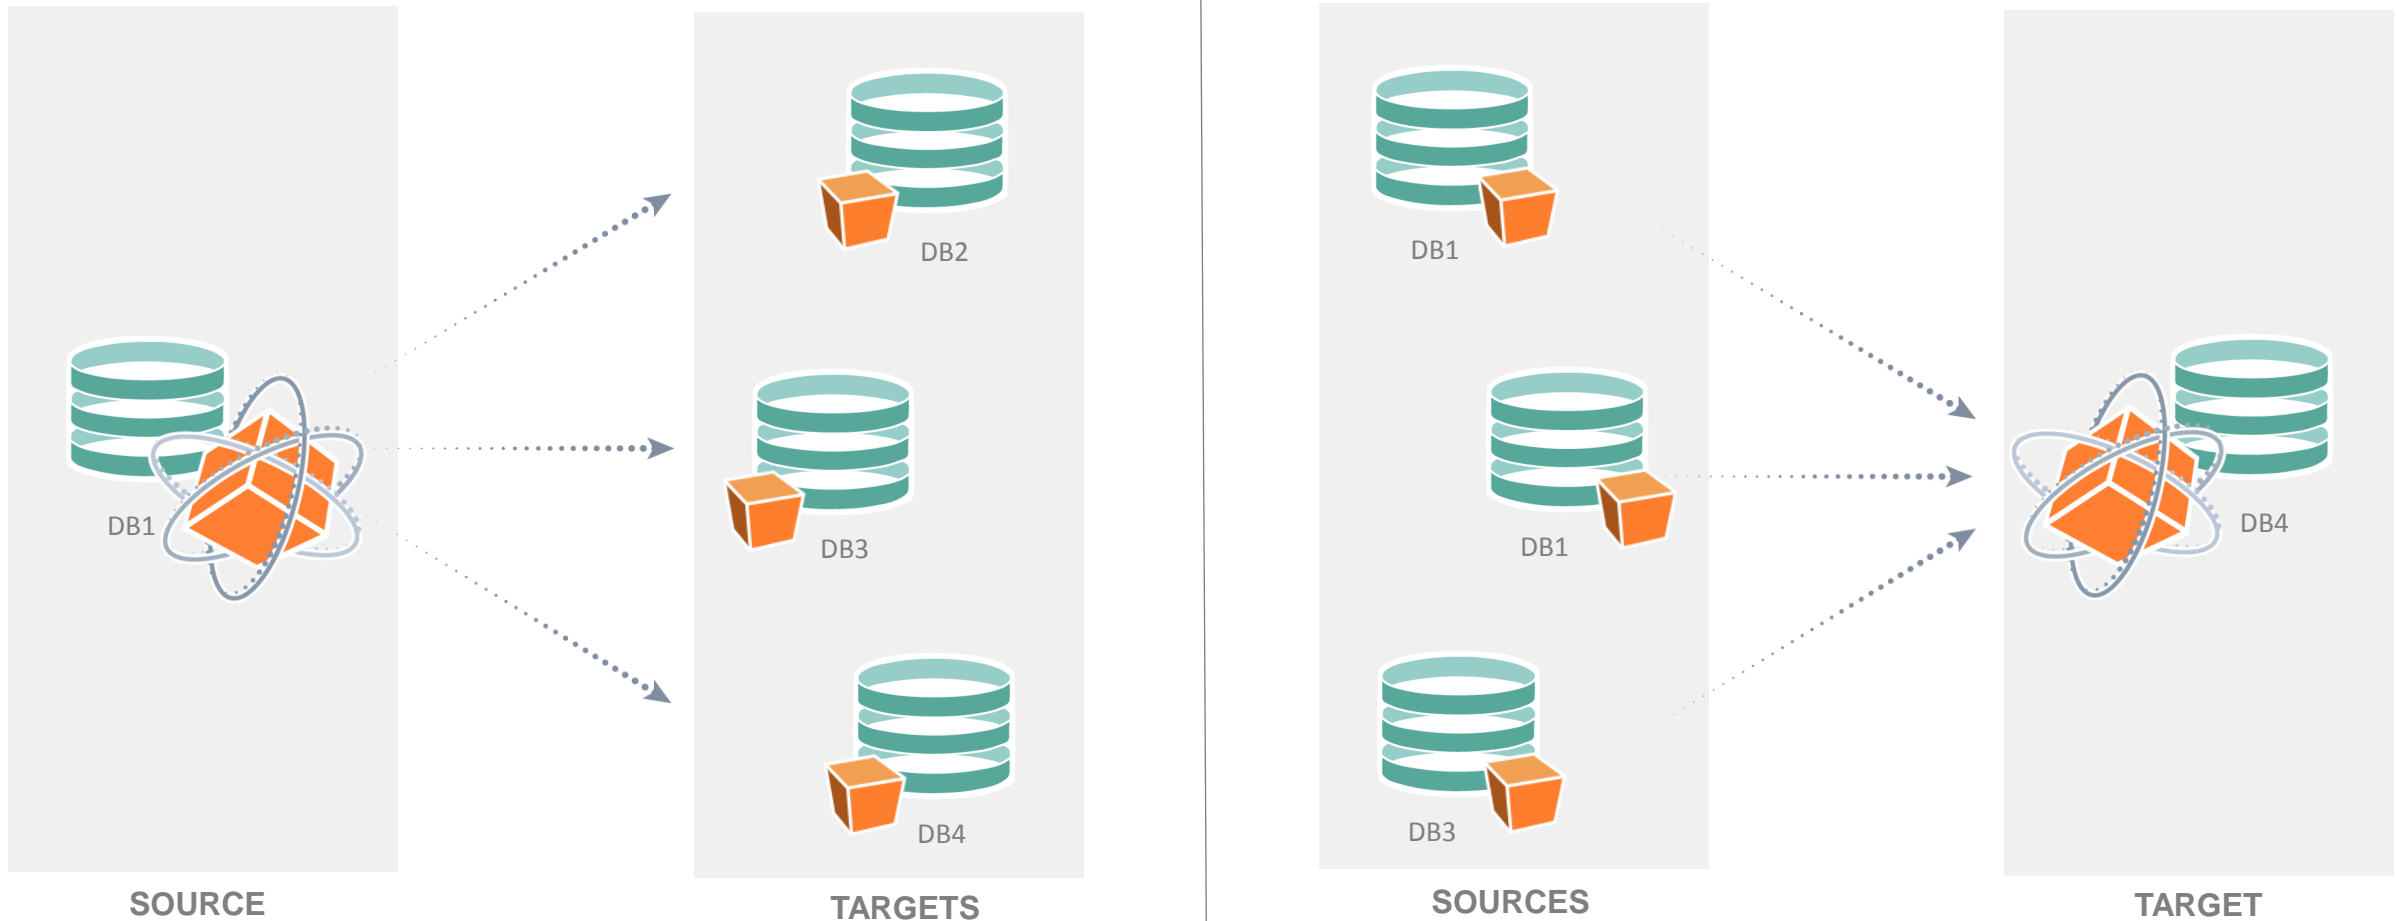
Central Hub Architecture





Flexible Deployment Options

➤ Hub can be co-located with source server, target server or on a mid-tier server





UNI-DIRECTIONAL

Real-Time Reporting
Migrations (reverse post migration)



BI-DIRECTIONAL

Active / Active Standby
High Availability



PEER-TO-PEER

Multi-way Active / Active
Geographical Distribution



BROADCAST

Data Distribution



INTEGRATION / CONSOLIDATION

Data Warehouse / Data Lake



CASCADING

Data Marts







HVR's Technology Benefit

- Least resource utilization thanks to efficiency
 - Log-based Change Data Capture for least impact on the source
 - Dense compression for lowest possible bandwidth utilization
 - Leverage target system direct load for low latency and fast performance
- Innovative architecture to ensure no data loss
 - Automatic retry to minimize latency and maximize up-time
 - Initial load integrated with real-time change replication for zero transaction loss
- Ease of use and automation for productivity and low TCO
 - Graphical User Interface for ease of use, even in complex setups
 - Design once, reuse many times for productivity
 - Automatic monitoring for reactive management



Wide Heterogeneous Support



GREENPLUM
DATABASE



Amazon
Redshift



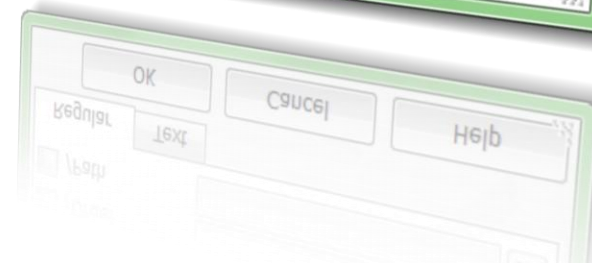
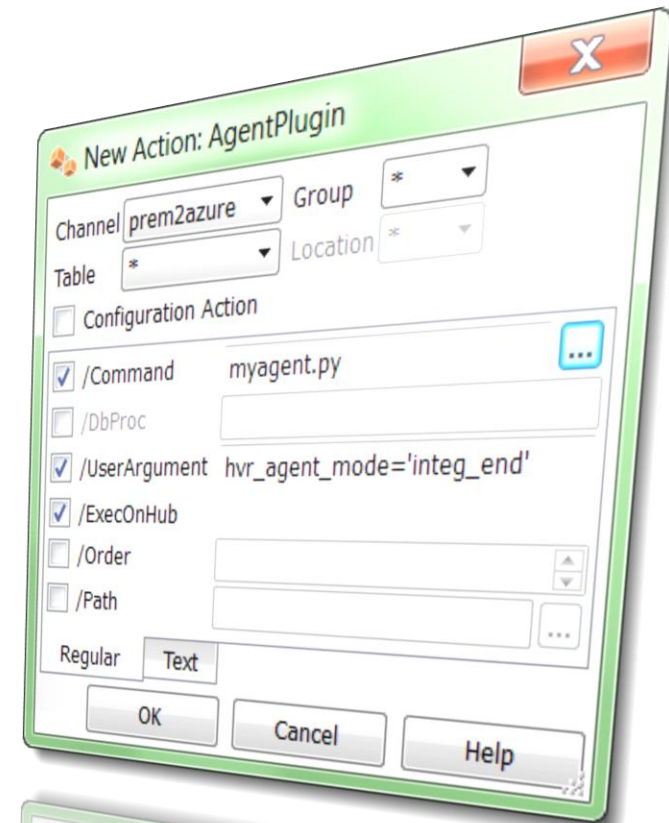
**Files: Avro, CSV,
HDFS, S3, XML**





➤ Leverage orchestration and centralized logging

```
189 def main(argv):
190     env_load();
191     #env_var_print();
192
193     if ((argv == "refr_write_end" or argv == "integ_end") and
194         os.getenv('HVR_FILE_NAMES') != ''):
195         file_loc_processing()
196         if (file_counter > 0) :
197             print "Successfully transmitted {0:d} file(s)".format(file_counter)
198
199     # just recovery stage
200     if (argv == "integ_begin") :
201         file_loc_processing()
202         if (file_counter > 0) :
203             print "Successfully transmitted {0:d} file(s)".format(file_counter)
204
```





Heterogeneous Initial Loads

- Create and load target tables
- Use to initiate real-time replication
- Use independent of replication
- Schedule during off-work hours

Table Name	Target	State	Rows on ord	Duration
mssql	Finished	30000	5.1	
mssql	Finished	10	0.1	
mssql	Finished	30000	1.0	
mssql	Finished	100000	3.8	
new_order	mssql	Finished	9000	0.2
order_line	mssql	Finished	299786	9.1
orders	mssql	Finished	30000	1.0
stock	mssql	Finished	100000	7.5
warehouse	mssql	Finished	1	0.1

Heterogeneous Data Validation and Repair

- Protect against human error
- Validate the data before you complete your migration
- Compare against mixed technologies

HVR Compare for channel orcl2ms

Location

Location	Class	Node
<input type="checkbox"/> mssql	sqlserver	
<input checked="" type="checkbox"/> ord	oracle	

Location

Location	Class	Node
<input checked="" type="checkbox"/> mssql	sqlserver	
<input type="checkbox"/> ord	oracle	

Table Name Base Table Name

All Tables

Options

Bulk Granularity

Row by Row Granularity

Verbose

Apply All

Compare Db Sequences

Schedule Compare Jobs

Taskname cmp

Compare Schedule Close Help

hvrcompare -qb -P4 -r orcl -l mssql -h sqlserver \sqlhub orcl2ms

Compare Result

Table Name	Target	State	Rows on ord	Rows on Target	Duration
customer	mssql	Different	30000	30000	3.2
district	mssql	Different	10	10	0.0
history	mssql	Different	56367	50230	1.0
item	mssql	Identical	100000	100000	0.8
new_order	mssql	Different	9460	9415	0.0
order_line	mssql	Different	560089	497398	4.2
orders	mssql	Different	56110	49815	0.6
stock	mssql	Different	100000	100000	2.0
warehouse	mssql	Different	1	1	0.1

Finished

Compare Result

Table Name	Target	State	Rows on ord	Rows on Target	Duration
customer	mssql	Identical	30000	30000	0.5
district	mssql	Identical	10	10	0.0
history	mssql	Identical	45764	45764	0.2
item	mssql	Identical	100000	100000	0.5
new_order	mssql	Identical	9462	9462	0.0
order_line	mssql	Identical	454077	454077	2.8
orders	mssql	Identical	45462	45462	0.2
stock	mssql	Identical	100000	100000	1.0
warehouse	mssql	Identical	1	1	0.0

Close

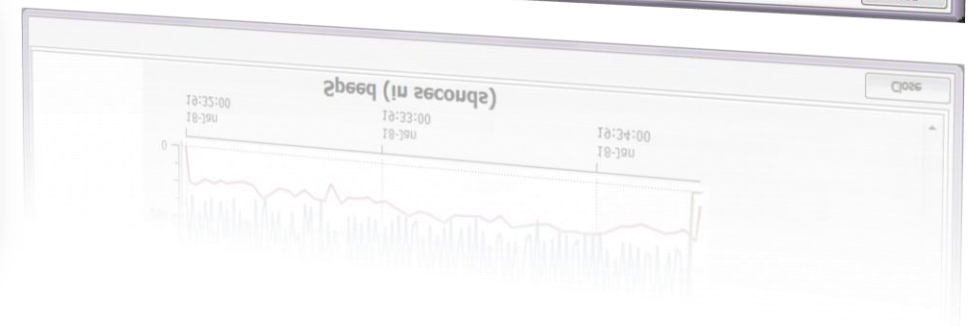
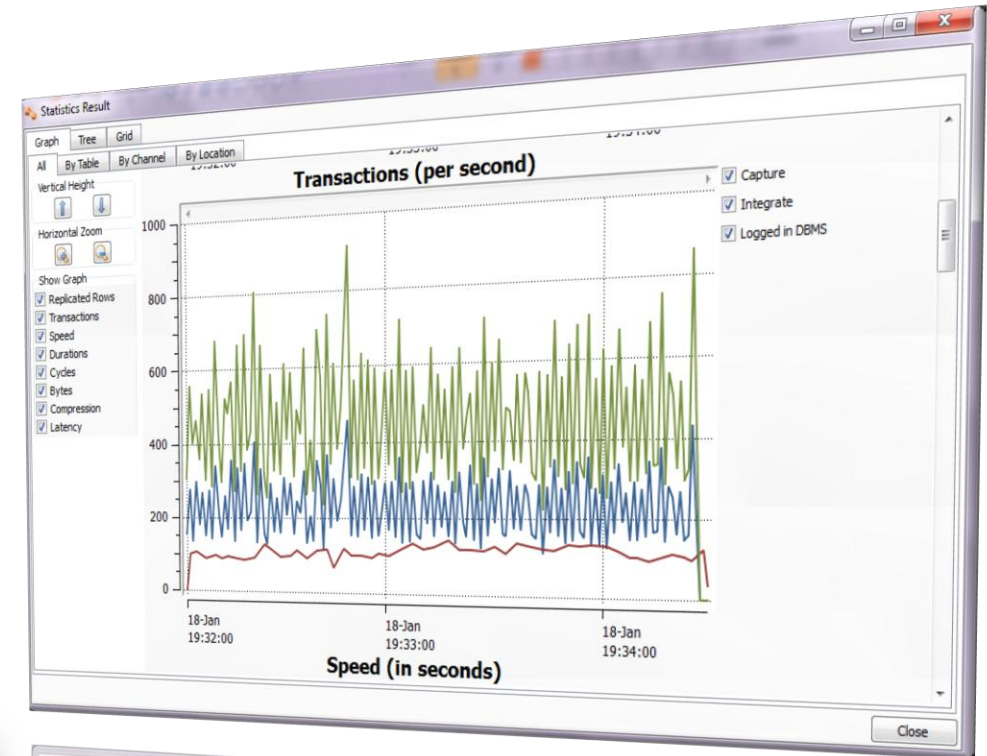


Built-In Monitoring and Alerting

- Real-time monitoring of HVR processes
- Automatic alerting and notification
- Integration with enterprise monitoring solutions
- Report on replication statistics

Job	State	Retries	Recent Error	Latency
orcl2ms-cap-orcl	SUSPEND			42m 5s
orcl2ms-cmp-orcl-mssql	PENDING			
orcl2ms-integ-mssql	SUSPEND			

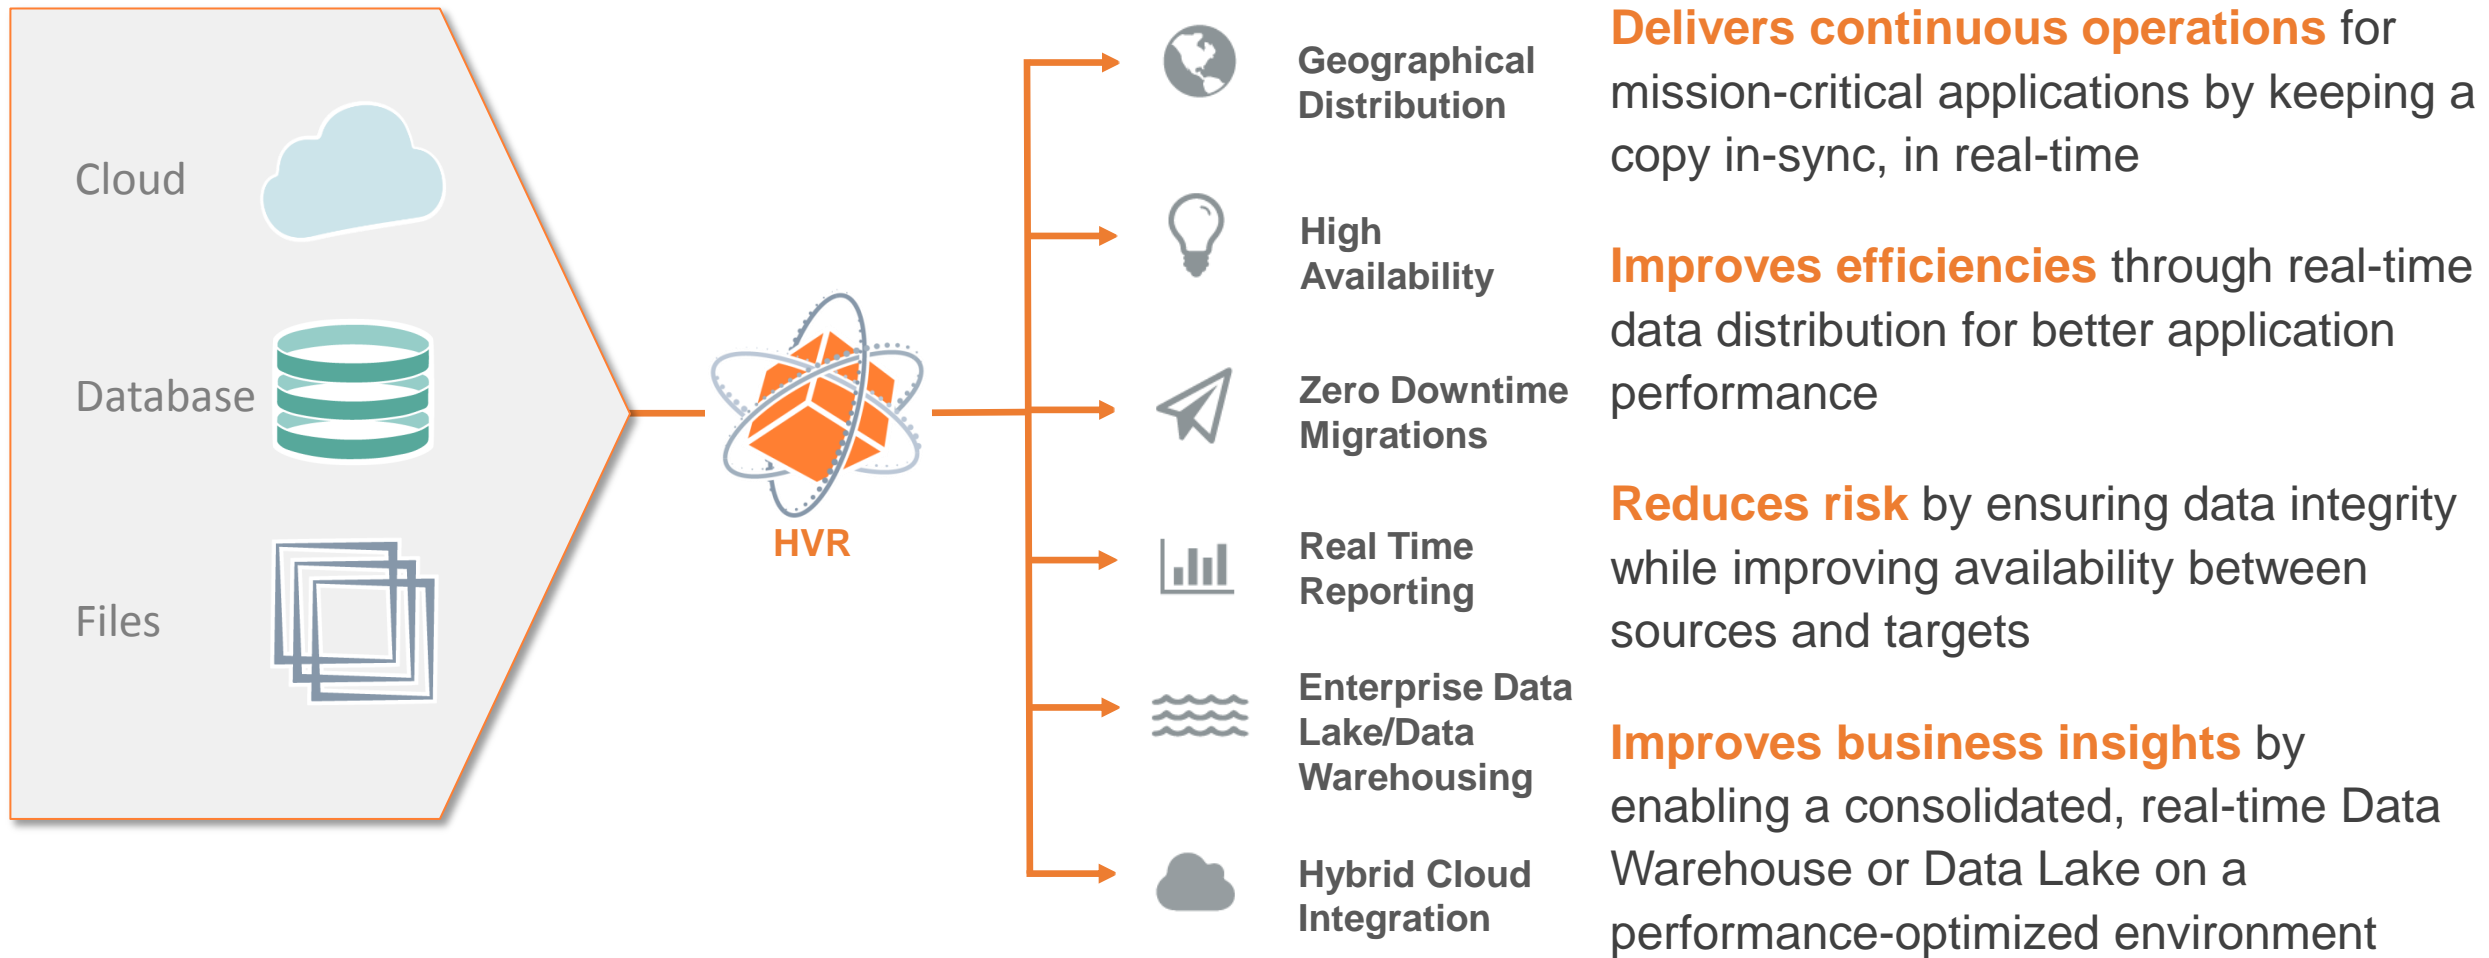
Job	State	Retries	Recent Error	Latency
orcl2ms-cap-orcl	RUNNING			1s
orcl2ms-integ-mssql	RUNNING			3s





HVR USE CASE

Overview





HVR USE CASE

Geographical Distribution



➤ Solution

- Log-based CDC and integration at multiple sites
- Efficient network utilization
- Active-Active with conflict detection and resolution
- Filter data for different requirements
- Centralized design, monitoring and manageability

➤ Benefits

- Optimum application performance
- Data available where you need it, when you need
- Real-time data, synchronized for 24x7 business continuity





HVR USE CASE

High Availability



➤ Solution

- Committed transactions are only replicated to the target
- Active-Active to fully utilize DR and load balance workload
- Conflict detection and resolution to support active sites

➤ Benefits

- No distance restrictions
- Different database versions
- Different database vendors
- Target is fully read/write





➤ Solution

- Capture from Active Data Guard Standby
- Committed transactions are only replicated to the target(s)

➤ Benefits

- No distance restrictions
- Combination of synchronous for local and asynchronous for long distance replication





HVR USE CASE

Zero Downtime Migrations/Upgrades

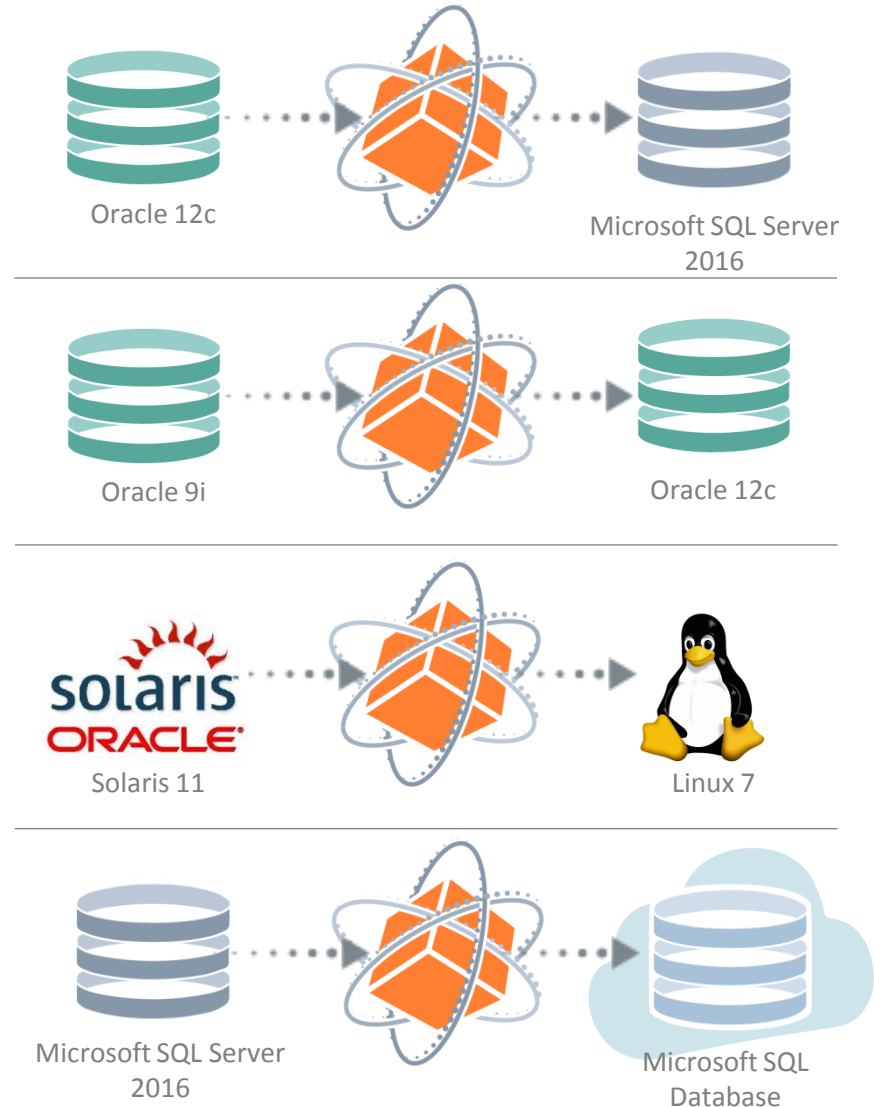


➤ Solution

- Heterogeneous initial load and table creation for instantiation of target
- Change data capture to keep source systems on-line
- Mitigate risk with failback option
- Data validation and repair to verify data consistency
- Active-Active for phased switchover and keep both systems synchronized

➤ Benefits

- Reduce and even eliminate downtime
- All in one solution
- Automatic reporting and notification





Uni-directional



- Minimize downtime by keeping source on-line
- Allows for extended testing period
- Downtime incurred only when switching users to new system
- Data validation before switchover

Uni-directional with Failback



- Minimize downtime by keeping source on-line
- Allows for extended testing period
- Failback to old system to mitigate risk of failure on new system
- Downtime incurred only when switching users to new system
- Data validation before switchover

Active-Active



- Eliminate downtime and keep both systems on-line
- Allows for extended testing period
- Phased migration approach, move users across in batches
- Account for conflicts and resolution
- No downtime incurred
- On-line data validation



HVR USE CASE

Real-Time Reporting



➤ Solution

- Non-intrusive log-based CDC from major mega vendors
- Filter on data that is only required for reporting
- Delivery to low costing platform
- Optimized target structures for faster reporting

➤ Benefits

- Access to real-time updated data to make quicker decisions
- Remove vendor lock-in and reduce license costs





HVR USE CASE

Real-Time Data Warehousing/Data Lake




➤ Solution

- Non-intrusive log-based CDC from heterogeneous sources
- Transactions streaming in near real-time into the data warehouse
- Augment with 3rd party ETL tools for 'micro-batch' operations

➤ Benefits

- Remove nightly 'batch window' operations
- Low impact on OLTP source systems
- Real-time data for quicker decision making
- Optimized reporting environments



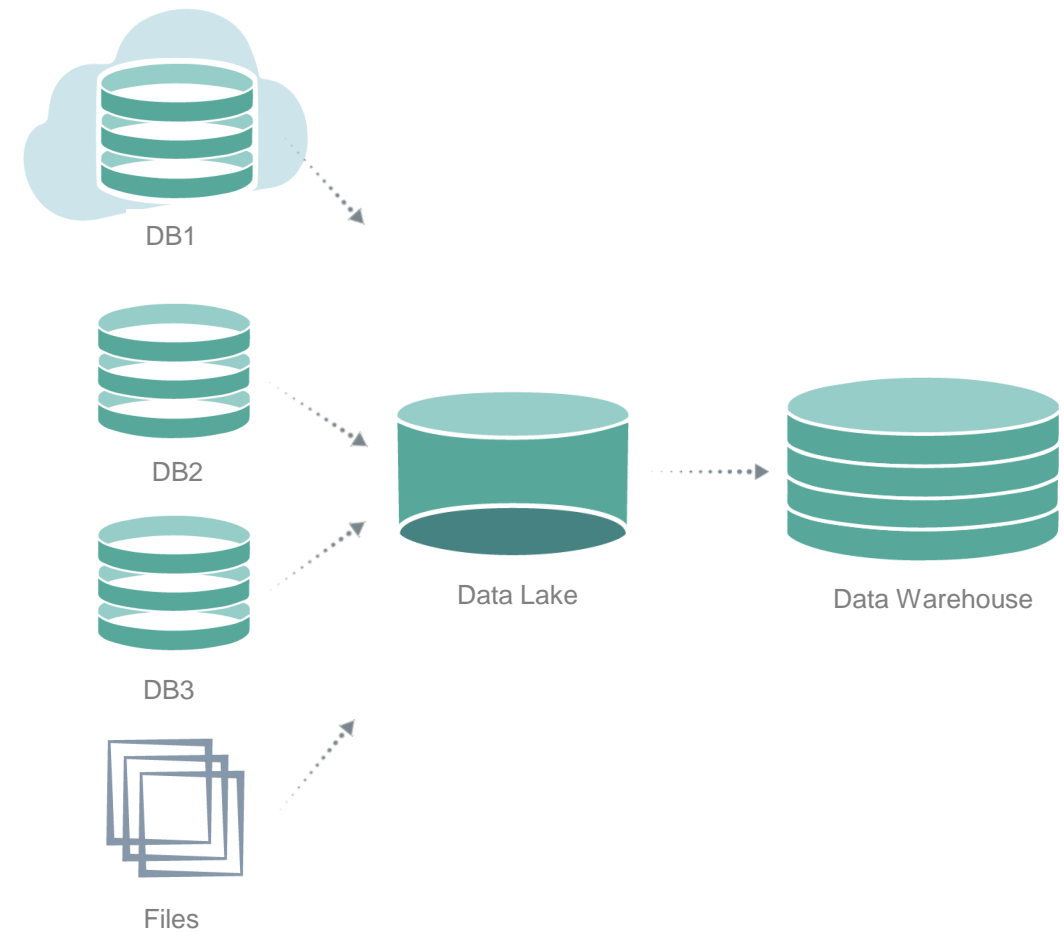



➤ Solution

- Data consolidation from various heterogeneous environments
- Optimized, parallel data integration
- Supports multiple Big Data platforms
- File replication

➤ Benefits

- 360 degree view of the data
- Scalable data retrieval
- Low-impact and real-time onboarding of data






- Real-Time Integration into Hadoop
 - Support for HDFS (files)
 - Support for HIVE (SQL)
 - Support for HBASE (NoSQL)
 - Support for Kafka (publish/subscribe)
 - HVR Agent Plugin API (any other target)
- Always in sync – low latency
 - Non-intrusive real-time log-based change data capture
 - Efficient large block transfer with dense data compression
- High speed Throughput
 - Load data transactions row-by-row, or in batch (burst)



HVR USE CASE

Hybrid Cloud Integration



- Available where you need it
 - On AWS and Azure Marketplace
- Always in sync – low latency
 - Non-intrusive real-time log-based change data capture
 - Efficient large block transfer with dense data compression
- Secure
 - Encrypted network traffic
- Low migration risk
 - Synchronize on-premises and cloud for testing and initial setup
 - Synchronize cloud and on-premises as a fallback option
 - Inter-cloud data integration



CASE STUDIES

Customer Success Stories

Customer Use Cases




Moved Enterprise ERP into a cross division **Data Lake** spanning 100+ sources

WELLINGTON
MANAGEMENT®

Replaced Shareplex for service, support, functionality and **lower TCO**



Cloud data integration and migration from on-prem



High availability and real-time reporting on mission-critical maintenance data.



Data integration for flight plans for global airlines. Half the world's airplanes wouldn't fly without it.

gemalto
security to be free

Saving customers' money and optimizing data availability by performing **zero downtime migrations** with HVR.

CEB | SHL Talent Measurement

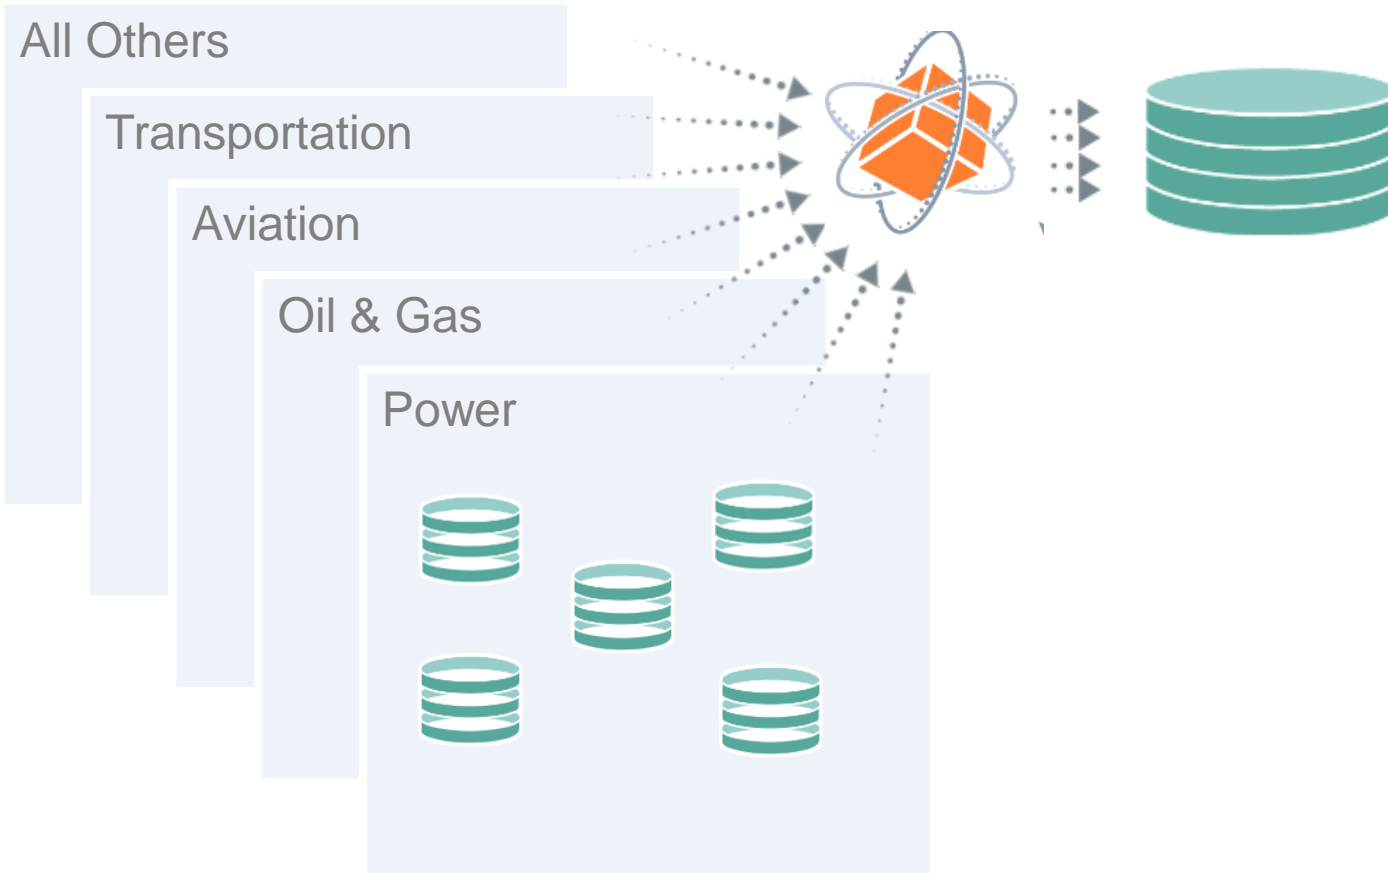
Improving system utilization by enabling **Passive Failover Clusters**

Kiewit

Real-time data integration from diverse databases into an EDW on Teradata for up-to-date consolidated data.



Case Study | GE Corporate Data Lake




Objective: Cross-divisional Data Lake

- Dozens of mostly ERP sources
 - SAP on Oracle and DB2i
 - Oracle Applications
 - Industry-specific applications
- Target Greenplum
- Customer Benefits
- Efficient and flexible in a complex setup
- Performs well for MPP target
- Cost-effective to setup and maintain



Case Study | Lufthansa Flight Planning



Objective: Distribute flight planning data

➤ Critical: airplanes don't take off without it

- Over (60) Airlines use LIDO, including China Southern, China Airlines, Emirates, British Airways, Air France, UPS

➤ Secure: Confidential data cannot be shared

➤ In use 10+ years


Customer Benefits

- Global real-time data replication
- Secure and reliable
- Complex routing and filtering, for some airlines with cascading setup



Case Study | Wellington Real-Time Data

ORACLE
EXADATA
OPTIMIZED



Objective: Real-Time Data for Traders

- Critical: supports core business of trading
- High volume
 - 10+ Oracle Exadata 11g and 12c RAC sources and targets
 - Busiest systems up to 100 MB of redo per second (rate of 8.5 TB/day)
 - Data Warehouse processing up to 50k transactions/s

Customer Benefits

- Outperformed previous Shareplex solution
- HVR is true partner to be successful
- Flexible technology to leverage Cloud in near future and lower dependency on Oracle



Less than 100MB our streamlined download includes everything:

- **Comprehensive** heterogeneous platform support, secure, high speed real-time database, big data and file replication
- Built in heterogeneous database **data validation and repair**
- Heterogeneous **initial loads**, schema creation, auto mapping
- **Heterogeneous DDL** replication
- Built in graphical **monitoring**, alerting and scheduling
- Learn once, **use everywhere**



DATA,
12, avenue Raspail 94250 GENTILLY, France

+33 1 57 19 59 38

contact@data.fr

www.data.fr

www.hvr-software.com