



# Insights Data Storage

## Highly Scalable Network Intelligence Storage

### INSIGHTS DATA STORAGE BENEFITS

- Big data platform designed for storing Sandvine's unique network and QoE intelligence in large-scale deployments
- High availability options to ensure the appropriate level of mission-critical redundancy, ranging from a single node up to multi-node clusters
- Advanced reporting and analytics with parallelized query execution for timely business decisions
- Flexible API integration for enriching existing data lakes and big data systems

**Insights Data Storage is the highly scalable data storage component of Sandvine's Insights Portfolio. Insights Data Storage provides long-term data storage for Sandvine's unique network intelligence and quality of experience (QoE) metrics, enabling rapid data visualization as well as extensive data exploration.**

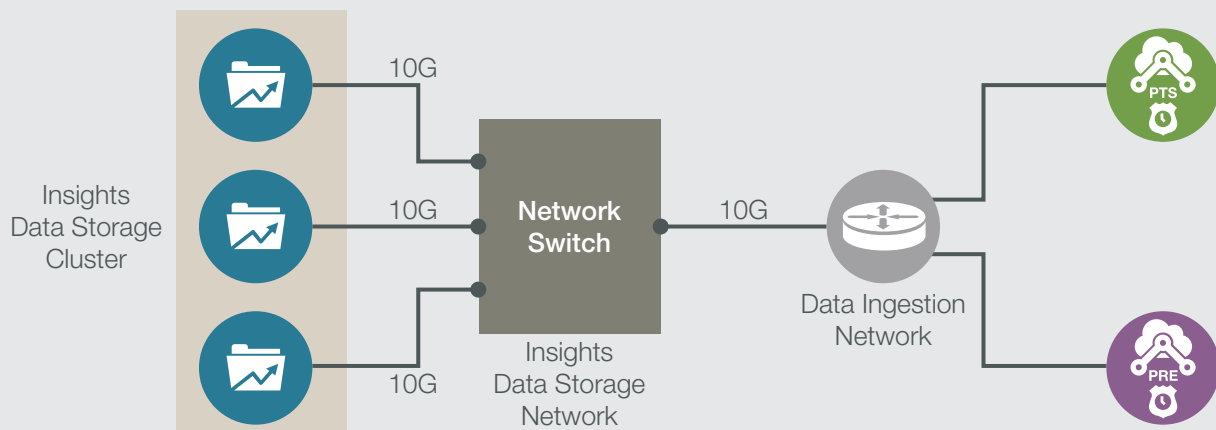
Insights Data Storage is designed for large-scale, carrier-class networks, where it can be used in conjunction with big data systems or where it fulfills the role of Sandvine's contextual network intelligence data lake. It uses a columnar database designed for exploratory analytics use cases, providing storage and retrieval of large volumes of granular information with segmentation and complex data structures (e.g., per-subscriber, per-location, and per-protocol reporting).

Insights Data Storage includes a number of capabilities specifically designed to simplify interoperability with existing big data solutions deployed by network operators. Sandvine's network intelligence, foundational to all use cases, is organized in a fixed schema, supplying easy-to-use data for a network operator's analytics requirements with enrichment for actionable intelligence without complex integration projects.

Sandvine's Insights Product Family visualizes the Active Network Intelligence use cases deployed in telecommunications networks, empowering operators to solve problems and spot opportunities to attract more subscribers, generate revenue, improve QoE, and reduce CAPEX. Insights Data Storage provides the foundation for managing Sandvine's network intelligence, allowing network operators to retain the right data, ensuring a high quality network experience is delivered.

Figure 1

Typical deployment of Insights Data Storage using a dedicated network switch for the Insights Data Storage cluster, with 10 Gbps physical links for data ingestion and cluster communication.



## Designed for Carrier-Class Network Intelligence Storage

Network operators running their business on network intelligence need the highest levels of performance and availability. Insights Data Storage delivers both and ensures the right intelligence is available to make the right business decisions.

### Clustering

Flexible clustering options optimize performance and high availability by combining many Insights Data Storage instances together to scale storage to meet carrier-class needs. By enabling a practically limitless number of nodes to operate as a unit, allowing for RAID-like striping and massive parallel processing for fast data retrieval, the largest networks in the world can maintain full network visibility.

### High Availability

A built-in failure recovery function ensures that Insights Data Storage is reliable and data is always available.

### Hierarchies

Hierarchies allow operators to segment network reporting from local regions while maintaining comprehensive site-wide visibility. It also allows operators to optimize different parts of a cluster based on scale requirements (e.g., smaller databases for smaller regions, reporting into a larger database for the entire network).

### Simple Fixed Schema

Unique KPI enrichment for context, including per-subscriber, location, application, service plan, device, and many more options, ensures the intelligence is actionable in near real-time.

### ODBC Connector for Big Data Connectivity

ODBC access is provided to the underlying data within the Insights database. This interface can be used by both Sandvine and the operator's third-party visualization and reporting tools to create custom reports or perform custom data analysis, depending on the operator's requirements. Sandvine provides ultimate flexibility for leveraging the value of the underlying data from the network.

### Zero Administration

Self-diagnostics and automated clean-up of the database is supported without the need for additional tools (e.g., no need for re-indexing or refreshing of data). Additionally, the installation self-tunes based on available resources and defined roles. Zero downtime is supported since the database continues to run query requests while rebalancing is in progress (e.g., during upgrades or when taking down a node for maintenance).

### Elastic Data Storage

Insights Data Storage automatically rebalances data when nodes are added, removed, replaced, or repaired, and clones databases from one cluster to another.

### Advanced Compression

It replaces slower disk I/O with fast CPU cycles through aggressive compression, reducing the disk space needed for storage relative to traditional databases and allowing smaller disks to be used for storage.

### Distributed Query Execution

It takes advantage of massive parallel processing for faster query execution and reporting by leveraging all nodes of the Insights Data Storage cluster.

## ABOUT SANDVINE

Sandvine helps organizations run world-class networks with Active Network Intelligence, leveraging machine learning analytics and closed-loop automation to identify and adapt to network behavior in real-time. With Sandvine, organizations have the power of a highly automated platform from a single vendor that delivers a deep understanding of their network data to drive faster, better decisions. For more information, visit [sandvine.com](http://sandvine.com) or follow Sandvine on Twitter at [@Sandvine](https://twitter.com/Sandvine).



**USA**  
2055 Junction Avenue  
Suite Number 105  
San Jose,  
CA, 95131  
USA

**EUROPE**  
Svärdfiskgatan 4  
432 40 Varberg,  
Halland  
Sweden  
T. +46 340.48 38 00

**CANADA**  
408 Albert Street,  
Waterloo,  
Ontario N2L 3V3,  
Canada  
T. +1 519.880.2600

**ASIA**  
RMZ Ecoworld,  
Building-1, Ground Floor,  
East Wing Devarabeesanahalli,  
Bellandur, Outer Ring Road,  
Bangalore 560103, India  
T. +91 80677.43333

Copyright ©2020 Sandvine Corporation. All rights reserved. Any unauthorized reproduction prohibited. All other trademarks are the property of their respective owners.

This documentation, including all documentation incorporated by reference herein such as documentation provided or made available on the Sandvine website, are provided or made accessible "AS IS" and "AS AVAILABLE" and without condition, endorsement, guarantee, representation, or warranty of any kind by Sandvine Corporation and its affiliated companies ("Sandvine"), and Sandvine assumes no responsibility for any typographical, technical, or other inaccuracies, errors, or omissions in this documentation. In order to protect Sandvine proprietary and confidential information and/or trade secrets, this documentation may describe some aspects of Sandvine technology in generalized terms. Sandvine reserves the right to periodically change information that is contained in this documentation; however, Sandvine makes no commitment to provide any such changes, updates, enhancements, or other additions to this documentation to you in a timely manner or at all.