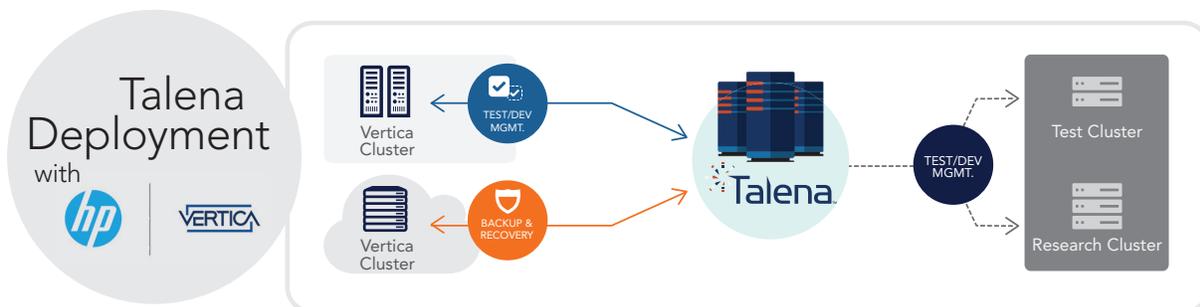# Talena Deployment With HPE Vertica

HPE Vertica is an ideal database for large, rapidly growing data sets that require fast query performance. The database has seen great traction for applications in gaming, social media, entertainment and financial services. In most cases, these applications are critical to the success of the business. Downtime or data loss can severely impact revenue and your business reputation. Hence, constant availability and recoverability of your HPE Vertica infrastructure is critical to sustaining business operations.

To ensure data recoverability from user error or application corruption, a flexible backup and recovery solution that goes beyond the limited capabilities of today's solutions is required. Having multiple replicas protects against hardware failures but won't prevent data loss if a user accidentally deletes data in a table or an application corrupts data in your Vertica database. Implementing the wrong solution can have catastrophic consequences both in terms of data recoverability as well as making copies of production data for non-production use.

An enterprise grade backup and recovery solution for Vertica must address common issues encountered by customers.

1. The solution must be designed with both recoverability and backup in mind, versus just the latter. No matter how well a solution can back up data, if data recovery is inflexible, manual or onerous, it can result in downtime and data loss and won't meet the service level agreement needs of enterprise applications

2. IT best practices require that multiple point-in-time backup copies be stored on a system different from the production system. The ability to easily recover individual tables from these backups and to a cluster of different size and configuration is essential to protect against user error or application corruption. Storage reduction techniques such as de-duplication and compression are essential to store these backups efficiently.

3. An ideal backup solution must also natively facilitate database cloning for non-production use cases such as test, development & disaster recovery. In case of test & development, making a full clone is not required and can pose compliance risks especially if the database contains personally identifiable information (PII) or other confidential data. Cloning must include the ability to mask confidential data and create smaller databases (via sampling or filtering of data).

Organizations using HPE Vertica deploy the Talena software to provide self-service access to production data for sandbox use and automate the backup and recovery process. In addition to HPE Vertica, Talena also supports data stored in HDFS, Hive and Impala, as well as other NoSQL data stores like Cassandra. Via the Talena metadata catalog, Talena FastFind™, companies can quickly identify data sets they want to recover from backups. Backups are stored in a de-duplicated and compressed format and are erasure coded to protect against hardware failures. Talena makes it very easy for development and data science teams to quickly gain access to production data in sandboxes to support application iteration needs. With its ability to mask and filter data consistently, Talena supports the ability to provide clean data for non-production use.

KEY BENEFITS

Rapidly iterate applications
with self-service access to production data on your sandboxes & DR clusters.

Eliminate scripting
and other manual processes to copy data

Flexible backup and recovery
allows backup and recovery of individual tables or the whole database to the original or alternate location of a different size or configuration

Storage optimization
is typically 5x through compression, de-duplication and erasure coding, saving CapEx and OpEx

Reduce Compliance Risk
through masking of confidential and personal data

Please contact us for more information at **info@talena-inc.com** or visit us at **www.talena-inc.com**.