



Release Notes

Product	VoltDB
Version	8.4.16
Release Date	November 28, 2022

This document provides information about known issues and limitations to the current release of VoltDB. If you encounter any problems not listed below, please be sure to report them to support@voldb.com. Thank you.

Upgrading From Older Versions

The process for upgrading from the recent versions of VoltDB is as follows:

1. Shutdown the database, creating a final snapshot (using **voltadmin shutdown --save**).
2. Upgrade the VoltDB software.
3. Restart the database (using **voltdb start**).

For DR clusters, see the section on "Upgrading VoltDB Software" in the *VoltDB Administrator's Guide* for special considerations related to DR upgrades. If you are upgrading from versions before V6.8, see the section on "Upgrading Older Versions of VoltDB Manually" in the same manual.

For customers upgrading from V7.x or earlier releases of VoltDB, please see the *V7.0 Upgrade Notes*.

For customers upgrading from V6.x or earlier releases of VoltDB, please see the *V6.0 Upgrade Notes*.

For customers upgrading from V5.x or earlier releases of VoltDB, please see the *V5.0 Upgrade Notes*.

For customers upgrading from V4.x or earlier releases of VoltDB, please see the *V4.0 Upgrade Notes*.

Changes Since the Last Release

Users of previous versions of VoltDB should take note of the following changes that might impact their existing applications.

1. Release V8.4.16 (November 28, 2022)

1.1. Log4J replaced by reload4J

VoltDB does not use any of the components implicated in the published CVEs related to Log4J. However, to avoid any confusion, VoltDB has replaced the Log4J library with reload4J, a drop-in replacement that replicates the log4J namespace and functionality, but eliminates all known security vulnerabilities.

1.2. Recent improvements

The following limitations in previous versions have been resolved.

- There was an issue where, if a stored procedure queued more than 200 SQL statements before calling `voltExecuteSQL()` and at least one of the statements was a `SELECT` statement that returned data, the result buffer could become corrupted causing one or more nodes to crash. This issue has been resolved.
- There was a problem where, if a properties file in the database root was corrupted, the database would issue a fatal error with no explanation. The error now identifies the corrupted file and the names of the missing properties.
- In the situation where a cluster failed or was forcibly shutdown while a node was being added or removed, attempting to restart the cluster could result in an error claiming there were "incomplete command logs", even if command logging was not enabled. This was caused by an incomplete snapshot left by the interrupted cluster expansion. The issue has been resolved.
- Additional information is now logged if the SQL compiler encounters an unexpected error while processing a data definition language (DDL) statement.

2. Release V8.4.15 (November 12, 2021)

2.1. IMPORTANT: Limit partition row feature has been removed from VoltDB V11.0

The `LIMIT PARTITION ROWS` feature was deprecated in VoltDB V9 and subsequently removed in V11. This is a change to the VoltDB schema syntax that is not forward compatible.

This means that if your database schema still contains the `LIMIT PARTITION ROWS` syntax, you need to remove the offending clause before upgrading to the latest major release. Fortunately, there is a simple process for doing this. You can use the `ALTER TABLE {table-name} DROP LIMIT PARTITION ROWS` statement to correct the table schema while the database is running and with no impact to the database contents.

2.2. Recent improvements

The following limitations in previous versions have been resolved.

- There was a rare condition where the VoltDB network process could report an index out of bounds error, causing the cluster to hang. This condition is now caught. As a consequence of the error, one of the nodes will stop, but the cluster as a whole will continue and not be deadlocked.
- There was an issue where using the `CAST` function to convert a `VARCHAR` column to a `BIGINT` could generate incorrect values if the number in question had more than 18 digits. This issue has been resolved.

3. Release V8.4.14 (June 7, 2021)

3.1. New license improvements

This release includes a number of improvements to the licensing and management of VoltDB software. These improvements include:

- Support for a new, improved license format
- The `voltadmin show license` command for displaying information about the current license for a running VoltDB cluster
- A new selector, `LICENSE`, for the `@SystemInformation` system procedure to provide similar information programmatically
- A new `voltadmin inspect` command used by VoltDB product support to display summary information about the cluster operating environment

3.2. Recent improvement

The following limitation in previous versions has been resolved.

- Previously, if the snapshot rate limit was set (using the Java property `SNAPSHOT_RATELIMIT_MEGABYTES`), requesting a CSV formatted snapshot could raise an illegal argument exception stating that "requested permits must be positive" and the resulting snapshot files would be empty. This only affected CSV formatted snapshots. This problem has been resolved.

4. Release V8.4.13 (March 9, 2021)

4.1. Ability to cancel a shutdown snapshot

The **voltadmin shutdown --save** command pauses the database, flushes buffers for services for export and database replication, and saves a final snapshot of the database contents. However, if remote systems are not available to acknowledge export or DR buffers, the shutdown may hang and you need to CTRL-C to exit from the command. However, the shutdown process is still pending and the database not in normal operating mode. The **voltadmin shutdown --cancel** command has been added that lets you cancel a pending shutdown and resume normal operations.

4.2. Recent improvements

The following limitations in previous versions have been resolved.

- There was a potential situation where, if a cluster used for cross datacenter replication (XDCR) suffered one or more node failures, then was shutdown and restarted using command logs to recover, replication might later fail with a "replica ahead of master" error. This underlying issue was related to recovery using the failed node's command logs which did not match the current state of the remote cluster. This problem has been resolved.
- A similar problem could occur if a node becomes detached from the cluster (for example, due to network issues) and does not immediately fail but times out. The result was that the remote cluster might stop replication, also reporting the "replica ahead of master" error. This issue has also been resolved.
- There was a rare condition where workloads containing lots of complex multi-partition read transactions interspersed with single-partition writes could result in a deadlock, stalling the database and generating log messages reporting a "possible multipartition transaction deadlock". This issue has been resolved.
- The **snapshotconverter** utility lets you generate CSV files from VoltDB snapshot files. These files can be used to recover and reload data from individual tables through the **csvloader** utility. However, for certain data — such as XDCR tables, tables defined with `MIGRATE`, or views with no `COUNT(*)` column — the **snapshotconverter** utility includes hidden columns in its output, which can be confusing. A new command flag has been added, `--filter-hidden`, that lets you exclude these hidden columns from the utility's output.

5. Release V8.4.12 (January 22, 2021)

5.1. Improved performance for DR DROP command

The **voltadmin DR DROP** command removes a cluster from a cross datacenter replication conversation. Previously, this procedure could take up to 90 seconds. The command has been optimized to significantly reduce the time it takes to complete.

5.2. Recent improvements

The following limitations in previous versions have been resolved.

- There was an issue where a stream could stop writing data to its export target after having more than two billion rows inserted into any one partition. The problem surfaced only after the necessary number of records (approximately 2.15 billion) were written to the export connector and the database was saved, shutdown, restarted, and restored. After the snapshot was restored, no further records were written to the target by the export connector.

This issue has now been resolved. In fact, upgrading VoltDB using the standard **voltadmin shutdown --save** command, installing this release, and then restarting the database will automatically circumvent the issue.

- Due to issues in the underlying library used, it was possible for the JSON functions to return results in a different order on different servers, causing a hash mismatch error. This inconsistency, and the resulting issue, have now be resolved.
- The VoltDB bulk loader (available in the client API and used in the loader utilities such as csvloader) has been optimized to remove an unnecessary regular expression evaluation of string columns. This change produces a noticeable improvement in load times for large data sets.
- The **snapshotconvert** utility has been corrected to interpret null values as end-of-file, rather than reporting an error. At the same time, general error handling has been enhanced and extended to report more detailed information when a failure occurs.
- There was a problem with the Kinesis importer where the importer could fail with a "no class found" error. This issue has been resolved.
- There was a rare situation where if a schema failed causing a deadlock, subsequent attempts to rejoin nodes to the cluster would fail. This issue has been resolved.
- Previously, the bulkloader interface, which is used by VoltDB data utilities such as csvloader and is available through the Java API, did not correctly account for the additional data structures required by cross datacenter replication (XDCR) or TTL with migrate. As a result, attempting to bulk load data into an XDCR cluster or a table with MIGRATE TO TARGET and USING TTL could cause the cluster to crash. This issue has been resolved.
- There was an issue with earlier releases of the JDBC export connector. If the connector property `createable` was set to true but the target database (for example, Oracle) did *not* support the CREATE TABLE... IF NOT EXISTS clause, the export connector would repeatedly attempt to create the table and fail, although the table did exist, resulting in innumerable spurious error messages. This issue has been resolved.
- In earlier releases, if an attempt was made to restore a snapshot to an empty database, but the specified unique ID did not exist or was misspelled, subsequent attempts to restore the snapshot would claim to succeed but would not restore any data. This issue has been resolved.
- Previously, if two clusters established cross datacenter replication (XDCR), including a table defined with USING TTL and MIGRATE TO TARGET on one cluster but not on the other, the cluster could crash with a fatal error while processing the XDCR binary logs. This issue has been resolved.
- There was a rare condition where the digest of the synchronization snapshot for database replication (DR or XDCR) could exceed the 50MB limit for network packets. When this happened, nodes on the consumer cluster would fail. The size of the digest files triggering this condition have been reduced to eliminate the issue.

6. Release V8.4.11 (June 25, 2020)

6.1. Recent improvements

The following limitations in previous versions have been resolved.

- There was an issue where a schema or configuration change on a running database could stall if there was an active export queue connected to a JDBC target. Closing the export connection took too long and could potentially cause a deadlock with the requested change. This issue has been resolved.
- There was a rare case where, if a replica cluster in passive DR is promoted and then a node fails, other nodes could subsequently fail when the cluster tries to reassign partition leadership. The symptom for this particular failure scenario was that the failing nodes reported an error processing an invocation dispatcher request. This issue has now been resolved.

7. Release V8.4.10 (June 15, 2020)

7.1. Java 11 support

VoltDB now supports the use of Java 11 as well as Java 8 for production use of the VoltDB server and clients.

7.2. Improved handling of SSL/TLS connections in the JDBC interface

The handling of secure export connections (using SSL/TLS) through the JDBC interface has been improved. Specifically, the requirement for a truststore when using a commercial certificate has been removed.

7.3. Change to Kafka export default behavior

The Kafka `acks` property determines whether VoltDB waits for acknowledgement of receipt from the Kafka brokers. Previously, the default was set to "1" (one), but the recommendation was to set it to "all" to protect against the loss of records if the Kafka brokers fail. The default has changed to match the recommended setting.

Existing customers who use Kafka export but do not explicitly set the `acks` property may notice a slight change in export latency. The new default of "all" is the recommended setting. However, if you are willing to accept less durability on the part of the Kafka brokers, you can explicitly set the property back to "1" to replicate previous behavior.

7.4. Security Notice

The following change has been made to improve security and eliminate potential threats:

- Previously, you could not enable SSL/TLS and Kerberos authentication at the same time. This limitation has been removed.

7.5. Recent improvements

The following limitations in previous versions have been resolved.

Restore:

- Using VoltDB V6.5, taking a snapshot of a cluster with database replication (DR) enabled could create a snapshot that cannot be restored in later versions of VoltDB (V7, V8 or V9). The restore operation would fail with an error. This issue has been resolved and the current version of VoltDB can now restore the problematic snapshots.
- There was an issue where restoring a snapshot, either manually or automatically on startup, could fail if the directory contained multiple snapshots and the unique IDs for the snapshots were similar. (That is, one ID matched the starting characters of another ID, such as *test* and *testme*.) When this happened, the restore would fail with an error stating that a table had an "inconsistent transaction ID." This issue has been resolved. VoltDB now performs an exact match on the selected snapshot's unique identifier.

- Previously, attempting to restore a snapshot created on a standalone cluster to a cluster configured for cross datacenter replication (XDCC) would fail with a misleading error message (indicating that the configuration could not be updated). This issue has been resolved, the snapshot is restored, and no error is reported.
- There was an obscure issue where if a database cluster was restored from snapshots multiple times, the fourth restore command could stall and eventually timeout. For this unusual situation to occur the database schema must contain views, must be K safe, and nodes must have stopped and rejoined between each snapshot restore operation. This issue has been resolved.
- When manually restoring a snapshot, you must specify a path to the snapshot files. Previously, if the path did not exist on a node, the restore would fail, even if sufficient files for a full backup existed on the other nodes. The restore now ignores the missing path.
- There was a problem where using the VoltDB Management Center to restore a snapshot would not work. Using the **Restore** button on the web interface would actually start two restore operations, which would then cause constraint violations. This issue has been resolved.

Adding/Rejoining Nodes:

- When rejoining a node to a running cluster, the system clock on the rejoining node must be within the limits for clock skew on the cluster, just like when starting the cluster for the first time. If not, the rejoin operation will fail. Previously, there was an issue where if a rejoin failed due to clock skew, subsequent attempts to rejoin nodes would fail even if the clock skew had been corrected. This issue has been resolved.
- VoltDB 8.2 introduced an issue that could cause a cluster to hang or crash when attempting to add nodes using elastic scaling. Under certain conditions, where a database has views, it was possible for the cluster to hang or to crash reporting a deserialization error while attempting to pass copies of the current database contents to the joining nodes. This issue has been resolved.
- Previously, while adding nodes to the cluster on the fly, if a table exceeded its LIMIT PARTITION ROWS setting, the elastic expansion of the cluster would fail along with one of the existing nodes. This issue has been resolved.
- There was an issue where, if a database had no persistent tables (but may have had streams defined with or without export), a node rejoining the cluster could end up with inconsistent internal settings related to export and/or database replication. This issue has been resolved.
- Under certain rare error conditions, an attempt to elastically add nodes to a cluster could fail, crashing the cluster and reporting an illegal argument exception. This issue has been resolved.
- There was an issue with elastic expansion of a cluster. If, while adding nodes "on the fly", ongoing transactions within one of the partitions being rebalanced generates a constraint violation, it could result in the rebalance operation reporting a fatal "failed to delete tuple" error, causing the server process to exit. This issue has now been resolved.
- In previous releases, attempting to rejoin a node to a cluster could fail with the error "Bad message length exception" if the database schema contained too many tables (that is, thousands of table declarations, regardless of actual data volume). This issue has now been resolved.

Monitoring:

- The New Relic latency graph data has been adjusted to improve accuracy.

- The New Relic node count report was accurate while the cluster was running, but did not "zero out" for periods while the cluster was down. This report now reports zero running nodes for those intervals when the cluster was stopped.

Export:

- There was an issue where changing the Kafka export configuration on a running database could cause the update to hang. The issue was triggered by a Kafka export configuration specifying an invalid Kafka broker. Attempting to update the configuration to change the invalid broker specification would cause the update process to hang. This issue has been resolved.

SQL:

- VoltDB V8.4.9 introduced an obscure but critical bug where a SQL query with a complex ORDER BY clause and two separate column expressions each using the DECODE() function could produce incorrect results. This issue has been resolved.

8. Release V8.4.9 (April 28, 2020)

8.1. Recent improvements

The following limitations in previous versions have been resolved:

- Previously, if a GROUP BY query included an arithmetic expression combining an aggregate function and a parameter cast to a specific datatype with the CAST() function, the query failed to compile. For example:

```
SELECT a, CAST(? AS INTEGER) + COUNT(b)
FROM TABLE c GROUP BY a;
```

This issue has been resolved for the CAST() function. However, combining other parameterized functions with an aggregate may also cause a compilation error. The workaround, until the more general case is fixed, is to put the parameterized function in a subquery.

- Previously, there was a rare case where an ad hoc query using parameters in a complex combination of CAST() and SUM() operations would give different results than the same query using constant values. For example, the following fragment with the arguments "A" and 1 gave different results than expected if the placeholders were replaced by the SQL constants 'A' and 1:

```
SUM( DECODE(column1, CAST(? AS VARCHAR), CAST(? AS INTEGER), NULL))
```

This issue has been resolved.

- There was an issue where SNMP traps would fail if the server had been running for around 24 days. After 24 days, whenever an SNMP trap was triggered, VoltDB would report that the trap failed with an illegal argument exception error. This issue has been resolved.
- All SNMP traps are now recorded in the log file. Previously, SNMP traps were sent, but not recorded. Log entries provide additional confirmation that SNMP events are triggered.
- There was an edge case where the database could stop processing snapshots if there were too many tables with too much data. The issue involved a race condition and did not surface all the time, so the database might run for an extended period of time before an error was triggered. But the more tables or the larger volume of data in the database, the more likely a snapshot might fail. And once it failed, no subsequent snapshots could be taken. This issue has been resolved.

- As part of the overall hardening of the LTS release, a number of rare circumstances causing deadlocks were identified, tracked down, and resolved. These corner cases included a deadlock during a node rejoin, when a node failed while processing a read-only multi-partition procedure, and as a result of nodes failing while the cluster is trying to re-establish a quorum from an earlier node failure. All these issues have now been resolved.

9. Release V8.4.8 (January 16, 2020)

9.1. New Relic enhancements

The current release improves the content and structure of VoltDB performance and management data provided to the New Relic monitoring suite.

10. Release V8.4.7 (December 16, 2019)

10.1. Security Notice

The following libraries used by VoltDB have been updated to ensure the latest security and performance patches are applied:

- Commons-compress 1.19
- Dom4j 2.1.1
- Jackson 2.9.10
- Jetty 3.27
- Kafka client 0.10.2.2
- Netty 4.1.43
- Openssl 1.0.2t
- Scala 2.11.12
- Tomcat 7.0.96

10.2. Recent improvements

The following limitations in previous versions have been resolved:

- If a node rejoining a cluster was stopped with @StopNode but later successfully rejoined the cluster, there was a chance the other nodes in the cluster did not recognize the successful rejoin and would continue reporting the error "REJOIN: No stream snapshot ack message was received in the past 10 minutes or the thread was interrupted." This issue is now resolved.
- Under certain rare conditions, if a node failed on a K-safe cluster while a node was being added, a snapshot was in progress, or the @SwapTables system procedure was executing, the command logs for that cluster could be rendered incomplete. If, after this event, the cluster stopped and restarted before the next command log snapshot was taken, the command logs could not be replayed beyond the point of the node failure. This issue has now been resolved.
- There was an issue with file export when using TSV (tab-separated value) format and exporting data containing quotations marks or backslashes. The export connector incorrectly attempted to "escape" the output, although TSV format does not support quoting or escaping. The result was incorrect output in the export file. This issue has been resolved. The export connector no longer attempts to escape special characters in the output.
- Export data optionally contains six columns of metadata, including a timestamp identifying when the row was exported. Previously, this timestamp was mistakenly set 12 years prior to the actual date. This issue has been resolved.

- This release contains a substantial number of fixes and improvements to the handling of export data. In particular, it addresses several edge cases associated with the durability of export queues during unusual system failure scenarios. Taken together these changes significantly improve the reliability, durability, and recoverability of export in VoltDB.

11. Release V8.4.6 (September 11, 2019)

11.1. Recent improvements

The following limitations in previous versions have been resolved:

- There was issue where, in rare cases, if a schema update or configuration change failed, subsequent attempts to update the schema or configuration would also fail. This could only happen if the original update failed with an unhandled exception (such as a reference to a missing Java method), at which point subsequent update attempts reported that another update was still in progress. This issue has been resolved.
- There was an edge case where, if a multi-partition transaction encounters a deadlock during a rejoin, the rejoin — and any subsequent rejoins — would fail repeatedly reporting a possible deadlock. This issue has been resolved.

12. Release V8.4.5 (July 3, 2019)

12.1. Recent improvements

The following limitation in previous versions has been resolved:

- There was a rare edge case where a query could return incomplete results. If an index on a table included two columns and the WHERE clause of a query on the table included both an IN clause applied to one column and a less than or equal to (\leq) evaluation of the second column, fewer than expected rows were selected. This issue has been resolved.

13. Release V8.4.4 (June 13, 2019)

13.1. Improved handling of multi-partition transactions with large intermediate result sets

VoltDB limits each transaction to 50 MB of results. In fact, each transaction fragment is limited to 50MB. However, for multi-partition transactions, this means each partition can return up to 50MB of data to the coordinator. For large clusters with many unique partitions, it is possible for all this data to exceed the allocated Java heap for the coordinator, causing that node to fail with an out of memory error.

To avoid this situation, VoltDB now limits the amount of data each fragment can return, based on the current maximum heap size and number of partitions. By default, each partition in a multi-partition transaction is only allowed to return the lesser of 65% of the maximum heap size divided by the number of unique partitions ($\text{sitesperhost} * \text{number of nodes} / k+1$) or 50MB. The limit is further reduced for read-only multi-partition transactions to accommodate for the fact that multiple read-only transactions can be run simultaneously. If, at runtime, a fragment exceeds the limit, it throws an exception and the transaction rolls back gracefully.

You can adjust this per-partition multi-partition response limit by setting the environment variable `MP_MAX_TOTAL_RESP_SIZE`. You can either set it as the percentage of max heap to use in the calculation (by using the percent sign) or as a specific number of bytes (by using an integer value with no suffix). For example, to allow a maximum of only 50% of the allowable heap size, you can set the following environment variable before starting the server:

```
$ export MP_MAX_TOTAL_RESP_SIZE="50%"
$ voltdb start -D ~/mydb ...
```

You can set `MP_MAX_TOTAL_RESP_SIZE` as an environment variable or as a Java system property through `VOLTDB_OPTS`:

```
$ export VOLTDB_OPTS="-DMP_MAX_TOTAL_RESP_SIZE=50%"
$ voltdb start -D ~/mydb ...
```

13.2. Additional improvements

The following limitations in previous versions have been resolved:

- There was an issue with elastic expansion of a cluster. If, while adding nodes "on the fly", ongoing transactions within one of the partitions being rebalanced generates a constraint violation, it could result in the rebalance operation reporting a fatal "failed to delete tuple" error, causing the server process to exit. This issue has now been resolved.
- There was a edge case where a K-safe cluster could run out of memory, caused by too many long-running and/or failing multi-partition transactions. In the case of K-safe clusters, the multi-part coordinator must maintain a log of its in-flight transactions. If too many transactions linger or fail and rollback, the log itself can grow out of control, ultimately exceeding available memory before it can be pruned and causing the server process to fail. This rare condition is now avoided by increasing the priority of log management, ensuring it is pruned in a timely fashion.
- There was an obscure issue where if a database cluster was restored from snapshots multiple times, the fourth restore command could stall and eventually timeout. For this unusual situation to occur the database schema must contain views, must be K safe, and nodes must have stopped and rejoined between each snapshot restore operation. This issue has been resolved.
- There was an issue in the VoltDB Enterprise Edition where, if you attempted to add a node on the fly (known as elastic scaling) but did not have a license for database replication (DR), the operation would cause the cluster to crash and interfere with restarting the cluster from command logs. This issue has been resolved.
- There was an issue where the export subsystem was repeatedly reporting a warning that it received an export message with a signature "which does not exist on this node." Although this message was innocuous and did not indicate any real problem, it was misleading and could fill up the logs for no reason. This issue has now been resolved.
- There was an edge case where attempting to update the schema while a snapshot is being saved could hang the database. This issue has been resolved. Attempts to change the schema or configuration during a snapshot are no longer allowed and must wait until the snapshot is complete.

14. Release V8.4.3 (March 24, 2019)

14.1. Support for multiple schema and classes files when initializing the database root

The `voltdb init` command now allows you to specify multiple files as arguments to the `--schema` and `--classes` flags. Separate multiple files with commas. You can also use the asterisk (*) as a wildcard character. For example, the following command initializes a root directory with two schema files, plus all the JAR files from one folder and another JAR file from the current working directory:

```
$ voltdb init -D ~/mydb \
  --schema=customers.sql,companies.sql \
  --classes=storedprocs/*.jar,myfunctions.jar
```

It is also possible to specify multiple schema and classes files when configuring VoltDB for use in Kubernetes. See in the readme file in the `tools/kubernetes/` subfolder where VoltDB is installed for details.

14.2. Additional improvements

The following limitations in previous versions have been resolved:

- There was an issue introduced in VoltDB 8.2 when the USING TTL clause was added to allow you to automatically delete old records from tables based on the specified column value. Accidentally the USING TTL clause was also allowed on CREATE STREAM statements, although it has no application to streams. This issue has now been resolved.
- Under certain rare conditions, the @Quiesce system procedure could return control to the calling program before all export and DR data is successfully processed. If this occurs it was possible for an orderly shutdown (that is, **voltadmin shutdown** without the **--force** argument) to stop the cluster before all pending DR or export data was made durable. This rare race condition has now been resolved.

15. Release V8.4.2 (February 14, 2019)

15.1. Improvements to the New Relic plugin

The VoltDB plugin for the New Relic monitoring application has been updated and enhanced. See the README in the `tools/monitoring/newrelic` folder where the VoltDB Enterprise Edition software is installed for details.

15.2. Fixes to the Kubernetes support

In VoltDB 8.4 and 8.4.1, a component of the Kubernetes support was missing from the kit. As a consequence, the readiness probe did not operate as expected. This issue has been resolved.

16. Release V8.4.1 (January 29, 2019)

16.1. Improvements to Kubernetes support

Support for running VoltDB in Kubernetes, and in particular resilience during cluster startup and XDCR durability, has been improved in several ways. These changes improve the reliability of XDCR in Kubernetes environments and especially environments using glusterfs. However, these changes are equally applicable to other environments where node stability is questionable during startup. Specific improvements include:

- Errors and stack traces appearing during startup have been eliminated.
- Files written to the database root directory during startup are now protected against corruption due to nodes failing during the startup process.

16.2. Additional improvements

The following limitations in previous versions have been resolved:

- There was an issue in database replication (DR) where dropping a partitioned table could cause the consumer cluster to hang. This could only happen if the cluster was not paused and DR drained before the schema change and there was a simultaneous transaction being processed involving the table and other non-partitioned tables. This issue has now been resolved.
- ASSUMEUNIQUE indexes are designed specifically for partitioned tables, where a partitioned procedure cannot verify the uniqueness of an index on columns other than the partitioning column. They do not apply to replicated tables. However, there was an issue where defining an ASSUMEUNIQUE index on a replicated table could result in a database failure. This issue has been resolved.

As part of this fix, use of CREATE ASSUMEUNIQUE INDEX on a replicated table will automatically be converted to CREATE UNIQUE INDEX. So to create an ASSUMEUNIQUE index on a partitioned table

you must partition the table before issuing the `CREATE ASSUMEUNIQUE INDEX` statement or define the index as a constraint in the `CREATE TABLE` statement itself.

- Similarly, declaring an index that should be `ASSUMEUNIQUE` (because it is on a non-partitioning column of a partitioned table) as `UNIQUE` could cause problems in previous releases. VoltDB correctly warned that the declaration of the `UNIQUE` index was invalid. But a later transaction could result in an "invalid use of `UNIQUE`" error, stopping the database. This issue has been resolved.
- The `LOG10()` function can result in invalid values, depending on the argument to the function. Normally, using such functions to define an index on a table that already contains data is not allowed. However, `LOG10()` was inadvertently allowed in this situation. This issue has been resolved and use of `LOG10()` is no longer allowed when creating indexes on non-empty tables.
- Recent changes to the export connector programming interface incorrectly changed the behavior of the `onBlockStart()` method. For custom export clients, `onBlockStart()` should be invoked only once per block. However, starting in VoltDB 8.1, multiple calls to `onBlockStart()` could be initiated before `onBlockCompletion()`. This change in behavior has been corrected.
- There is a Java environment variable that is intended to set the maximum time interval between when data is inserted into a stream and when it is flushed to the export connector. However, in previous releases, setting that variable, `MAX_EXPORT_BUFFER_FLUSH_INTERVAL`, did not have any affect. This issue has now been resolved. You can now set the export flush interval in milliseconds, with 1000 being the current lowest practical value.

17. Release V8.4 (December 29, 2018)

17.1. Export improvements

A number of improvements have been made to export in VoltDB. In particular three areas have been addressed:

- **Errors when draining export** — a number of issues related to the last few records in an export buffer not being able to drain have been resolved.
- **Export Statistics** — The `@Statistics` system procedure has a new selector, `EXPORT`, that provides detailed information about the progress of export for each stream and partition.
- **Gaps in export queues** — Exporting data from VoltDB is an asynchronous process, so that export does not impact the performance of the database itself if there are delays communicating with the external target. If the remote system cannot keep up, VoltDB queues the overflow data to disk. This both buffers the pending content and provides durability in case of database failure.

However, export data is only durable if the database servers restart with the same disk-based content. Similarly, in a K-safe environment, all copies of a partition buffer export data. But when a node stops and rejoins the cluster, there can be a gap in its export buffer for the time it is down. If the export target is not accepting data fast enough and the cluster nodes stop and rejoin frequently, it is possible for a gap to occur in the currently active export queue.

In that past, it was assumed the data was lost and export proceeded to the next available record. The database now handles export gaps more intelligently. First it queries all copies of the partition to see if any contain the missing records. If the records are not found but nodes are currently missing from the cluster, export waits for the nodes to rejoin to see if they have the missing export records. If so, export continues where it left off. If not and the cluster is complete, export resumes, logging the fact that certain export records may be lost. If, for any reason, export is paused at a gap in the queue and you wish to continue (for example, you cannot rejoin a failed node but want export to continue) you can use the `voltadmin export release` command to have export resume at the next available record.

17.2. New interactive command line utility, `voltsql`

There is a new interactive command line utility, **voltsql**, for accessing VoltDB databases that performs the same functions as **sqlcmd** but adds command completion to simplify entry of SQL commands. As you type, **voltsql** lists possible keywords, validating the syntax and reducing the typing required. The new utility is a preview release and does require a few additional Python libraries. See the description of `voltsql` in the *Using VoltDB* manual for details.

17.3. Support for the `ROW_NUMBER()` window function

The new `ROW_NUMBER()` window function provides the ordinal position of the current row in the group defined by the `PARTITION BY` clause. See the description of window functions for the `SELECT` statement in the *Using VoltDB* manual for details.

17.4. Protection against `TRUNCATE TABLE` in single-partitioned procedures

`TRUNCATE TABLE` is intended to efficiently delete all tuples in the specified database table. However, previously it was possible to execute a `TRUNCATE TABLE` statement in a single-partitioned procedure, resulting in only the tuples from that partition being deleted. To avoid this indeterminate behavior, `TRUNCATE TABLE` is no longer permitted in a single-partitioned stored procedure.

17.5. `COUNT(*)` no longer required for single table views

You no longer need to include `COUNT(*)` as an explicit column in a view of a single table. For example, if you want a view of the maximum population, by state, of the table `City`, the `CREATE VIEW` statement might look like this:

```
CREATE VIEW maxbystate (state, population)
  AS SELECT state, MAX(population) FROM city
  GROUP BY state;
```

The `COUNT(*)` column is still required for views that join multiple tables. See the description of `CREATE VIEW` in the *Using VoltDB* manual for details.

17.6. More control of TTL processing

The `USING TTL` clause to the `CREATE TABLE` statement now supports two additional arguments: `BATCH_SIZE` and `MAX_FREQUENCY`. These arguments let you adjust how many records are deleted in each cycle and how frequently the process checks for outdated records. These values can be set per table. See the documentation of `CREATE TABLE` in the *Using VoltDB* manual for details.

17.7. Improved `voltadmin stop` command

The **voltadmin stop** command has been enhanced to perform an orderly stoppage of the specified node. Previously, the stop command simply stopped the server process on the node. Now, the command migrates all partition and export leadership before stopping the node. This avoids a number of problems:

- Less disruption to ongoing transactions
- No lost connection errors returned to client applications
- Better management of export buffers, avoiding gaps in the pending queues

You can still force the command to stop the node immediately if necessary by adding the `--force` argument. See the description of `voltadmin` in the *Using VoltDB* manual for details.

17.8. New quick reference information available

The VoltDB documentation now includes an online quick reference guide that lists the syntax for SQL statements, DDL, SQL functions, command line utilities and system procedures. Each item contains a link to its full documentation. Give it a try and tell us what you think.

17.9. Deprecation of Java 7 support

Support for the use of Java 7 in client applications is deprecated. The recommended version of Java for VoltDB servers and clients is version 8.

17.10. Support for Ubuntu 18.04 added

VoltDB now supports Ubuntu 18.04 as a server operating system. Officially supported platforms for VoltDB servers currently are CentOS 6.6, CentOS 7.0, RHEL 6.6, RHEL 7.0, and Ubuntu 14.04, 16.04 and 18.04. See the section on in "Operating System and Software Requirements" the *Using VoltDB* manual for details.

17.11. VoltDB Community Edition available in Maven

The VoltDB Community Edition Jar file is now available for distribution through the Maven central repository.

17.12. Security Notice

The following change has been made to improve security and eliminate potential threats:

- Enabling SSL encryption on VoltDB interfaces now takes advantage of OpenSSL to significantly improve performance if it is installed on the server. For the improvement to take effect, both the VoltDB server and Java client library must be upgraded to the latest version. The new client library also has a dependency on the Netty libraries `io.netty:netty-all:4.1.32.Final` and `io.netty:netty-tcnative-boringssl-static:2.0.20.Final`. If OpenSSL is not installed or the latest (8.3.4 or later) VoltDB client library used, the server falls back to using the builtin Java SSL client.

17.13. Additional improvements

In addition to the new features and capabilities described above, the following limitations in previous versions have been resolved:

- There was an issue with database replication where replication could stop and report that the replica was "ahead" of the master. This issue has been resolved.
- Previously, VoltDB did not allow you to create an index using a unary minus. That is, a column expression preceded by a minus sign. For example "CREATE INDEX myindex ON mytable (product, -price)". This issue has been resolved.
- In a related issue, VoltDB previously allowed DESC (descending) in the index definition, although it had no effect on the index created. The DESC keyword is no longer allowed.
- There was an issue in the Java client API where an idle client application (that is, an application that creates a client instance but makes no calls to VoltDB for an extended period of time) could eventually use up all available CPU. This issue has been resolved.
- In database replication (DR), if a transaction executed a TRUNCATE TABLE statement followed by inserts, the inserted data was present on the producer cluster but was not properly replicated to the consumer. This issue has been resolved.

18. Release V8.3.3 (November 21, 2018)

18.1. Improved Kubernetes support

The scripts that support running VoltDB in a Kubernetes environment have been rewritten to use ConfigMaps for storing configuration, schema, and stored procedure classes, separating this information from the Docker image. It is now possible to reuse a single Docker image in multiple configurations. See the readme in the `/tools/kubernetes` directory where VoltDB is installed for details.

18.2. Additional improvement

The following limitation in previous versions has been resolved:

- VoltDB introduced two major changes in recent releases: significant improvement in replicated table storage (V8.1) and leadership rebalancing (V8.3). One outcome of these changes is the discovery of edge cases where multi-partition transactions stall in what is referred to as a "deadlock". These deadlocks are rare — often only seen in lab tests. Rare as they are, any such case of a processing failure is critical and this release fixes two known instances of multi-part deadlocks.

19. Release V8.3.2 (November 8, 2018)

19.1. New command line option for advertising an alternate DR interface

A new command line option for the `voltldb start` command, `--drpublic`, lets you specify an alternate interface and (optionally) port, which the server then reports to database replication (DR) consumer clusters. This feature is helpful for cloud environments where the internal interfaces are not accessible from outside the hosted region, so the other DR clusters must use redirected interfaces and ports. Specify the public interface as an IP address or host name followed by an optional colon and port number. If you do not specify a port number, the publicly advertised port number is the same as the value for the internal `--replication` port. For example:

```
$ voltldb start --drpublic=some.external.addr
```

19.2. Additional improvements

The following limitations in previous versions have been resolved:

- There was an issue introduced in 8.1 where, when using export and nodes failed and then rejoined the cluster, export could lose track of a few records, causing the export connector to never completely drain its queue. Symptoms of this bug were that the export statistics would never reach zero and attempts to use the `voltadmin` commands `pause --wait` or `shutdown --save` would hang. This issue has been resolved.
- There was an issue where attempting to apply a schema change involving the "time to live" (TTL) feature through the `sqlcmd` utility could result in the command hanging. The root cause was an error applying the schema change. However, the error was not reported to the user. The `sqlcmd` utility now recognizes such situations, reports the error to the user, and returns to the command prompt. At the same time, the underlying issue with the TTL schema change has been corrected.

20. Release V8.3.1 (October 5, 2018)

20.1. Additional improvements

The following limitations in previous versions have been resolved:

- There was an issue introduced in 8.3 that interfered with database replication (DR) of multi-partition transactions. If the last query executed in a read/write transaction was a `SELECT` statement rather than a data manipulation statement (such as `INSERT`, `UPDATE`, or `DELETE`), the transaction might not be correctly replicated to the consumer cluster. This issue has been resolved.
- The syntax of the `ALTER TABLE` for adding a `USING TTL` statement to an existing table was incorrect in 8.3. The correct syntax is `ALTER TABLE table-name ADD USING TTL...` but the "ADD" keyword was missing previously. This issue has been resolved.

- Due to a change in the server logic for the VoltDB Management Center (VMC) introduced in V7.9, the database replication (DR) and Import tabs no longer showed up when those features were turned on in the database. This issue has been resolved.

21. Release V8.3 (September 21, 2018)

21.1. Leadership Rebalancing

In a K-safe cluster, individual nodes are assigned as *leaders* for each unique partition and coordinate executing transactions for that partition on any copies within the cluster. If a node fails, leadership can be reassigned to one of the remaining nodes. If multiple nodes fail, this means leadership for all partitions could end up congregating on only a few nodes. Previously, leadership was not redistributed when nodes rejoined the cluster.

Now the cluster rebalances leadership of the partitions when the cluster returns to its full complement of nodes. That is, as soon as all of the failed nodes complete the rejoin process.

21.2. New STARTS WITH clause optimizes text comparisons

There is a new clause available for SQL statements such as SELECT. The STARTS WITH clause does text comparisons that are equivalent to LIKE with a text string ending in a percent sign (%). That is, it matches string values starting with the specified argument. STARTS WITH is beneficial in compiled statements (such as stored procedures) because the clause STARTS WITH ? can use indexes on the column being evaluated, whereas an equivalent LIKE ? clause cannot. See the description of the STARTS WITH clause in the SELECT statement in the *Using VoltDB* manual for details.

21.3. FORMAT_TIMESTAMP() function

TIMESTAMP values are stored and expressed in Greenwich Mean Time (GMT). The new SQL function FORMAT_TIMESTAMP() lets you convert such values to a formatted text string in time zones other than GMT. See the description of FORMAT_TIMESTAMP() in the *Using VoltDB* manual for details.

21.4. Ability to selectively restore specific tables from a snapshot

The **voltadmin restore** command now supports the `--tables` and `--skiptables` arguments that let you either include or exclude data from specific tables when restoring a snapshot. Note that, for an empty database, all of the tables in the snapshot schema are created. The `--tables` and `--skiptables` arguments only control whether data is restored or not for the specified tables. See the description of the **voltadmin** utility in the *Using VoltDB* manual for details.

21.5. New information in @Statistics DRPRODUCER output

The DRPRODUCER selector for the @Statistics system procedure now has an additional field. The CONNECTION_STATUS column in the first VoltTable results tells you whether the connection to the consumer cluster is active ("UP") or the connection is broken ("DOWN"). See the description of the DRPRODUCER return values in the *Using VoltDB* manual for details.

21.6. Java ByteBuffer now accepted as input to VARBINARY columns

Previously, the VoltDB callProcedure method accepted strings or byte arrays as input to VARBINARY columns. Now you can also use ByteBuffer as an input datatype.

21.7. Security Notice

The following change has been made to improve security and eliminate potential threats:

- The JSON interface no longer accepts HTTP requests that use the TRACE method.

21.8. Additional improvements

In addition to the new features and capabilities described above, the following limitations in previous versions have been resolved:

- There is a condition where, under certain conditions, if the database is idle (that is, no read or write transactions are occurring) snapshots can get into a scheduling loop causing a CPU spike and preventing other threads from running. This occurs only when the database is configured with a large number of sites per host running on systems with slower disks and fewer CPU cores (for example, in virtualized environments). To avoid this condition, a new option, `DISABLE_IMMEDIATE_SNAPSHOT_RESCHEDULING`, has been added.

In normal database operation, this option is not needed. However, if your configuration matches these conditions and your database falls idle for any significant time, you can set this option to true when you start the database to circumvent the problem. You set the option as a Java environment variable on all the servers at start up using the `VOLTDDB_OPTS` environment variable and including the "-D" flag. For example:

```
export VOLTDDB_OPTS="-DDISABLE_IMMEDIATE_SNAPSHOT_RESCHEDULING=true"
```

- Previously, if you initialized a new database root directory with a configuration that enabled cross datacenter replication (XDCR), it was not possible to restore a snapshot after starting the new, empty database. The problem was that XDCR creates streams for logging XDCR conflicts and those streams were seen as an existing schema. This issue has been resolved and XDCR conflict streams are ignored for the purposes of snapshot restore.
- There was an issue introduced in 8.2 that affects database replication (DR) and time to live (TTL). If a replicated DR table is defined with TTL, the TTL delete procedure continuously generates DR binary logs. As a result, any attempt to pause the DR consumer or perform an orderly shutdown fails since the DR buffer never drains. Note that this issue was specific to databases with replicated tables declared as both DR tables and with TTL. This issue has now been resolved.

22. Release V8.2.2 (September 17, 2018)

22.1. Additional improvements

The following limitations in previous versions have been resolved:

- There was an issue in VoltDB 8.2 with the elastic expansion of clusters. After adding nodes to the cluster, it was possible for the rebalancing of the partitions to fail. This could only happen on VoltDB 8.2 or later when the schema includes a view on a partitioned table involved in the rebalancing. The failure could result in memory corruption, potentially leading to the cluster crashing at some future point in time. This issue has been resolved.
- There was an extremely rare case where in a K-safe cluster, if a node failed and the remaining nodes are busy, the intra-cluster failure and repair messages could be processed out of order, causing errors in synchronization. The major difficulty with this rare error condition is that it would not be detected or reported by the cluster at the time. However, it could result in later transactions generating a hash mismatch. The issue was exacerbated with SSL/TLS enabled on the internal ports, since there could be a backlog of messages from the failed host requiring decryption before being delivered. This issue has been resolved.
- When loading data into variable-width columns with the bulkloader methods (either through a utility such as `csvloader` or in a custom Java application), there was a memory leak that occurred whenever the input data exceeded the maximum size of the column. This issue was specific to `VARCHAR` and `VARBINARY`

columns larger than 64 bytes and only affects VoltDB V8.1 or later. In extreme cases, where many such exceptions occurred, the process could potentially run out of memory. This issue has been resolved.

- There was an issue with time to live (TTL) introduced in 8.2, where if too many records scheduled to be deleted had the exact same value for the TTL column, TTL would fail to delete the records and report an unexpected error condition. This issue has been resolved.
- There was an issue introduced in VoltDB 8.1 related to database replication (DR) and TRUNCATE TABLE statements applied to replicated DR tables. Use of TRUNCATE TABLE on a producer cluster to clear a replicated table could result in memory corruption and failure of the consumer cluster. This issue has been resolved.

It should be noted that use of TRUNCATE TABLE on *any* tables in an XDCR environment (as opposed to passive DR) is not recommended, even with this issue resolved. The current implementation cannot guarantee that the two XDCR clusters might not suffer an undetectable conflict since a TRUNCATE TABLE statement does not log the specific rows that are deleted.

23. Release V8.2.1 (August 6, 2018)

23.1. Kubernetes Support

The software kit now includes support for running VoltDB under Kubernetes and Docker. See the read me file in the `tools/kubernetes/` folder where VoltDB is installed, or the documentation (available as HTML and PDF).

23.2. Additional improvements

The following limitations in previous versions have been resolved:

- Normally, if the size of any input value exceeds the size of a variable-width column (such as VARCHAR or VARBINARY), it generates a SQL exception and the statement is rejected. There was an issue introduced in V8.1 where, under certain circumstances, when bulk loading data and the input exceeded the width of a VARCHAR or VARBINARY column, a fatal exception was generated and the database stopped. This did *not* happen during normal INSERT or UPDATE statements; only when using utilities such as csvloader or using the bulkloader methods in the Java client API. It also only occurred in certain operating system-specific environments. This issue has been resolved.
- Due to a change in the server logic for the VoltDB Management Center (VMC) introduced in V7.9, the database replication (DR) and Import tabs no longer showed up when those features were turned on in the database. This issue has been resolved.
- A query can have 1,025 parameters (or placeholders) at most. However, there was an issue where entering a query with too many parameters through the @AdHoc system procedure could result in a runtime error, crashing the database server. This includes queries entered through JDBC prepared statements, which use @AdHoc implicitly. This issue has been resolved.

24. Release V8.2 (July 12, 2018)

24.1. New TTL feature automates deleting old data

A new feature, "time to live" (TTL), allows you to define an expiration timestamp for individual tables. Once the TTL value is exceeded, the records from that table are automatically deleted from the database. This makes the processing of streaming data easier by automating the deletion of old data.

You define the expiration timestamp with the new USING TTL {value} ON COLUMN {column-name} clause in the CREATE TABLE statement. You can also monitor the performance of TTL processing using the new

TTL selector for the @Statistics system procedure. See the documentation of CREATE TABLE and @Statistics in the *Using VoltDB* manual for details.

24.2. Support for reading the username and password from a file for the VoltDB command line utilities

When using scripts to manage a secure database, the command line utilities (such as **sqlcmd** and **voltadmin**) require a username and password. Previously, there was no easy way to do this without either using Kerberos or hardcoding the information into the script itself. Now you can save the username and password into a properties file — accessible only to the user running the script — and then reference that file in the script using the new `--credentials` argument. See the description of the command line utilities in the *Using VoltDB* manual for details.

24.3. New option to create cluster-wide unique file names on file export

The file export connector writes export data to files on each server in a cluster. By default, the files are unique per server, but not necessarily across the cluster as a whole. You can now set the property `uniqueNames` to `true` in the export configuration to ensure that all files are unique cluster wide. See the description of the file export connector in the *Using VoltDB* manual for details.

24.4. Improved performance when restoring snapshots

Restoring the database from a snapshot can take time, particularly for large databases with many views. This release improves the performance of snapshots by storing the contents of certain views as part of the snapshot, eliminating the need to rebuild the views on the fly when restoring the snapshot. Note that not all views can be saved in the snapshot; views containing partitioned tables but no partitioning column in the GROUP BY clause must still be rebuilt. Also, snapshots created on earlier versions of VoltDB will still have their views rebuilt as part of the restoration process. But for new snapshots with views of replicated tables or partitioned tables with at least one partitioning column in the GROUP BY column, restoring the snapshots should be noticeably faster.

24.5. New sqlcmd directive describes the columns in a table

The **sqlcmd** utility now supports the `DESCRIBE table-name` directive. `DESCRIBE` lists the columns of the specified table, stream, or view and related information, such as datatype and size. See the description of **sqlcmd** in the *Using VoltDB* manual for details.

24.6. Security Notice

The following change has been made to improve security and eliminate potential threats:

- Ability to "hide" the username and password for command line utilities in a separate credentials file. (See description above.)
- When enabling SSL with a user-generated certificate, you need to specify both a keystore and truststore. When using a commercial certificate a local truststore should not be needed. However, previously VoltDB still required one. Specifying a truststore is no longer required when using a commercial certificate.

24.7. Additional improvements

In addition to the new features and capabilities described above, the following limitations in previous versions have been resolved:

- There was an issue with database replication (DR) where a large multi-partition transaction could produce binary logs from each partition on the producer that fell under the 50MB limit on inter-cluster communication. But when aggregated on the consumer for replay, the transaction exceeds the limit. The symptom was that nodes on the cluster would report a "bad message length" exception, causing nodes to be expelled from the cluster until the cluster itself failed. The possibility of an excessively large transaction still exists,

but now the producer cluster rejects the transaction and replication and consistency between the clusters is maintained.

- There was an issue using the VoltDB Management Center (VMC) if the VoltDB http port was set to port 80. Port 80 is the default port for web browsers and if the browser did not send a port number VMC would incorrectly assume a default of 8080 and not operate properly. This issue has been resolved.
- VoltDB reserves the maximum negative value of numeric datatypes as null. (For example, -128 for a TINY integer.) So users should not be allowed to use these reserved values in the context of a given datatype. However, previously the compiler silently accepted such constants and interpreted them as null. This issue has been resolved. The compiler now throws an error when evaluating numeric values equal to a datatype's maximum negative value.
- Certain SQL functions (such as NOW and PI) that take no argument can be entered with or without parentheses. However, these functions were not interpreted consistently in the selection list of a SELECT statement. If the parentheses were left off, NOW was interpreted as the function and PI was interpreted as a column reference. These functions are now both interpreted as functions, whether with or without parentheses.

Note: this is a slight change of behavior. If you use a column with the name PI, you will now have to fully qualify the name to have it be interpreted as a column rather than a function. For example, in the following statement, the first item in the section list is interpreted as the function PI and the second as the column PI:

```
SELECT pi, a.pi FROM a WHERE...
```

- There was a rare condition involving database replication (DR), where replication could break if a producer cluster suffered a network partition. If the production cluster split into two segments due to network issues, a race condition could result in the consumer cluster querying the smaller segment of the cluster for topology information after the separation but before the smaller segment was shutdown by VoltDB's network partition detection. If this occurred, the consumer cluster would wait for the smaller segment and fail to poll the larger, surviving segment. This issue has been resolved.

25. Release V8.1.2 (June 14, 2018)

25.1. Recent improvement

The following limitation in V8.1 has been resolved:

- VoltDB 8.1 introduced a performance improvement to the data loading utilities and bulkloader API. Unfortunately, this feature also introduced a potential error condition where, if the loader encounters a runtime error, such as an input value exceeding the maximum width of a VARCHAR column, rather than rolling back the transaction it could crash the database cluster. This issue has now been resolved. Because this bug can cause the database to stop, we strongly recommend all customers using 8.1 or 8.1.1 upgrade to 8.1.2 at their earliest convenience.

26. Release V8.1.1 (June 7, 2018)

26.1. Recent improvement

The following limitation in V8.1 has been resolved:

- VoltDB 8.1 introduced an issue that could interfere with the resilience of a K-safe cluster. If a node failed while processing a multi-partition transaction, it was possible for the remaining nodes in the cluster to suffer a deadlock. When this happened, the warning "possible multipartition transaction deadlock detected" was reported and all subsequent multi-partition transactions would hang, along with certain system operations such as snapshots and command log truncation. This issue has now been resolved. Customers using V8.1 are strongly encouraged to update to V8.1.1 at their earliest convenience.

27. Release V8.1 (May 26, 2018)

27.1. Improved performance of export during schema changes

VoltDB now does a better job of managing the interaction of export and ongoing schema changes. Previously, export associated with the original schema had to drain before export using the new schema could begin. Now export from before and after a schema change are managed independently and in parallel.

27.2. Better memory management for replicated tables

The storage of replicated tables has been reorganized in this release. Previously, each partition retained a copy of the replicated tables. Now, all of the partitions on a server share a single copy of the tables. Applications with sizeable replicated tables or a high sites-per-host count should notice a significant reduction in the amount of memory required by VoltDB after the upgrade.

27.3. Improved bulk loading of replicated tables

The default process for bulk loading replicated tables — either through the loader utilities such as `csvloader` or through the bulkloader API — has been improved. When bulk loading data into a replicated table using the default load procedure and performing inserts (not upserts), the load process can be as much as three times faster, according to testing.

27.4. New optimization for limit/offset query performance

There is a limit (50 megabytes) to the amount of data any query can return. When reading large volumes of data from a VoltDB database, use of the `LIMIT` and `OFFSET` clauses to "page" through the data is recommended. However, as the `OFFSET` value increases significantly, each query can take incrementally longer to execute. This release introduces a new feature where, if there is an index on the appropriate columns of the table, the query is optimized eliminating the penalty associated with large offsets.

27.5. Improved ad hoc query performance

Ad hoc queries that perform read-only operations on replicated (non-partitioned) tables have traditionally been executed by a single partition within the database, assuming such queries are a small percentage of the overall workload. However, when they are *not* a small percentage, that one partition can get over extended, resulting in increased latency. This release changes the execution model to "round robin" read-only queries of replicated tables to more evenly distribute the workload.

27.6. New system procedure @Ping

A new procedure has been added to the list of supported system procedures for VoltDB, `@Ping`. It returns a value of zero (0) if the database is up and operational. The `@Ping` system procedure is a lightweight procedure and does not require any interaction between cluster nodes, which makes it a better choice than other system procedures (such as `@Statistics`) if all you need to do is check if the database is running. See the description of `@Ping` in the *Using VoltDB* manual for details.

27.7. Additional improvements

In addition to the new features and capabilities described above, the following limitations in previous versions have been resolved:

- There was an issue where the **`voltadmin stop node`** command would not properly authenticate with the server when using Kerberos security. This issue has been resolved.
- There was an issue where, if a database was upgraded to a new version of VoltDB using **`voltadmin shutdown --save`**, restarted, but then crashes unexpectedly (for example, using **`kill -9`**), the database could not restart a second time. This issue has been resolved.

- Since VoltDB V7.7, it was possible for the response from a multi-partition read transaction to be "lost" during a node failure. This could only happen on K-safe clusters, where the node that failed was the multi-partition initiator (that is, the node responsible for coordinating multi-partition transactions), the node failed before the transaction completed, and the procedure call was invoked on a *different* node of the cluster. Under these conditions, the calling application might not receive a response from the invocation. Note that this issue only occurred on K-safe clusters, for read-only multi-partition queries, and for non-topology-aware clients. This issue has now been resolved.
- There was a sporadic problem with authentication of the VoltDB Management Center (VMC) on slow network connections. Starting with VoltDB 7.9, the JSON interface uses stay-alive connections with a timeout period of 10 seconds. VMC sends requests every 5 seconds. However, on slow networks the VMC calls could be delayed beyond the timeout period, forcing the user to re-authenticate manually. The timeout period has been extended to alleviate the unexpected timeouts.
- There was an issue with Kafka import where, if the database cluster was paused and then resumed, it was possible for certain Kafka records that were being processed when the database was paused to be lost. This issue has been resolved.
- There was a longstanding and somewhat obscure bug involving partitioned views joined to a derived table from a sub SELECT statement. If a partitioned view did not include the table's partition column and was joined to another table derived from a sub selection (that is a SELECT statement in parentheses in the FROM clause of the main SELECT) the query could result in unexpected behavior at runtime, including possibly crashing the database. This issue has been resolved.

28. Release V8.0 (February 6, 2018)

28.1. TLS/SSL encryption for intra-cluster communication

VoltDB now supports encrypting communication on the internal port, the port used for communication between nodes in the cluster, using TLS/SSL encryption. Note that encrypting the internal port automatically adds latency to any operations that require inter-node communication, such as K-safety and multi-partition procedures. The actual impact depends on the configuration and application workload. It is strongly recommended you benchmark your application before enabling internal TLS/SSL on production systems. See the chapter on "Security" in the *Using VoltDB* manual for details.

28.2. New behavior for placement groups

Placement groups, or rack-aware provisioning, was introduced in VoltDB 5.5. Placement groups let you specify where each node is located, so in a virtualized K-safe environment multiple copies of a partition are distributed onto distinct hardware, racks, etc. However, changes in VoltDB 7.0 to optimize K-safe partitioning in all cases ended up superceding placement groups and invalidating the rack-aware positioning.

This unintentional side effect has been corrected and placement groups once again provide rack-aware provisioning. However, the algorithm for interpreting placement groups has changed. Where before you could use a hierarchical list of names separated by periods (such as rack1.switch3.server5) the new algorithm focuses on the first name only and subnames are largely ignored.

Use of simple (non-hierarchical) placement names is recommended. In addition, the following rules apply to the top-level names:

- There must be more than one top-level group specified for the cluster
- The same number of nodes must be included in each group
- The number of partition copies (that is, K+1) must be a multiple of the number of top-level groups

28.3. Kafka 0.10.2 is now the default for Kafka import and kafkaloader

The default for the Kafka import connector and the kafkaloader command line utility has changed to support Kafka 0.10.2 and later, including the recent 1.0.0 release. Earlier versions of Kafka (0.8.2) are still supported through configuration options and an alternate kafkaloader8 utility.

28.4. Support for common table expressions

VoltDB now supports common table expressions, including recursive common table expressions. See the description of the SELECT statement in the *Using VoltDB* manual for details.

28.5. Deprecated features removed from the product

The following features, that had previously been deprecated, have now been removed from the product as of VoltDB 8.0:

- "Fast" read consistency (<consistency> element)
- Old, non-elastic partitioning (<cluster> elastic attribute)
- @ProcInfo Java annotation for specifying procedure partitioning
- Old shell commands (**volt** **add**, **create**, **recover**, and **rejoin**)

We are also deprecating the VoltDB Deployment Manager.

28.6. VoltDB Deployment Manager is deprecated

The VoltDB Deployment Manager was designed as a console for deploying VoltDB clusters. However, it has not met its goals for ease-of-use and flexibility. Therefore, we are deprecating it and it will be removed from the product in a future release. In its place we recommend using one of the existing frameworks for managing distributed systems such as (but not limited to) Chef, Puppet, Docker, and Kubernetes.

28.7. Security Notice

The following changes have been made to improve security and eliminate potential threats:

- Ability to enable TLS/SSL encryption on the internal port. (See description above.)

28.8. Additional improvements

In addition to the new features and capabilities described above, the following limitations in previous versions have been resolved:

- There was an issue using the SQL Query tab of the web-based VoltDB Management Console (VMC) to insert or filter records if any text fields in the query contained multiple consecutive spaces. (For example, two or more leading spaces or multiple spaces between two words.) Some of the spaces were interpreted as the UTF-8 character for a non-breaking space (\u00A0) rather than ASCII code 32, causing incorrect data insertion or filtering. This issue has been resolved.
- In the past it was possible for queries containing both a FULL or RIGHT OUTER JOIN and a GROUP BY operation on a floating point (FLOAT) column to produce incorrect results. This issue has been resolved.
- Previously, certain combinations of UNION, ORDER BY, and LIMIT clauses in a single query could produce incorrect results. This issue has been resolved.
- There was an issue with the Nagios script for monitoring replica clusters when using database replication (DR). The script could generate a series of false alarms if the replica database was idle for more than two minutes. The alarms incorrectly reported that replication was falling behind. This issue has been resolved.

Known Limitations

The following are known limitations to the current release of VoltDB. Workarounds are suggested where applicable. However, it is important to note that these limitations are considered temporary and are likely to be corrected in future releases of the product.

1. Command Logging

- 1.1. Do not use the subfolder name "segments" for the command log snapshot directory.

VoltDB reserves the subfolder "segments" under the command log directory for storing the actual command log files. Do not add, remove, or modify any files in this directory. In particular, do not set the command log snapshot directory to a subfolder "segments" of the command log directory, or else the server will hang on startup.

2. Database Replication

- 2.1. Some DR data may not be delivered if master database nodes fail and rejoin in rapid succession.

Because DR data is buffered on the master database and then delivered asynchronously to the replica, there is always the danger that data does not reach the replica if a master node stops. This situation is mitigated in a K-safe environment by all copies of a partition buffering on the master cluster. Then if a sending node goes down, another node on the master database can take over sending logs to the replica. However, if multiple nodes go down and rejoin in rapid succession, it is possible that some buffered DR data — from transactions when one or more nodes were down — could be lost when another node with the last copy of that buffer also goes down.

If this occurs and the replica recognizes that some binary logs are missing, DR stops and must be restarted.

To avoid this situation, especially when cycling through nodes for maintenance purposes, the key is to ensure that all buffered DR data is transmitted before stopping the next node in the cycle. You can do this using the @Statistics system procedure to make sure the last ACKed timestamp (using @Statistics DR on the master cluster) is later than the timestamp when the previous node completed its rejoin operation.

- 2.2. Avoid bulk data operations within a single transaction when using database replication

Bulk operations, such as large deletes, inserts, or updates are possible within a single stored procedure. However, if the binary logs generated for DR are larger than 45MB, the operation will fail. To avoid this situation, it is best to break up large bulk operations into multiple, smaller transactions. A general rule of thumb is to multiply the size of the table schema by the number of affected rows. For deletes and inserts, this value should be under 45MB to avoid exceeding the DR binary log size limit. For updates, this number should be under 22.5MB (because the binary log contains both the starting and ending row values for updates).

- 2.3. Database replication ignores resource limits

There are a number of VoltDB features that help manage the database by constraining memory size and resource utilization. These features are extremely useful in avoiding crashes as a result of unexpected or unconstrained growth. However, these features could interfere with the normal operation of DR when passing data from one cluster to another, especially if the two clusters are different sizes. Therefore, as a general rule of thumb, DR overrides these features in favor of maintaining synchronization between the two clusters.

Specifically, DR ignores any resource monitor limits defined in the deployment file when applying binary logs on the consumer cluster. DR also ignores any partition row limits defined in the database schema when applying binary logs. This means, for example, if the replica database in passive DR has less memory or fewer unique partitions than the master, it is possible that applying binary logs of transactions that succeeded on the master could cause the replica to run out of memory. Note that these resource monitor and tables row limits

are applied on any original transactions local to the cluster (for example, transactions on the master database in passive DR).

2.4. Different cluster sizes can require additional Java heap

Database Replication (DR) now supports replication across clusters of different sizes. However, if the replica cluster is smaller than the master cluster, it may require a significantly larger Java heap setting. Specifically, if the replica has fewer unique partitions than the master, each partition on the replica must manage the incoming binary logs from more partitions on the master, which places additional pressure on the Java heap.

A simple rule of thumb is that the worst case scenario could require an additional $P * R * 20\text{MB}$ space in the Java heap, where P is the number of sites per host on the replica server and R is the ratio of unique partitions on the master to partitions on the replica. For example, if the master cluster is 5 nodes with 10 sites per host and a K factor of 1 (i.e. 25 unique partitions) and the replica cluster is 3 nodes with 8 sites per host and a K factor of 1 (12 unique partitions), the Java heap on the replica cluster may require approximately 320MB of additional space in the heap:

```
Sites-per-host * master/replace ratio * 20MB  
8 * 25/12 * 20 = ~ 320MB
```

An alternative is to reduce the size of the DR buffers on the master cluster by setting the `DR_MEM_LIMIT` Java property. For example, you can reduce the DR buffer size from the default 10MB to 5MB using the `VOLTDB_OPTS` environment variable before starting the master cluster.

```
$ export VOLTDB_OPTS="-DDR_MEM_LIMIT=5"  
  
$ voltdb start
```

Changing the DR buffer limit on the master from 10MB to 5MB proportionally reduces the additional heap size needed. So in the previous example, the additional heap on the replica is reduced from 320MB to 160MB.

3. Cross Datacenter Replication (XDCR)

3.1. Avoid replicating tables without a unique index.

Part of the replication process for XDCR is to verify that the record's starting and ending states match on both clusters, otherwise known as *conflict resolution*. To do that, XDCR must find the record first. Finding uniquely indexed records is efficient; finding non-unique records is not and can impact overall database performance.

To make you aware of possible performance impact, VoltDB issues a warning if you declare a table as a DR table and it does not have a unique index.

3.2. When starting XDCR for the first time, only one database can contain data.

You cannot start XDCR if both databases already have data in the DR tables. Only one of the two participating databases can have preexisting data when DR starts for the first time.

3.3. During the initial synchronization of existing data, the receiving database is paused.

When starting XDCR for the first time, where one database already contains data, a snapshot of that data is sent to the other database. While receiving and processing that snapshot, the receiving database is paused. That is, it is in read-only mode. Once the snapshot is completed and the two database are synchronized, the receiving database is automatically unpaused, resuming normal read/write operations.

3.4. A large number of multi-partition write transactions may interfere with the ability to restart XDCR after a cluster stops and recovers.

Normally, XDCR will automatically restart where it left off after one of the clusters stops and recovers from its command logs (using the **voltodb recover** command). However, if the workload is predominantly multi-partition write transactions, a failed cluster may not be able to restart XDCR after it recovers. In this case, XDCR must be restarted from scratch, using the content from one of the clusters as the source for synchronizing and recreating the other cluster (using the **voltodb create --force** command) without any content in the DR tables.

3.5. Avoid using TRUNCATE TABLE in XDCR environments.

TRUNCATE TABLE is optimized to delete all data from a table rather than deleting tuples row by row. This means that the binary log does not identify which rows are deleted. As a consequence, a TRUNCATE TABLE statement and a simultaneous write operation to the same table can produce a conflict that the XDCR clusters cannot detect or report in the conflict log.

Therefore, do not use TRUNCATE TABLE with XDCR. Instead, explicitly delete all rows with a DELETE statement and a filter. For example, `DELETE * FROM table WHERE column=column` ensures all deleted rows are identified in the binary log and any conflicts are accurately reported. Note that `DELETE FROM table` without a WHERE clause is *not* sufficient, since its execution plan is optimized to equate to TRUNCATE TABLE.

3.6. Exceeding a LIMIT PARTITION ROWS constraint can generate multiple conflicts

It is possible to place a limit on the number of rows that any partition can hold for a specific table using the LIMIT PARTITION ROWS clause of the CREATE TABLE statement. When close to the limit, transactions on either or both clusters can exceed the limit simultaneously, resulting in a potentially large number of delete operations that then generate conflicts when the the associated binary log reaches the other cluster.

3.7. Use of the VoltProcedure.getUniqueId method is unique to a cluster, not across clusters.

VoltDB provides a way to generate a deterministically unique ID within a stored procedure using the getUniqueId method. This method guarantees uniqueness *within the current cluster*. However, the method could generate the same ID on two distinct database clusters. Consequently, when using XDCR, you should combine the return values of VoltProcedure.getUniqueId with VoltProcedure.getClusterId, which returns the current cluster's unique DR ID, to generate IDs that are unique across all clusters in your environment.

3.8. XDCR cannot be used with deprecated export syntax.

You cannot use cross datacenter replication (XDCR) with the deprecated export syntax, that is the EXPORT TABLE statement. To use XDCR with export, you must use the current CREATE STREAM syntax for declaring the source streams for export targets.

4. TTL

4.1. Use of TTL (time to live) with replicated tables and Database Replication (DR) can result in increased DR activity.

TTL, or time to live, is a feature that automatically deletes old records based on a timestamp or integer column. For replicated tables, the process of checking whether records need to be deleted is performed as a write transaction — even if no rows are deleted. As a consequence, any replicated DR table with TTL defined will generate frequent DR log entries, whether there are any changes or not, significantly increasing DR traffic.

Because of the possible performance impact this behavior can have on the database, use of TTL with replicated tables and DR is not recommended at this time.

5. Export

5.1. Synchronous export in Kafka can use up all available file descriptors and crash the database.

A bug in the Apache Kafka client can result in file descriptors being allocated but not released if the `producer.type` attribute is set to "sync" (which is the default). The consequence is that the system eventually runs out of file descriptors and the VoltDB server process will crash.

Until this bug is fixed, use of synchronous Kafka export is not recommended. The workaround is to set the Kafka `producer.type` attribute to "async" using the VoltDB export properties.

- 5.2. Changing the `enabled="true"` attribute of an export configuration does not properly stop or start the export connector.

With the latest version of export, the proper way to stop or start an export connector on a running database is to add or remove the `<configuration>` element for that connector from the database configuration file. Simply changing the `enabled` attribute from "true" to "false" (or vice versa) does not properly stop or start the connector and additional rows may be leaked to the export target after the configuration is updated.

6. Import

- 6.1. Data may be lost if a Kafka broker stops during import.

If, while Kafka import is enabled, the Kafka broker that VoltDB is connected to stops (for example, if the server crashes or is taken down for maintenance), some messages may be lost between Kafka and VoltDB. To ensure no data is lost, we recommend you disable VoltDB import before taking down the associated Kafka broker. You can then re-enable import after the Kafka broker comes back online.

- 6.2. Kafka import can lose data if multiple nodes stop in succession.

There is an issue with the Kafka importer where, if multiple nodes in the cluster fail and restart, the importer can lose track of some of the data that was being processed when the nodes failed. Normally, these pending imports are replayed properly on restart. But if multiple nodes fail, it is possible for some in-flight imports to get lost. This issue will be addressed in an upcoming release.

7. SQL and Stored Procedures

- 7.1. Comments containing unmatched single quotes in multi-line statements can produce unexpected results.

When entering a multi-line statement at the `sqlcmd` prompt, if a line ends in a comment (indicated by two hyphens) and the comment contains an unmatched single quote character, the following lines of input are not interpreted correctly. Specifically, the comment is incorrectly interpreted as continuing until the next single quote character or a closing semi-colon is read. This is most likely to happen when reading in a schema file containing comments. This issue is specific to the `sqlcmd` utility.

A fix for this condition is planned for an upcoming point release

- 7.2. Do not use assertions in VoltDB stored procedures.

VoltDB currently intercepts assertions as part of its handling of stored procedures. Attempts to use assertions in stored procedures for debugging or to find programmatic errors will not work as expected.

- 7.3. The `UPPER()` and `LOWER()` functions currently convert ASCII characters only.

The `UPPER()` and `LOWER()` functions return a string converted to all uppercase or all lowercase letters, respectively. However, for the initial release, these functions only operate on characters in the ASCII character set. Other case-sensitive UTF-8 characters in the string are returned unchanged. Support for all case-sensitive UTF-8 characters will be included in a future release.

8. Client Interfaces

- 8.1. Avoid using decimal datatypes with the C++ client interface on 32-bit platforms.

There is a problem with how the math library used to build the C++ client library handles large decimal values on 32-bit operating systems. As a result, the C++ library cannot serialize and pass Decimal datatypes reliably on these systems.

Note that the C++ client interface *can* send and receive Decimal values properly on 64-bit platforms.

9. SNMP

9.1. Enabling SNMP traps can slow down database startup.

Enabling SNMP can take up to 2 minutes to complete. This delay does not always occur and can vary in length. If SNMP is enabled when the database server starts, the delay occurs after the server logs the message "Initializing SNMP" and before it attempts to connect to the cluster. If you enable SNMP while the database is running, the delay can occur when you issue the **voltadmin update** command or modify the setting in the VoltDB Management Center Admin tab. This issue results from a Java constraint related to secure random numbers used by the SNMP library.

10. VoltDB Management Center

10.1. The VoltDB Management Center currently reports on only one DR connection.

With VoltDB V7.0, cross datacenter replication (XDCR) supports multiple clusters in an XDCR network. However, the VoltDB Management Center currently reports on only one such connection per cluster. In the future, the Management Center will provide monitoring and statistics for all connections to the current cluster.

Implementation Notes

The following notes provide details concerning how certain VoltDB features operate. The behavior is not considered incorrect. However, this information can be important when using specific components of the VoltDB product.

1. VoltDB Management Center

1.1. Schema updates clear the stored procedure data table in the Management Center Monitor section

Any time the database schema or stored procedures are changed, the data table showing stored procedure statistics at the bottom of the Monitor section of the VoltDB Management Center get reset. As soon as new invocations of the stored procedures occur, the statistics table will show new values based on performance after the schema update. Until invocations occur, the procedure table is blank.

2. SQL

2.1. You cannot partition a table on a column defined as ASSUMEUNIQUE.

The ASSUMEUNIQUE attribute is designed for identifying columns in partitioned tables where the column values are known to be unique but the table is not partitioned on that column, so VoltDB cannot verify complete uniqueness across the database. Using interactive DDL, you can create a table with a column marked as ASSUMEUNIQUE, but if you try to partition the table on the ASSUMEUNIQUE column, you receive an error. The solution is to drop and add the column using the UNIQUE attribute instead of ASSUMEUNIQUE.

2.2. Adding or dropping column constraints (UNIQUE or ASSUMEUNIQUE) is not supported by the ALTER TABLE ALTER COLUMN statement.

You cannot add or remove a column constraint such as UNIQUE or ASSUMEUNIQUE using the ALTER TABLE ALTER COLUMN statement. Instead to add or remove such constraints, you must first drop then add the modified column. For example:

```
ALTER TABLE employee DROP COLUMN empID;  
ALTER TABLE employee ADD COLUMN empID INTEGER UNIQUE;
```

2.3. Do not use UPDATE to change the value of a partitioning column

For partitioned tables, the value of the column used to partition the table determines what partition the row belongs to. If you use UPDATE to change this value and the new value belongs in a different partition, the UPDATE request will fail and the stored procedure will be rolled back.

Updating the partition column value may or may not cause the record to be repartitioned (depending on the old and new values). However, since you cannot determine if the update will succeed or fail, you should not use UPDATE to change the value of partitioning columns.

The workaround, if you must change the value of the partitioning column, is to use both a DELETE and an INSERT statement to explicitly remove and then re-insert the desired rows.

2.4. Ambiguous column references no longer allowed.

Starting with VoltDB 6.0, ambiguous column references are no longer allowed. For example, if both the *Customer* and *Placedorder* tables have a column named *Address*, the reference to *Address* in the following SELECT statement is ambiguous:

```
SELECT OrderNumber, Address FROM Customer, Placedorder  
. . .
```

Previously, VoltDB would select the column from the leftmost table (*Customer*, in this case). Ambiguous column references are no longer allowed and you must use table prefixes to disambiguate identical column names. For example, specifying the column in the preceding statement as *Customer.Address*.

A corollary to this change is that a column declared in a USING clause can now be referenced using a prefix. For example, the following statement uses the prefix *Customer.Address* to disambiguate the column selection from a possibly similarly named column belonging to the *Supplier* table:

```
SELECT OrderNumber, Vendor, Customer.Address  
FROM Customer, Placedorder Using (Address), Supplier  
. . .
```

3. Export

3.1. Export Overflow and Durability

VoltDB uses the export overflow directory on disk to queue pending export data before it is sent and acknowledged by the export target. Some of this data is retained, even after the queue is flushed, to ensure durability across database sessions. Up to 64MB may be retained per stream, per partition for this purpose. You should take this disk space requirement into consideration when configuring a database with multiple export streams.

Note that because some data is retained for durability purposes after it is written, you cannot use disk usage of this directory as an indicator of whether export is drained or not. Use the @Statistics system procedure with the EXPORT selector to determine the current state of the export streams.

4. Runtime

4.1. File Descriptor Limits

VoltDB opens a file descriptor for every client connection to the database. In normal operation, this use of file descriptors is transparent to the user. However, if there are an inordinate number of concurrent client

connections, or clients open and close many connections in rapid succession, it is possible for VoltDB to exceed the process limit on file descriptors. When this happens, new connections may be rejected or other disk-based activities (such as snapshotting) may be disrupted.

In environments where there are likely to be an extremely large number of connections, you should consider increasing the operating system's per-process limit on file descriptors.

4.2. Use of Resources in JAR Files

There are two ways to access additional resources in a VoltDB database. You can place the resources in the `/lib` folder where VoltDB is installed on each server in the cluster or you can include the resource in a subfolder of a JAR file you add using the sqlcmd **LOAD CLASSES** directive. Adding resources via the `/lib` directory is useful for stable resources (such as third-party software libraries) that do not require updating. Including resources (such as XML files) in the JAR file is useful for resources that may need to be updated, as a single transaction, while the database is running.

LOAD CLASSES is used primarily to load classes associated with stored procedures and user-defined functions. However, it will also load any additional resource files included in subfolders of the JAR file. You can remove classes that are no longer needed using the **REMOVE CLASSES** directive. However, there is no explicit command for removing other resources.

Consequently, if you rename resources or move them to a different location and reload the JAR file, the database will end up having multiple copies. Over time, this could result in more and more unnecessary memory being used by the database. To remove obsolete resources, you must first reinitialize the database root directory, start a fresh database, reload the schema (including the new JAR files with only the needed resources) and then restore the data from a snapshot.

4.3. Servers with Multiple Network Interfaces

If a server has multiple network interfaces (and therefore multiple IP addresses) VoltDB will, by default, open ports on all available interfaces. You can limit the ports to a single interface in two ways:

- Specify which interface to use for internal and external ports, respectively, using the **--internalinterface** and **--externalinterface** arguments when starting the database process with the **voltadb start** command.
- For an individual port, specify the interface and port on the command line. For example **voltadb start --client=32.31.30.29:21212**.

Also, when using an IP address to reference a server with multiple interfaces in command line utilities (such as **voltadmin stop node**), use the `@SystemInformation` system procedure to determine which IP address VoltDB has selected to identify the server. Otherwise, if you choose the wrong IP address, the command might fail.