

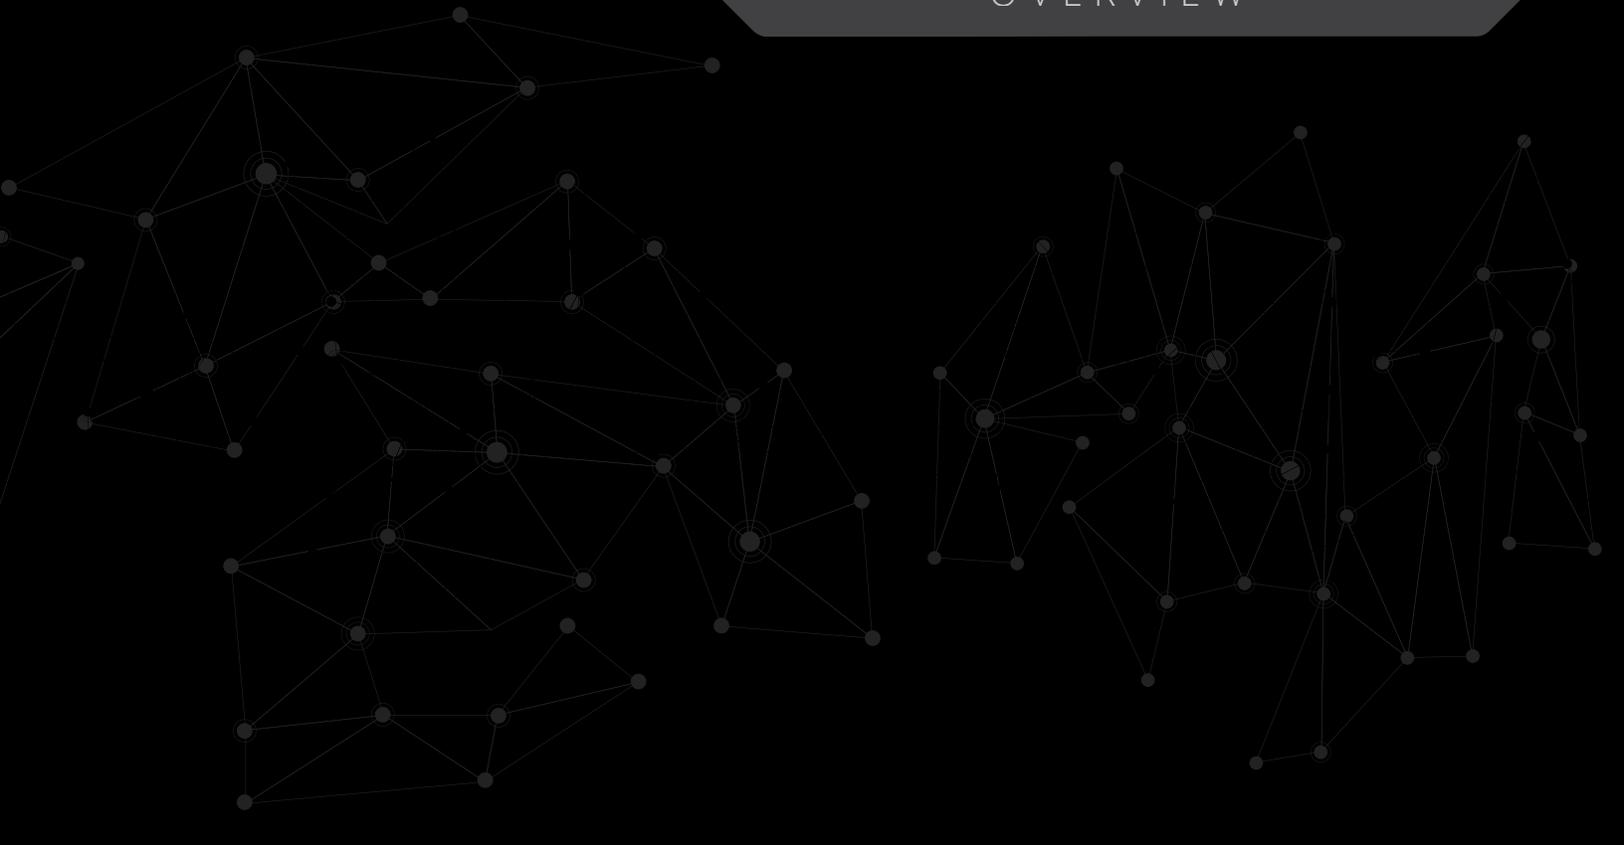


sandvine[®]

Intelligent Broadband Networks

Subscriber Policy Broker

OVERVIEW



Subscriber Policy Broker: Overview

The Subscriber Policy Broker (SPB) is a requisite component of Sandvine's policy control platform deployment, and fulfills two key functions:

Persistent Data Store

Provides long-term storage for the Sandvine deployment. To maximize performance and efficiency, storage is split between two separate databases:

- A subscriber database that stores subscriber profile information (e.g., subscriber attributes, session mapping information, location, etc.)
- A statistics database that enables long-term historical reporting

Subscriber Provisioning

The SPB enables subscriber-aware policy control by making subscriber information, session details, and related subscriber attributes available to the Sandvine Policy Traffic Switch (PTS) and the Sandvine Service Delivery Engine (SDE) in real time, both via notifications and a query interface, through a high-performance, in-memory database.

Different Versions for Different Needs

The SPB is available in two product versions, the SPB Data Manager and the SPB Insight Engine.

SPB Data Manager

The SPB Data Manager is well-suited for relatively smaller policy control deployments that focus on straightforward reporting use cases using simple reporting structures (e.g., basic per subscriber reporting based on element IDs).

SPB Insight Engine

The SPB Insight Engine is designed for relatively larger deployments, and for deployments in which the SPB is used in conjunction with big data systems or in which it fulfills the role of a data lake.

The Insight Engine uses a columnar database designed for analytical and exploratory use cases that depend on storing and retrieving granular information with segmentation and complex data structures (e.g., reporting per-subscriber, per-location, per-protocol).

The Insight Engine scales elastically, acting as one virtual cluster, and includes a number of features specifically to enable natively big data functions and interoperability.

Both SPB versions are completely virtualized, and run on Red Hat Enterprise Linux (RHEL) or Community Enterprise Operating System (CentOS).

Please speak with a Sandvine representative to determine which version of the SPB is best for your needs.



High Performance, High Availability

Performance and availability top the list of demands when it comes to databases, and the SPB delivers both through a combination of four closely related technology features.

| | |
|-------------------|--|
| Clustering | <p>Clustering contributes to optimizing SPB performance and availability by combining many SPBs together to scale performance and storage.</p> <p>For both the Data Manager and the Insight Engine, clustering is an important element of high availability and fault tolerance, and allows operators to segment network locations into individual data homes.</p> <p>In the Insight Engine, clustering enhances performance by enabling a practically limitless number of nodes to operate as a single giant instance, allowing for RAID-like striping and massive parallel processing for fast reporting.</p> <p>To simplify deployment and management, clusters include automatic load balancing and automatic discovery of new elements.</p> |
| Roles | <p>SPB roles segment SPB functions (e.g., statistics storage, subscriber session management, subscriber profile management) to enable independent scaling, redundancy, and performance optimizations based on network needs.</p> |
| Hierarchies | <p>Hierarchies allow operators to segment network reporting from local regions all the way up to comprehensive site-wide visibility. In this manner, operators can optimize different parts of an SPB cluster based upon specific needs and scale (e.g., smaller databases for smaller regions, reporting up into a larger database for the entire network).</p> |
| High Availability | <p>A range of technologies ensures the SPB is a reliable, always-on data store; these features ensure fault management, in general, and minimized downtime during upgrade procedures.</p> |

Other SPB Features

Beyond those features described above, the SPB includes a number of additional features designed to simplify deployment, to ensure smooth operation, and to enable meaningful network insight.

The table below lists some of these features, and indicates to what SPB roles the features apply.

| Feature | Description | Subscriber Session and Profile Manager | Data Manager | Insight Engine |
|-----------------------|---|--|--------------|----------------|
| Top Talkers | Automatically identifies the network's top N subscribers, in terms of bandwidth usage. | | ✓ | ✓ |
| Web Services API | Simplifies integration with operational support systems and other external systems via a SOAP or HTTP REST API to set and/or retrieve data from the system. | ✓ | ✓ | ✓ |
| Subscriber Attributes | Attributes are type-value pairs such as <code>servicelevel="gold"</code> , <code>tethered="true"</code> , <code>browser="chrome"</code> , and <code>masqueradeflag="true"</code> ; attributes are used as conditions in SandScript policy and to enable deep business intelligence reporting. | ✓ | | |

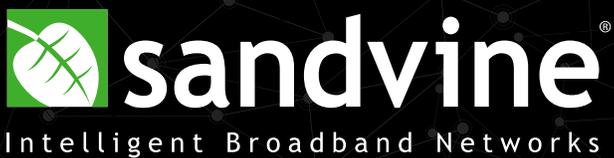
| Feature | Description | Subscriber Session and Profile Manager | Data Manager | Insight Engine |
|--------------------------------|--|--|--------------|----------------|
| Subscriber Attribute Archiver | <p>Enables historical tracking of subscriber-to-IP address mapping and of subscriber attribute values over time.</p> <p>Once the attributes are archived, they can be used in queries that join against subscriber statistics, further enhancing reporting and/or insights into subscriber trends.</p> | ✓ | | |
| Inactive Subscriber Management | <p>Optimizes performance by automatically managing subscribers who have not been provisioned for a configurable number of days. Perform inactive subscriber management in these ways:</p> <ul style="list-style-type: none"> Automatically unassign all the IP assignments from a subscriber after the configured threshold of the number of days of inactivity (default 30 days) is reached. Automatically change the status of a subscriber to inactive after the configured threshold of the number of days of inactivity (default 180 days) is reached. Automatically delete a subscriber after the configured threshold of the number of days of inactivity (default 240 days) is reached. | ✓ | | |
| Attribute Indexing | <p>Adds indexing capabilities to one or more subscriber or session attribute definitions, enabling efficient retrieval of the collection of subscribers with particular attribute definition values.</p> <p>This feature allows for use of non-identifying keys to retrieve subscriber information by outside systems.</p> | ✓ | ✓ | ✓ |
| Element Classifiers | <p>Element classifiers represent unique measurements across one or more dimensions that can be written to the SPB by Sandvine elements, enabling powerful and unique reporting and insights.</p> <p>This capability allows operators to create their own custom report dimensions within SandScript to extend well beyond the data logged by default.</p> | | ✓ | ✓ |
| Auto-Discovery | <p>Allows data from new elements to be added automatically into the database schema so that no additional configuration (beyond pointing the element to the database) is required.</p> | ✓ | ✓ | ✓ |

| Feature | Description | Subscriber Session and Profile Manager | Data Manager | Insight Engine |
|--------------------------------|--|--|--------------|----------------|
| Zero Administration | Provides self-diagnostics and automated clean-up of the SPB database without the need for additional tools (e.g., no need for re-indexing or refreshing of data). Additionally, the installation self-tunes based upon available resources and defined roles. | ✓ | ✓ | ✓ |
| Advanced Compression | 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. | | | ✓ |
| Columnar Storage and Execution | Analytical column store intelligently organizes data on the disk, delivering extremely high data query performance by reducing the amount of data needed to be retrieved for advanced queries and lowering the I/O load on disks. | | | ✓ |
| Distributed Query Execution | Takes advantage of massive parallel processing for faster query execution and reporting by leveraging all nodes within the SPB cluster. | | | ✓ |
| Top N Aggregation | Joins and aggregates encoded data for efficient execution over billions of records, no matter how many nodes are within a cluster or group of clusters. | | | ✓ |
| Elastic Data Store | Automatically rebalances data when nodes are added, removed, replaced, repaired, etc. on-the-fly, and clones databases from one cluster to another. | | | ✓ |
| Zero-Downtime Database | Database continues to support query requests while rebalancing is in progress (e.g., during upgrades or when taking down a node for maintenance). | | | ✓ |



This document is published by Sandvine Incorporated ULC©. All rights reserved. No part of this publication may be reproduced, copied or transmitted in any form or by any means, or stored in a retrieval system of any nature, without the prior permission of Sandvine Incorporated ULC. Although the greatest care has been taken in the preparation and compilation of this whitepaper, no liability or responsibility of any kind, including responsibility for negligence is accepted by Sandvine Incorporated ULC.

© Copyright: Sandvine Incorporated ULC 2017



408 Albert St. Waterloo, Ontario Canada N2L 3V3

www.sandvine.com