ALIBABA CLOUD

Alibaba Cloud

AnalyticDB for PostgreSQL
Release Notes

Document Version: 20220216

Alibaba Cloud

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	**Danger:** Resetting will result in the loss of user configuration data.
Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	**Warning:** Restarting will cause business interruption. About 10 minutes are required to restart an instance.
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	**Notice:** If the weight is set to 0, the server no longer receives new requests.
Note	A note indicates supplemental instructions, best practices, tips, and other content.	**Note:** You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings**> **Network**> **Set network type**.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click **OK**.
`Courier font`	Courier font is used for commands	Run the `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`
{} or {a|b}	This format is used for a required value, where only one item can be selected.	`switch {active|stand}`

Table of Contents

1.AnalyticDB for PostgreSQL V6.0 available for commercial use

AnalyticDB for PostgreSQL V6.0 has been available for commercial use since its public review was completed on September 30, 2019. When you purchase an AnalyticDB for PostgreSQL V6.0 instance or renew the subscription of an existing AnalyticDB for PostgreSQL V6.0 instance, you must pay the required fees before you can use it. The prices for V6.0 and V4.3 of the same instance type are the same. We recommend that you renew the subscription of your instance before it expires to ensure service availability.

For more information, visit AnalyticDB for PostgreSQL

2.Release notes

This topic describes the release notes for AnalyticDB for PostgreSQL and provides links to the relevant references.

You can update your instances to the latest minor version in the AnalyticDB for PostgreSQL console. For more information about how to update the minor version, see Upgrade the engine version.

January 2022

Console

Feature	Description	References
Monitoring and alerting	The monitoring and alerting feature is upgraded to visually display the running status of instances, including their performance and load metrics. This feature is upgraded in the following aspects: • Instance-level monitoring is provided on the **Performance & Load Metrics** page to display the running status of instances. • Node-level monitoring is provided on the **Node Monitoring** page to display metrics of the coordinator node and compute nodes to help you identify the nodes that consume great amounts of resources and take appropriate measures. In addition, compute nodes can be sorted based on their resource usage. • The overall instance status is updated in the **Instance Status** section of the **Basic Information** page, including health status, connection status, and storage status.	API operations: • DescribeDBClusterPerformance • DescribeHealthStatus Metric details: • Performance parameters

Minor version: V6.3.7.0

Release date: January 7, 2022

Category	Feature	Description	References
New feature	pg_hint_plan plug-in	The pg_hint_plan plug-in is supported. This plug-in uses hints to intervene in and optimize execution plans, and registers the optimized SQL patterns and hints. Then, optimized execution plans are automatically generated for SQL statements that use the same SQL patterns to improve execution efficiency.	Updating
	Query cache	The query cache feature is supported to cache query results. This feature allows repeated queries to be executed within milliseconds	Updating

Category	Feature	Description	References
Optimized feature	LLVM JIT	The LLVM JIT feature provided by PostgreSQL 14 is supported.	None
	Backup and restoration	Parallel archiving is supported to improve backup performance.	Back up data
Fixed issue	The issue that causes ineffective ORDER BY clauses in subqueries when the primary query contains only a LIMIT clause is fixed.		
	The issue that causes rebalance failures during data recovery across compute nodes due to a missing mirror host in the pg_hba.conf file of the primary compute node is fixed.		
	The issue that causes partitioned tables to fail to be created on the primary coordinator node in a multi-coordinator architecture is fixed.		

December 2021

Console

Feature	Description	References
Serverless mode	AnalyticDB for PostgreSQL in Serverless mode is released. The Serverless mode provides features such as separation of computing and storage resources, elastic scaling within seconds, and real-time data share across instances. These features are implemented based on the resource pooling and massive storage capabilities of cloud infrastructures as well as the massively parallel processing (MPP), online and offline data processing, and Serverless technologies.	• Serverless mode • Create an AnalyticDB for PostgreSQL instance

New zones for Basic Edition	The following new zones are available for AnalyticDB for PostgreSQL Basic Edition instances:
	• China
	  ◦ China (Hangzhou): Zone I
	  ◦ China (Shenzhen): Zone E
	• Asia Pacific
	  ◦ Singapore (Singapore): Zone A
	  ◦ Indonesia (Jakarta): Zone A
		Basic Edition

Minor version: V6.3.6.1

Release date: December 6, 2021

Category	Feature	Description	References

Category	Feature	Description	References
New feature	Materialized view	The query rewrite feature is supported for materialized views to improve performance for JOIN operations, aggregate functions, subqueries, common table expressions (CTEs), and high-concurrency SQL statements.	• Query rewrite for materialized views (Beta) • Use real-time materialized views to accelerate queries that contain variable parameters
		Real-time materialized views are supported for partitioned tables.	Real-time materialized views
	Partitioned table	The INSERT ON CONFLICT statement is supported for partitioned tables.	Use INSERT ON CONFLICT to overwrite data
		The COPY ON CONFLICT statement is supported for partitioned tables.	Use COPY ON CONFLICT to overwrite data
	Scheduled task	The following features are provided for scheduled tasks after the pg_cron plug-in is updated to V1.4: • Scheduled tasks can be executed across databases. • Existing scheduled tasks can be modified. • The names of scheduled tasks can be specified. • The cron.job_run_details table is added to record the execution information of scheduled tasks.	Use the pg_cron plug-in to configure scheduled tasks
	Data readable during compute node specification changes	In versions earlier than V6.3.6.1, the data of each table is redistributed when the compute node specifications are being changed, and tables are not readable or writable. In V6.3.6.1 and later, all tables are readable and only the tables that incur data redistribution are not writable during this process. This improvement reduces the impact of compute node specification changes on your business.	None
Fixed issue	The issue that causes `Unexpected gang size error` messages in specific scenarios is fixed.		
	The issue that causes frequent `FATAL: reader could not find writer proc entry` messages in specific scenarios is fixed.		

Fixed issue Category	Feature	Description	References
	The issue that causes the Laser engine to incur left join errors in specific scenarios is fixed.		
	The issue that causes data restoration failures for Basic Edition instances due to an invalid synchronous_standby_names parameter is fixed.		

November 2021

Console

Feature	Description	References
Instance data restoration across zones and editions	Instance data can be restored across zones and editions. This way, you can change the zones and editions of your AnalyticDB for PostgreSQL instances to suit your business growth and leverage the new zones and services released by Alibaba Cloud.	• [Notice] Instance data restoration across zones and editions supported • Restore data
Compute node scale-down	Compute nodes can be scaled down to reduce costs during off-peak hours.	• [Notice] Compute node scale-down supported • Change the specifications of compute nodes

October 2021

Minor version: V6.3.5.2

Release date: October 13, 2021

Category	Feature	Description	References
New feature	Auto-merge	Auto-merge is a process that runs in the backend of AnalyticDB for PostgreSQL to automatically sort data. This feature checks the status of data in a table on a regular basis. It sorts new unordered data and merges the data with the existing ordered data. Auto-merge supports only append-optimized (AO) tables that are configured with sort keys.	Auto-merge

Category	Feature	Description	References
	Automatic closing of idle connections	Connections can be closed 6 hours after they become idle. Note	None
Optimized feature	Sorting acceleration	Sorting acceleration is supported when data is updated.	Configure sorting acceleration
	Laser engine	• The performance of the Laser engine on coordinator nodes is improved. • The performance of the Laser engine in scenarios where ORDER BY clauses are executed on a large amount of data is improved. • The performance of the Laser engine in scenarios where the COUNT DISTINCT operator is used is improved. • The issue that may trigger pointer exceptions in specific hash join scenarios is fixed. • The issue that may cause errors in specific hash left join, hash right join, or hash full join scenarios is fixed. • The issue that bitmap index scans for an append-optimized column-oriented table take a long period of time is fixed. • The issue that may cause executions to fail when a large number of partitioned tables are involved is fixed.	Use the Laser computing engine
	Maximum number of user connections	The maximum number of user connections is increased based on instance specifications. • 2 cores, 16 GB: from 300 to 550. • 4 cores, 32 GB: from 300 to 750. • 8 cores, 64 GB: from 400 to 950.	Limits

Category	Feature	Description	References
	Auto-vacuum and auto-analyze	The auto-vacuum and auto-analyze features are supported for multi-coordinator AnalyticDB for PostgreSQL instances.	• Configure scheduled maintenance tasks to clear junk data • Use the ANALYZE statement to collect statistics on AnalyticDB for PostgreSQL • Multi-coordinator architecture
	Real-time materialized view	• Real-time materialized views are supported for INSERT ON CONFLICT DO UPDATE and COPY ON CONFLICT DO UPDATE statements. • The issue that real-time materialized views fail to refresh in a multi-coordinator architecture is fixed. • The issue that real-time materialized views cannot support replicated tables is fixed. • The issue that real-time materialized views cannot support UPSERT clauses is fixed.	Real-time materialized views
	PL/Java	By default, the PL/Java feature is disabled.	Create and use PL/Java UDFs
Fixed issue	The issue that the column name cannot contain `totalrows` or `totaldeadrows` is fixed.		
	The issue that the data type cannot be converted from INT96 to TIMESTAMP in Parquet-formatted Object Storage Service (OSS) foreign tables is fixed.		
	The issue that may cause archiving exceptions is fixed. Archiving exceptions may occur because archiving programs are missing or Python modules fail to be imported after a primary/secondary switchover.		
	The issue that may cause memory leaks if data is accessed or a source file is changed while OSS foreign tables are being scanned is fixed.		

September 2021

Console

Feature	Description	References
Diagnostics and optimization	The diagnostics and optimization feature is supported to display slow queries and SQL distribution statistics in the console.	• [Notice] Diagnostics and optimization feature released • View slow SQL queries • View SQL distribution statistics
Basic Edition	AnalyticDB for PostgreSQL Basic Edition instances are supported to reduce costs for data storage and entry-level instances. Data storage costs are reduced for small and micro enterprises and individual users. Entry-level instance costs are reduced for large and medium-sized enterprises.	• [Notice] AnalyticDB for PostgreSQL Basic Edition Released • Basic Edition • Create an AnalyticDB for PostgreSQL instance

July 2021

Minor version: V6.3.4.0

Release date: July 13, 2021

Category	Feature	Description	References
New feature	Laser engine	By default, the Laser engine is enabled in AnalyticDB for PostgreSQL.	Use the Laser computing engine
	Parallel query	By default, parallel query is enabled for individual tables in AnalyticDB for PostgreSQL.	Configure parallel query
	Sorting acceleration	The sorting acceleration feature is supported. After you execute a `SORT <tablename>` statement, the system sorts the data of the specified table. Then, AnalyticDB for PostgreSQL pushes operators such as SORT down to the storage layer so that queries are accelerated based on the physical order of data. This feature can accelerate queries that contain SORT, AGG, and JOIN operators based on sort keys.	Configure sorting acceleration
	Backup and restoration	The backup and restoration feature is supported for AnalyticDB for PostgreSQL instances that have multiple coordinator nodes.	• Overview • Back up data • Restore data

Category	Feature	Description	References
	Minor version query	Global User Configuration (GUC) parameters are supported to query the minor version of AnalyticDB for PostgreSQL. The following statement is used: `show adbpg_version;`	Query the minor engine version
Optimized feature	Real-time materialized view	Performance is enhanced for real-time materialized views. The aggregate functions MAX and MIN can be used in DELETE and UPDATE statements. CTEs can be used in change statements executed on base tables.	Real-time materialized views
	Distributed transaction for instances that have multiple coordinator nodes	When coordinator nodes execute distributed transactions, each coordinator node directly generates distributed transaction IDs instead of requesting these IDs from the global transaction manager (GTM). This eliminates single-point bottlenecks of the GTM and improves distributed transaction performance in the event of highly concurrent queries.	None
Fixed issue	The issue that may cause deadlock detection failures when more than two distributed transactions are initiated in the same session of a secondary coordinator node is fixed.		
	In previous versions, Panic errors may occur in the checkpoint process on the secondary server after you truncate an AO table multiple times in a transaction and then abort the transaction. This issue is fixed.		

May 28, 2021

Category	Feature	Description	References

Category	Feature	Description	References
New feature	Auto-vacuum	The auto-vacuum feature is supported to automatically execute VACUUM statements. Auto-vacuum checks for tables that have a large number of INSERT, UPDATE, or DELETE operations. When necessary, auto-vacuum executes a VACUUM statement on the tables to vacuum junk data to accelerate queries. By default, when more than half of the rows of a table are deleted, auto-vacuum executes a VACUUM statement on the table to vacuum junk data. If an AnalyticDB for PostgreSQL instance has multiple coordinator nodes, auto-vacuum checks for change operations only on the primary coordinator node. Auto-vacuum cannot be triggered if change operations are performed on secondary coordinator nodes.	Configure scheduled maintenance tasks to clear junk data
	Auto-analyze	The auto-analyze feature is supported to automatically execute ANALYZE statements. Auto-analyze checks for tables that have a large number of INSERT, UPDATE, or DELETE operations. When necessary, auto-analyze executes an ANALYZE statement on the tables to collect their statistics after the operations. By default, when these operations are performed on more than 10% of the rows in a table, auto-analyze executes an ANALYZE statement on the table. If an AnalyticDB for PostgreSQL instance has multiple coordinator nodes, auto-analyze checks for change operations only on the primary coordinator node. Auto-analyze cannot be triggered if change operations are performed on secondary coordinator nodes.	Use the ANALYZE statement to collect statistics on AnalyticDB for PostgreSQL
	BRIN index	BRIN stands for Block Range Index and is supported for compressed AO tables. For example, BRIN indexes can be used to query large-range data from an ordered table or filter out data blocks that are not needed to reduce I/O. When large datasets are involved, BRIN indexes can provide similar or superior performance but require less physical storage and generation costs compared with other indexes such as B-tree indexes.	Manage indexes

Category	Feature	Description	References
	Parallel query (beta)	The rds_segment_expansion_coeff parameter is supported for session-level queries. The parameter value is of the INTEGER type, and the default value is 1. This parameter specifies a multiple of cores for a single query to implement parallel query and uses a query optimizer to further optimize the query. In scenarios with low queries per second (QPS), this parameter can linearly improve the performance of compute-intensive queries executed for longer than 3 seconds. Typical scenarios include aggregate queries for individual tables (TPC-H Q1 or TPC-H Q6) and join operations between large and small tables. However, this parameter provides only marginal improvements for I/O-intensive queries and high disk usage scenarios, and may reduce the performance for network-intensive queries.	None
	Real-time materialized view (beta)	Real-time materialized views are supported for AnalyticDB for PostgreSQL. Real-time materialized views provide features similar to those of materialized views, but do not require you to manually execute REFRESH statements when data changes.	Real-time materialized views
Optimized feature	Instance endpoints can be used to perform DDL operations between the BEGIN transaction block and the ROLLBACK or COMMIT transaction block.		
	Instance endpoints can be used to perform TEMP TABLE and TEMP VIEW operations for session-level queries.		

May 19, 2021

Category	Feature	Description	References
Optimized feature	I/O performance	The I/O performance of new instances is improved by 30% to 50%.	I/O performance optimization on compute nodes

February 8, 2021

Category	Feature	Description	References
	Multi-coordinator node	The number of connections and the I/O performance linearly increase with the number of coordinator nodes based on the scalability of compute nodes.	Endpoints of an instance and its primary coordinator node

Category	Feature	Description	References
New feature	PITR	Point-in-time recovery (PITR) is supported.	None
	Multi-replica	Two or more compute nodes are supported to synchronize data.	None
	Upgrade or downgrade across specifications	Clusters can be changed from the reserved mode to the elastic mode.	None
	Compatible with time formats in AnalyticDB for MySQL V2.0	Time formats in AnalyticDB for MySQL V2.0 are supported. Example: 2020-08-03T23:59:59.	None
	Indexes creation or deletion for JSON fields by executing ALTER TABLE statements	Indexes for JSON fields can be disabled by executing the ALTER TABLE statement.	ALTER TABLE
	BINARY type	The BINARY type is supported for the metadata of the protocol layer.	None
	Export of file headers during export from AnalyticDB for PostgreSQL to a single OSS object	File headers can be exported when you export data from AnalyticDB for PostgreSQL to a single OSS object by using an external table.	Use an external table to import data from OSS at a high speed
	Maximum number of rows in an object that can be generated when you export data from AnalyticDB for PostgreSQL to OSS by using an external table	If the number of exported rows exceeds the maximum number, extra rows are exported to one or more new objects. You can specify both the maximum size and maximum number of rows in an object. Written data that first triggers the limit is exported to a new object.	None
	SQL Plan module	Execution plans of slow SQL queries can be viewed in the AnalyticDB for PostgreSQL console.	None
	INSERT INTO SELECT ON DUPLICATE KEY UPDATE	This query is supported when the input values in the UPDATE column are constants or when the input values in the UPDATE column are those in the SELECT column.	None
	File format of OSS foreign tables	The ORC format is supported for OSS foreign tables.	None

Category	Feature	Description	References
	Priority of BATCH LOAD statements	The priority of BATCH LOAD statements can be specified by using hints.	None
Optimized feature	Performance of the LIMIT n statement	Performance is improved when you use the pushdown logic of the LIMIT n statement to filter data.	None
	Compatibility	The BOOL and BOOLEAN types are supported for table creation statements.	None
	Database naming conventions	Database names can start with an uppercase letter or an underscore (_).	None
Fixed issue	The issue that may cause index exceptions when indexes are added or removed at the same time is fixed.		
	The issue that may cause `java.lang.NullPointerException` errors when data is inserted to a table that does not exist is fixed.		
	The issue that may cause data query failures if partition fields exist in Parquet-formatted files when OSS partition tables are pushed down is fixed.		
	The issue that may cause `Premature end of Content-Length` errors when you access OSS is fixed.		
	The issue that may cause `com.alibaba.fastsql.sql.ast.expr.SQLMethodInvokeExpr cannot be cast to com.alibaba.fastsql.sql.ast.SQLName` errors when ADD KEY statements are executed is fixed.		
	The issue that may cause `no composite key to add` errors due to composite keys when tables are managed in hot or cold storage is fixed.		

December 28, 2020

Category	Feature	Description	References
New feature	Ganos	The spatial database engine Ganos can be integrated into AnalyticDB for PostgreSQL.	None
	Auto-analyze	The auto-analyze feature is supported. If the number of inserted or deleted rows exceeds a specified threshold, the system performs ANALYZE operations to collect and update statistics and avoid performance regression.	None
	Data export from partitioned external tables	Data can be exported from partitioned external tables.	None
	ZORDER BY clause	The ZORDER BY clause is supported.	None

October 12, 2020

Category	Feature	Description	References
New feature	JIT compilation feature of PostgreSQL 11	The just-in-time (JIT) compilation feature of PostgreSQL 11 is supported for the row-oriented Laser computing engine of AnalyticDB for PostgreSQL.	None
	FastANN	Update-in-place indexing is supported for FastANN.	None
		A new storage engine is supported for FastANN vector indexing.	None
	OSS foreign table	The COPY and UNLOAD statements supported by Amazon Redshift can be used to import data to or export data from OSS foreign tables.	None
		OSS foreign tables can be used to export data from on-premises tables to OSS.	None

August 25, 2020

Category	Feature	Description	References
New feature	Rough set filtering	The rough set filtering and sort key features are supported for column-oriented tables.	None
	MaxCompute foreign table	MaxCompute foreign tables can be used to map from MaxCompute strings to AnalyticDB for PostgreSQL data types.	None

July 3, 2020

Category	Feature	Description	References
New feature	Engine version and ORCA version update	The engine version is updated to 6.3.0. The ORCA version is updated to 3.86.0.	None
	OSS and MaxCompute foreign tables	OSS foreign tables can be used to read standard Snappy-compressed files in the CSV or TEXT format, Parquet files, and JSON files. Filter conditions can be pushed down to decrease the volume of data that is pulled from OSS.	None
		MaxCompute foreign tables are supported to access MaxCompute data.	None

May 6, 2020

Category	Feature	Description	References
New feature	OSS foreign table created based on the FDW feature	OSS foreign tables created based on the foreign data wrapper (FDW) feature can be used to scan ORC files.	None
	Real-time logging	The real-time transfer process of gptransfer can be written to logs.	None
	orafce extension	The orafce extension is supported to implement Oracle syntax and functions.	None
	Field naming conventions	Uppercase field names are supported.	None
	More features of gptransfer	More features are supported for gptransfer. Partitions can be exchanged between tables. If the size of a row exceeds the upper limit during transmission, the upper limit can be automatically modified for retry. Tables that have the SERIAL field are supported. Double quotation marks (") can be added to table names. External tables can be partitioned.	None
Fixed issue	The issue that may cause errors for the plancache direct dispatch feature when the ORCA query optimizer is enabled is fixed.		
	The issue that may cause failure to inherit tables by using gptransfer is fixed.		

April 17, 2020

Category	Feature	Description	References
New feature	intarray extension	The intarray extension is supported.	None
	gptransfer for individual databases or tables	gptransfer can be used to store separate index data in individual databases or tables.	None
	Integrated query of vectors and structured data	Integrated queries of vectors and structured data are supported.	None
	Temporary authorization token	Temporary authorization tokens can be created to pull data from OSS by using CREATE LIBRARY statements.	None
	Append-optimized column-oriented table	The Odyssey computing engine is supported for append-optimized column-oriented tables.	None
	HyperLogLog extension	The HyperLogLog extension is supported.	None

Category	Feature	Description	References
Fixed issue	The issue that may cause failure to migrate temporary tables and extension data tables by using gptransfer is fixed.		
	The issue that may cause errors for VACUUM FULL operations on the pg_class table is fixed.		

March 27, 2020

Category	Feature	Description	References
New feature	Lineage analysis	The lineage analysis feature is supported for Data Management (DMS).	None
	Fast decimal	The fast decimal feature is supported.	None
	FDW on OSS	The FDW feature is supported for OSS.	None
	pg_trgm extension	The pg_trgm extension is supported.	None
Fixed issue	The issue caused by invalidity of the OSS account is fixed.		
	The version incompatibility issue of the MADlib dynamic library is fixed.		
	The issue that may increase memory usage of coordinator nodes when a large number of transactions are executed by Data Transmission Service (DTS) is fixed.		

February 26, 2020

Category	Feature	Description	References
New feature	Laser vector computing engine	The Laser vector computing engine is supported. The session variable `enable_laser` can be set to true or off to specify whether to accelerate queries by using the Laser computing engine.	None
	PostGIS extension	The PostGIS extension is supported to analyze spatial data.	None
	Accelerated analysis of OSS foreign tables	Accelerated analysis of OSS foreign tables is implemented based on the PostgreSQL FDW framework.	None
	FastANN extension	The FastANN extension is supported for vector queries.	None
Fixed issue	The method used to import SDK code for OSS external tables is modified. This way, the OSS SDK versions can be managed in a centralized manner.		
	The issue that may cause core segmentation faults due to CANCEL operations when data is imported by using OSS external tables is fixed.		

Category	Feature	Description	References
	The issue that may cause failure to dump rds_superuser by using pg_dumpall is fixed.		

December 30, 2019

Feature	Description	References
Commercial use of AnalyticDB for PostgreSQL V6.0	AnalyticDB for PostgreSQL V6.0 has been available for commercial use since its public review was complete on September 30, 2019. When you purchase an AnalyticDB for PostgreSQL V6.0 instance or renew the subscription of an existing AnalyticDB for PostgreSQL V6.0 instance, you must pay the required fees before you can use it. The prices for V6.0 and V4.3 of the same instance type are the same. We recommend that you renew the subscription of your instance before it expires to ensure service availability.	• Features of AnalyticDB for PostgreSQL V6.0 • AnalyticDB for PostgreSQL pricing

December 26, 2019

Category	Feature	Description	References
New feature	tablefunc extension	The tablefunc extension is supported.	None
Fixed issue	The issue that may cause MADlib to be unavailable on some types of servers is fixed.		
	The issue that may cause failures when extensions are created is fixed.		
	The issue that may cause failures when temporary tables or indexes are created is fixed.		

December 21, 2019

Category	Feature	Description	References
New feature	postgres_fdw extension	The postgres_fdw extension is supported.	None
Fixed issue	The default maximum number of concurrent queries that are allowed in a resource queue is increased from 20 to 500.		

November 25, 2019

Category	Feature	Description	References
New feature	RoaringBitmap extension	The RoaringBitmap extension is supported.	None
	Instance restart is no longer required when the global deadlock detection mechanism of an instance is enabled or disabled.		

Category	Feature	Description	References
Fixed issue	The memory computing logic that is used in the out of memory (OOM) protection mechanism is revised to prevent OOM errors in the event of highly concurrent queries.		
	The superuser account is authorized to grant permissions to and revoke permissions from standard accounts.		

October 17, 2019

Category	Feature	Description	References
Fixed issue	The issue that may cause failures when instances are created is fixed.		

October 15, 2019

Category	Feature	Description	References
New feature	ORCA version	The ORCA version is updated to 3.77.0.	None
	Global deadlock detection mechanism	By default, the global deadlock detection mechanism is enabled.	None
Fixed issue	Online transaction processing (OLTP) is optimized. Point query, insert, update, and delete operations based on the partition key can be committed by using phase 1 to reduce network broadcast overheads.		

September 12, 2019

Category	Feature	Description	References
New feature	ORCA version	The ORCA version is updated to 3.67.0.	None
Fixed issue	By default, the query optimizer that comes with PostgreSQL is used for simple queries from individual tables.		
	The performance of instance scale-up and scale-out is improved.		

December 24, 2019 (V4.3)

New feature	Feature	Description	References
New feature	Auto-vacuum, auto-vacuum freeze, and auto-analyze	The auto-vacuum, auto-vacuum freeze, and auto-analyze features are supported. If the number of inserted or deleted rows exceeds a specified threshold, the system performs VACUUM or ANALYZE operations to avoid performance regression.	None

New feature	Feature	Description	References
Fixed issue	The issue that may cause failures when statements such as SET are executed during instance scaling is fixed.		
	The issue that may cause failures in distributed queries when compute nodes run different versions is fixed.		

December 18, 2019 (V4.3)

Category	Description
Fixed issue	The issue that may cause failures when read-only transactions are explicitly enabled during instance scaling is fixed.
	The issue that may cause an instance to break down due to deadlocks produced from a full resource queue is fixed.

November 13, 2019 (V4.3)

Category	Feature	Description	References
New feature	adbpg.sql script	The adbpg.sql script is supported to define functions to batch grant or revoke permissions. **Note** If you want to execute the script, submit a ticket.	None
Optimized feature	Metascan	The metadata of Metascan is optimized to include the block offset. The block offset reduces the disk I/O by enabling AnalyticDB for PostgreSQL to directly read blocks that meet the specified filter conditions.	None
Fixed issue	The rds_superuser account is authorized to grant permissions to and revoke permissions from non-superuser accounts.		

October 16, 2019 (V4.3)

Category	Description
Fixed issue	The issue that may cause the global deadlock detection mechanism to break down when the number of XIDs on compute nodes reaches the xid_warn_limit value is fixed.

October 15, 2019 (V4.3)

Category	Feature	Description	References
New feature	Global deadlock detection mechanism	By default, the global deadlock detection mechanism is enabled.	None
Optimized feature		The global deadlock detection mechanism is optimized.	None
	Instance scale-up and scale-out	The data transmission of tables that each have only a single index during instance scale-up and scale-out is optimized.	None
Fixed issue	The issue that may cause the HyperLogLog extension to occupy more memory resources than it has applied for is fixed.		

September 10, 2019 (V4.3)

Category	Feature	Description	References
New feature	Global deadlock detection mechanism	The global deadlock detection mechanism is supported. When you execute UPDATE or DELETE statements on a table, the table is no longer locked.	None
	uuid-ossp extension	The ORDER BY clause is supported for the uuid-ossp extension.	None
Optimized feature	gp_max_slices parameter	The default value of the gp_max_slices parameter is increased from 50 to 500.	None
	Instance scale-up and scale-out	The speed of instance scale-up and scale-out is increased.	None

September 2, 2019 (V4.3)

Category	Feature	Description	References
New feature	Instance migration	Data can be migrated from a V4.3 instance to a V6.0 instance.	None
	uuid-ossp extension	The uuid-ossp extension is supported.	None
	gp_max_slices parameter	The gp_max_slices parameter is supported to limit the maximum number of slices allowed in a query. This parameter helps prevent engine instability that may arise due to a large number of running processes.	None
Fixed issue	The issue that may cause failures when data is queried from a column-oriented table after an index is created for the LOB column is fixed.		
	The issue that may cause the gpexpand extension to abnormally exit during instance scale-up and scale-out is fixed.		