



Insights Data Storage

Highly scalable application and network intelligence storage

KEY BENEFITS

- Designed for storing Sandvine's unique network and application 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 and application intelligence, including and quality of experience (QoE) metrics, and enables rapid data visualization as well as extensive data exploration.

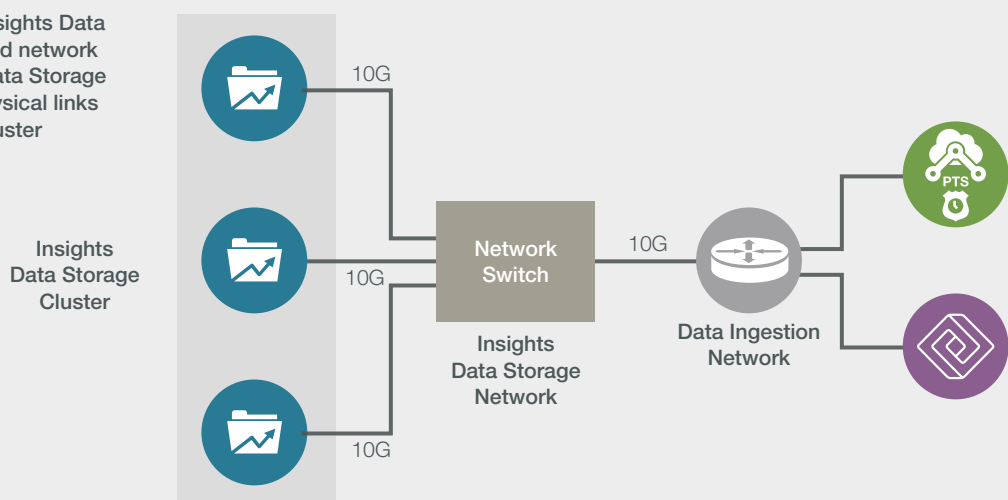
Insights Data Storage is designed for large-scale, carrier-class networks, where it can be used in conjunction with big data systems or where it fulfills the role of Sandvine's contextual network intelligence data lake. It uses a columnar database designed for exploratory analytics use cases, providing 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 several capabilities specifically designed to simplify interoperability with existing big data solutions deployed by network operators. Sandvine's network and application 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 Application and 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 and application 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 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 and network 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 Insights Data Storage reliability and continuous data availability.

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.

Cloud Transformation

Insights Data Storage can run as a cloud-native function (CNF) orchestrated by Kubernetes. All product components have been re-architected to be deployed as separate Docker containers. To provide address scalability and ensure reliability, the database component is deployed as a separate containerized service.

Data Export Options

Big data teams leverage different tools and techniques to generate actionable customer insights by harnessing and integrating new real-time data sources and types.

With Insights Data Storage, these big data teams can employ data export to create custom reports that:

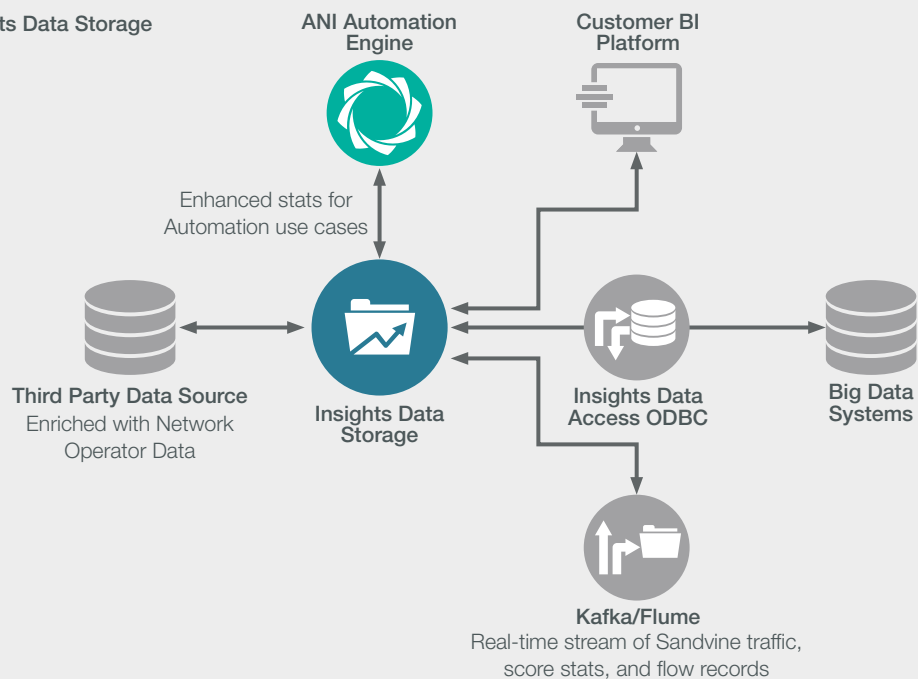
- Are highly personalized
- Blend Sandvine data with data pools from other sources
- Correlate network-centric KPIs with other business and operational KPIs
- Allow operators to experiment with different data visualizations

Data export options include:

- ODBC access to Insights Data Storage (access to traffic/score and use case-specific schema/tables)
- Processed/statistics data export from Insights Data Storage via Kafka

Figure 2

Data Export Options from Insights Data Storage



ABOUT SANDVINE

Sandvine's cloud-based Application and Network Intelligence portfolio helps customers deliver high quality, optimized experiences to consumers and enterprises. Customers use our solutions to analyze, optimize, and monetize application experiences using contextual machine learning-based insights and real-time actions. Market-leading classification of more than 95% of traffic across mobile and fixed networks by user, application, device, and location creates uniquely rich, real-time data that significantly enhances interactions between users and applications and drives revenues. For more information visit <http://www.sandvine.com> or follow Sandvine on Twitter @Sandvine.



USA
5800 Granite Parkway
Suite 170
Plano, TX 75024
USA

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

CANADA
410 Albert Street,
Suite 201, 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 ©2022 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.