

Network Working Group
Internet-Draft
Intended status: Informational
Expires: February 21, 2013

B. Pfaff
B. Davie, Ed.
Nicira, Inc.
August 20, 2012

The Open vSwitch Database Management Protocol
draft-pfaff-ovsdb-proto-00

Abstract

Open vSwitch is an open source software switch designed to be used as a vswitch (virtual switch) in virtualized server environments. A vswitch forwards traffic between different virtual machines (VMs) on the same physical host and also forwards traffic between VMs and the physical network. Open vSwitch is open to programmatic extension and control using OpenFlow and the OVSDB (Open vSwitch Database) management protocol. This document defines the OVSDB management protocol.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on February 21, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Requirements Language	3
1.2.	Terminology	3
2.	System Overview	4
3.	OVSDB Structure	5
3.1.	JSON Usage	5
3.2.	Schema Format	7
4.	Wire Protocol	11
4.1.	RPC Methods	11
4.1.1.	List_dbs	12
4.1.2.	Get_schema	12
4.1.3.	Transact	13
4.1.4.	Cancel	15
4.1.5.	Monitor	15
4.1.6.	Update Notification	17
4.1.7.	Monitor_Cancel	18
4.1.8.	Lock Operations	18
4.1.9.	Locked Notification	20
4.1.10.	Stolen Notification	20
4.1.11.	Echo	21
5.	Database Operations	21
5.1.	Notation	21
5.2.	Operations	25
5.2.1.	Insert	25
5.2.2.	Select	26
5.2.3.	Update	27
5.2.4.	Mutate	28
5.2.5.	Delete	28
5.2.6.	Wait	29
5.2.7.	Commit	30
5.2.8.	Abort	30
5.2.9.	Comment	30
5.2.10.	Assert	31
6.	IANA Considerations	31
7.	Security Considerations	31
8.	Acknowledgements	31
9.	References	32
9.1.	Normative References	32
9.2.	Informative References	32
	Authors' Addresses	32

1. Introduction

In virtualized server environments, it is typically required to use a vswitch (virtual switch) to forward traffic between different virtual machines (VMs) on the same physical host, and between VMs and the physical network. Open vSwitch [[OVS](#)] is an open source software switch designed to be used as a vswitch in such environments. Open vSwitch (OVS) is open to programmatic extension and control using OpenFlow [[OF-SPEC](#)] and the OVSDB (Open vSwitch Database) management protocol. This document defines the OVSDB management protocol.

The OVSDB management protocol uses JSON[RFC4627] for its wire format, and is based on JSON-RPC version 1.0 [[JSON-RPC](#)].

The schema of the Open vSwitch database is documented in [[DB-SCHEMA](#)]. This document specifies the protocol for interacting with that database for the purposes of managing and configuring Open vSwitch instances. The protocol specified in this document also provides means for discovering the schema in use, as described in [Section 4.1.2](#).

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in[RFC2119].

1.2. Terminology

UUID: Universally Unique Identifier. A 128-bit identifier that is unique in space and time [[DCE](#)].

OVS: Open vSwitch. An open source virtual switch.

OVSDB: The database that is used for the purpose of configuring OVS instances.

JSON: Javascript Object Notation [[RFC4627](#)].

JSON-RPC: JSON Remote Procedure Call [[JSON-RPC](#)].

Note that the JSON specification [[RFC4627](#)] provides precise definitions of a number of important terms such as JSON values, objects, arrays, numbers, and strings. In all cases, this document uses the definitions from [[RFC4627](#)].

2. System Overview

Figure 1 illustrates the main components of Open vSwitch and the interfaces to a control and management cluster. An OVS instance comprises a database server (ovsdb-server), a vswitch daemon (ovs-vswitchd), and a kernel module that performs fast path forwarding. The "management and control cluster" consists of some number of managers and controllers. Managers use the OVSDB management protocol to manage OVS instances. An OVS instance is managed by at least one manager. Controllers use OpenFlow to install flow state in OpenFlow switches. An OVS instance can support multiple logical datapaths, referred to as bridges. There is at least one controller for each OpenFlow bridge.

The OVSDB management interface is used to perform management and configuration operations on the OVS instance. Compared to OpenFlow, OVSDB management operations occur at a relatively long timescale. Examples of operations that are supported by OVSDB include:

- o Creation, modification and deletion of OpenFlow datapaths (bridges), of which there may be many in a single OVS instance;
- o Configuration of the set of controllers to which an OpenFlow datapath should connect;
- o Configuration of the set of managers to which the OVSDB server should connect;
- o Creation, modification and deletion of ports on OpenFlow datapaths;
- o Creation, modification and deletion of tunnel interfaces on OpenFlow datapaths;
- o Creation, modification and deletion of queues;
- o Configuration of QoS (quality of service) policies, and attachment of those policies to queues;
- o Collection of statistics.




Figure 1: Open vSwitch Interfaces

Further information about the usage of the OVSDB management protocol is provided in [\[DB-SCHEMA\]](#).

3. OVSDB Structure

This section outlines the overall structure of databases in OVSDB. As described here, the database is reasonably generic. For the complete and current description of the database schema as used in OVS, refer to [\[DB-SCHEMA\]](#). See also [Section 4.1.2](#) for information on how the OVSDB protocol may be used to discover the schema currently in use.

3.1. JSON Usage

OVSDB uses JSON [\[RFC4627\]](#) for both its schema format and its wire protocol format. The JSON implementation in Open vSwitch has the following limitations:

- o Null bytes (`\u0000`) SHOULD NOT be used in strings.
- o Only UTF-8 encoding is supported.

The descriptions below use the following shorthand notations for JSON values. Terminology follows [\[RFC4627\]](#).

`<string>` A JSON string. Any Unicode string is allowed. Implementations MAY disallow null bytes.

`<id>` A JSON string matching `[a-zA-Z_][a-zA-Z0-9_]*`. `<id>`s that begin with `_` are reserved to the implementation and MUST NOT be used by the user.

`<version>` A JSON string that contains a version number that matches `[0-9]+\.[0-9]+\.[0-9]+`

`<boolean>` A JSON true or false value.

`<number>` A JSON number.

`<integer>` A JSON number with an integer value, within the range `-(2**63)...+(2**63)-1`.

`<json-value>` Any JSON value.

`<nonnull-json-value>` Any JSON value except null.

`<error>` A JSON object with the following members:

<code>"error": <string></code>	required
<code>"details": <string></code>	optional

The value of the "error" member is a short string, specified in this document, that broadly indicates the class of the error. Most "error" strings are specific to contexts described elsewhere in this document, but the following "error" strings may appear in any context where an `<error>` is permitted:

"error": "resources exhausted" The operation requires more resources (memory, disk, CPU, etc.) than are currently available to the database server.

"error": "I/O error" Problems accessing the disk, network, or other required resources prevented the operation from completing.

Database implementations MAY use "error" strings not specified in this document to indicate errors that do not fit into any of the specified categories. Optionally, an `<error>` MAY include a "details" member, whose value is a string that describes the error in more detail for the benefit of a human user or administrator. This document does not specify the format or content of the "details" string. An `<error>` MAY also have other members that describe the error in more detail. This

document does not specify the names or values of these members.

3.2. Schema Format

An Open vSwitch configuration database consists of a set of tables, each of which has a number of columns and zero or more rows. A schema for the database is represented by `<database-schema>`, as described below.

`<database-schema>`

A JSON object with the following members:

"name": <code><id></code>	required
"version": <code><version></code>	required
"cksum": <code><string></code>	optional
"tables": { <code><id></code> : <code><table-schema></code> , ...}	required

The "name" identifies the database as a whole. It must be provided to most JSON-RPC requests to identify the database being operated on. The value of "tables" is a JSON object whose names are table names and whose values are `<table-schema>`s.

The "version" reports the version of the database schema. Because this is a recent addition to the schema format, OVSDB permits it to be omitted, but it SHOULD be present. Future versions of OVSDB will require it to be present. Open vSwitch semantics for "version" are described in [[DB-SCHEMA](#)]

The "cksum" optionally reports an implementation-defined checksum for the database schema.

The value of "tables" is a JSON object whose names are table names and whose values are `<table-schema>`s.

`<table-schema>`

A JSON object with the following members:

"columns": { <code><id></code> : <code><column-schema></code> , ...}	required
"maxRows": <code><integer></code>	optional
"isRoot": <code><boolean></code>	optional
"indexes": [<code><column-set></code> *]	optional

The value of "columns" is a JSON object whose names are column names and whose values are <column-schema>s.

Every table has the following columns whose definitions are not included in the schema:

"_uuid": This column, which contains exactly one UUID value, is initialized to a random value by the database engine when it creates a row. It is read-only, and its value never changes during the lifetime of a row.

"_version": Like "_uuid", this column contains exactly one UUID value, initialized to a random value by the database engine when it creates a row, and it is read-only. However, its value changes to a new random value whenever any other field in the row changes. Furthermore, its value is ephemeral: when the database is closed and reopened, or when the database process is stopped and then started again, each "_version" also changes to a new random value.

The "isRoot" boolean is used to determine whether rows in the table require strong references from other rows to avoid garbage collection. If "isRoot" is specified as true, then rows in the table exist independent of any references (they can be thought of as part of the "root set" in a garbage collector). If "isRoot" is omitted or specified as false, then any given row in the table may exist only when there is at least one reference to it, with refType "strong", from a different row (in the same table or a different table). This is a "deferred" action: unreferenced rows in the table are deleted just before transaction commit.

For compatibility with schemas created before "isRoot" was introduced, if "isRoot" is omitted or false in every <table-schema> in a given <database-schema>, then every table is part of the root set.

If "maxRows" is specified, as a positive integer, it limits the maximum number of rows that may be present in the table. This is a "deferred" constraint, enforced only at transaction commit time (see the "transact" request below). If "maxRows" is not specified, the size of the table is limited only by the resources available to the database server. "maxRows" constraints are enforced after unreferenced rows are deleted from tables with a false "isRoot".

If "indexes" is specified, it must be an array of zero or more <column-set>s. A <column-set> is an array of one or more strings, each of which names a column. Each <column-set> is a set of columns whose values, taken together within any given row, must be unique within the table. This is a "deferred" constraint, enforced only at transaction commit time, after unreferenced rows are deleted and dangling weak references are removed. Ephemeral columns may not be part of indexes.

<column-schema>

A JSON object with the following members:

"type": <type>	required
"ephemeral": <boolean>	optional
"mutable": <boolean>	optional

The "type" specifies the type of data stored in this column.

If "ephemeral" is specified as true, then this column's values are not guaranteed to be durable; they may be lost when the database restarts. A column whose type (either key or value) is a strong reference to a table that is not part of the root set is always durable, regardless of this value. (Otherwise, restarting the database could lose entire rows.)

If "mutable" is specified as false, then this column's values may not be modified after they are initially set with the "insert" operation.

<type>

The type of a database column. Either an <atomic-type> or a JSON object that describes the type of a database column, with the following members:

"key": <base-type>	required
"value": <base-type>	optional
"min": <integer>	optional
"max": <integer> or "unlimited"	optional

If "min" or "max" is not specified, each defaults to 1. If "max" is specified as "unlimited", then there is no specified maximum number of elements, although the implementation will enforce some limit. After considering defaults, "min" must be exactly 0 or exactly 1, "max" must be at least 1, and "max" must be greater

than or equal to "min".

If "min" and "max" are both 1 and "value" is not specified, the type is the scalar type specified by "key".

If "min" is not 1 or "max" is not 1, or both, and "value" is not specified, the type is a set of scalar type "key".

If "value" is specified, the type is a map from type "key" to type "value".

<base-type>

The type of a key or value in a database column. Either an <atomic-type> or a JSON object with the following members:

"type": <atomic-type>	required
"enum": <value>	optional
"minInteger": <integer>	optional, integers only
"maxInteger": <integer>	optional, integers only
"minReal": <real>	optional, reals only
"maxReal": <real>	optional, reals only
"minLength": <integer>	optional, strings only
"maxLength": <integer>	optional, strings only
"refTable": <id>	optional, uuids only
"refType": "strong" or "weak"	optional, only with "refTable"

An <atomic-type> by itself is equivalent to a JSON object with a single member "type" whose value is the <atomic-type>.

"enum" may be specified as a <value> whose type is a set of one or more values specified for the member "type". If "enum" is specified, then the valid values of the <base-type> are limited to those in the <value>.

"enum" is mutually exclusive with the following constraints.

If "type" is "integer", then "minInteger" or "maxInteger" or both may also be specified, restricting the valid integer range. If both are specified, then the maxInteger must be greater than or equal to minInteger.

If "type" is "real", then "minReal" or "maxReal" or both may also be specified, restricting the valid real range. If both are specified, then the maxReal must be greater than or equal to minReal.

If "type" is "string", then "minLength" and "maxLength" or both may be specified, restricting the valid length of value strings. If both are specified, then maxLength must be greater than or equal to minLength. String length is measured in characters (not bytes or UTF-16 code units).

If "type" is "uuid", then "refTable", if present, must be the name of a table within this database. If "refTable" is specified, then "refType" may also be specified. If "refTable" is set, the effect depends on "refType":

- * If "refType" is "strong" or if "refType" is omitted, the allowed UUIDs are limited to UUIDs for rows in the named table.
- * If "refType" is "weak", then any UUIDs are allowed, but UUIDs that do not correspond to rows in the named table will be automatically deleted.

"refTable" constraints are "deferred" constraints: they are enforced only at transaction commit time (see the "transact" request below). The other constraints on <base-type> are "immediate", enforced immediately by each operation.

<atomic-type>

One of the strings "integer", "real", "boolean", "string", or "uuid", representing the specified scalar type.

4. Wire Protocol

The database wire protocol is implemented in JSON-RPC 1.0 [[JSON-RPC](#)]. While the spec of JSON-RPC allows a range of transports, implementations of this specification SHOULD operate directly over TCP. See [Section 6](#) for discussion of the TCP port.

4.1. RPC Methods

The following subsections describe the RPC methods that are supported. As described in the JSON-RPC 1.0 specification, each request comprises a string containing the name of the method, a (possibly null) array of parameters to pass to the method, and a request ID, which can be used to match the response to the request. Each response comprises a result object (non-null in the event of a successful invocation), an error object (non-null in the event of an error), and the ID of the matching request. More details on each method, its parameters and results are described below.

The operations that may be performed on the OVS database using these methods (e.g., the "Transact" method) are described in [Section 5](#).

[4.1.1](#). List_dbs

This operation retrieves an array whose elements are the names of the databases that can be accessed over this management protocol connection.

The request object contains the following members:

- o "method": "list_dbs"
- o "params": []
- o "id": <nonnull-json-value>

The response object contains the following members:

- o "result": [<db-name>, ...]
- o "error": null
- o "id": same "id" as request

[4.1.2](#). Get_schema

This operation retrieves a <database-schema> that describes hosted database <db-name>.

The request object contains the following members:

- o "method": "get_schema"
- o "params": [<db-name>]
- o "id": <nonnull-json-value>

The response object contains the following members:

- o "result": <database-schema>
- o "error": null
- o "id": same "id" as request

4.1.3. Transact

This RPC method causes the database server to execute a series of operations in the specified order on a given database.

The request object contains the following members:

- o "method": "transact"
- o "params": [<db-name>, <operation>*]
- o "id": <nonnull-json-value>

The value of "id" MUST be unique among all in-flight transactions within the current JSON-RPC session. Otherwise, the server may return a JSON-RPC error.

The "params" array for this method consists of a <db-name> that identifies the database to which the transaction applies, followed by zero or more JSON objects, each of which represents a single database operation. [Section 5](#) describes the valid operations. The database server executes each of the specified operations in the specified order, except that if an operation fails, then the remaining operations are not executed. The set of operations is executed as a single atomic, consistent, isolated transaction. The transaction is committed only if every operation succeeds. Durability of the commit is not guaranteed unless the "commit" operation, with "durable" set to true, is included in the operation set. See [Section 5](#) for more discussion of the database operations.

The response object contains the following members:

- o "result": [<object>*]
- o "error": null
- o "id": same "id" as request

Regardless of whether errors occur in the database operations, the response is always a JSON-RPC response with null "error" and a "result" member that is an array with the same number of elements as "params". Each element of the "result" array corresponds to the same element of the "params" array. The "result" array elements may be interpreted as follows:

- o A JSON object that does not contain an "error" member indicates that the operation completed successfully. The specific members of the object are specified below in the descriptions of

individual operations. Some operations do not produce any results, in which case the object will have no members.

- o An <error> indicates that the matching operation completed with an error.
- o A JSON null value indicates that the operation was not attempted because a prior operation failed.

In general, "result" contains some number of successful results, possibly followed by an error, in turn followed by enough JSON null values to match the number of elements in "params". There is one exception: if all of the operations succeed, but the results cannot be committed, then "result" will have one more element than "params", with the additional element being an <error>. In this case, the possible "error" strings include the following:

"error": "referential integrity violation"

When the commit was attempted, a column's value referenced the UUID for a row that did not exist in the table named by the column's <base-type> key or value "refTable" that has a "refType" of "strong". (This can be caused by inserting a row that references a nonexistent row, by deleting a row that is still referenced by another row, by specifying the UUID for a row in the wrong table, and other ways.)

"error": "constraint violation"

A number of situations can arise in which the attempted commit would lead to a constraint on the database being violated. (See [Section 3.2](#) for more discussion of constraints.) These situations include:

- * The number of rows in a table exceeds the maximum number permitted by the table's "maxRows" value.
- * Two or more rows in a table had the same values in the columns that comprise an index.
- * A column with a <base-type> key or value "refTable" whose "refType" is "weak" became empty due to deletion(s), and this column is not allowed to be empty because its <type> has a "min" of 1. Such deletions may be the result of rows that it referenced being deleted (or never having existed, if the column's row was inserted within the transaction).


```
"error": "resources exhausted"
```

The operation requires more resources (memory, disk, CPU, etc.) than are currently available to the database server.

```
"error": "I/O error"
```

Problems accessing the disk, network, or other required resources prevented the operation from completing.

If "params" contains one or more "wait" operations, then the transaction may take an arbitrary amount of time to complete. The database implementation **MUST** be capable of accepting, executing, and replying to other transactions and other JSON-RPC requests while a transaction or transactions containing "wait" operations are outstanding on the same or different JSON-RPC sessions.

4.1.4. Cancel

The "cancel" method is a JSON-RPC notification, i.e. no matching response is provided. It instructs the database server to immediately complete or cancel the "transact" request whose "id" is the same as the notification's "params" value. The notification object has the following members:

- o "method": "cancel"
- o "params": [the "id" for an outstanding request]
- o "id": null

If the "transact" request can be completed immediately, then the server sends a response in the form described for "transact", above ([Section 4.1.3](#)). Otherwise, the server sends a JSON-RPC error response of the following form:

- o "result": null
- o "error": "canceled"
- o "id": the "id" member of the canceled request.

The "cancel" notification itself has no reply.

4.1.5. Monitor

The "monitor" request enables a client to replicate tables or subsets of tables within an OVSDB database by requesting notifications of

changes to those tables and by receiving the complete initial state of a table or a subset of a table. The request object has the following members:

- o "method": "monitor"
- o "params": [<db-name>, <json-value>, <monitor-requests>]
- o "id": <nonnull-json-value>

The <json-value> parameter is used to match subsequent update notifications (see below) to this request. The <monitor-requests> object comprises the name of the table to be monitored and an array of <monitor-request> objects. (For backward compatibility, a single <monitor-request> MAY be used instead of an array; it is treated as a single-element array.)

Each <monitor-request> is an object with the following members:

"columns": [<column>*]	optional
"select": <monitor-select>	optional

The columns, if present, defined the columns within the table to be monitored. <monitor-select> is an object with the following members:

"initial": <boolean>	optional
"insert": <boolean>	optional
"delete": <boolean>	optional
"modify": <boolean>	optional

The contents of this object specify how the columns or table are to be monitored, as explained in more detail below.

The response object has the following members:

- o "result": <table-updates>
- o "error": null
- o "id": same "id" as request

The <table-updates> object is described in detail in [Section 4.1.6](#). It contains the contents of the tables for which "initial" rows are selected. If no tables' initial contents are requested, then "result" is an empty object.

Subsequently, when changes to the specified tables are committed, the changes are automatically sent to the client using the "update" monitor notification (see [Section 4.1.6](#)). This monitoring persists until the JSON-RPC session terminates or until the client sends a

"monitor_cancel" JSON-RPC request.

Each <monitor-request> specifies one or more columns and the manner in which the columns (or the entire table) are to be monitored. The "columns" member specifies the columns whose values are monitored. It must not contain duplicates. If "columns" is omitted, all columns in the table, except for "_uuid", are monitored. The circumstances in which an "update" notification is sent for a row within the table are determined by <monitor-select>:

- o If "initial" is omitted or true, every row in the table is sent as part of the response to the "monitor" request.
- o If "insert" is omitted or true, "update" notifications are sent for rows newly inserted into the table.
- o If "delete" is omitted or true, "update" notifications are sent for rows deleted from the table.
- o If "modify" is omitted or true, "update" notifications are sent whenever when a row in the table is modified.

If there is more than one <monitor-request> in an array of them, then each <monitor-request> in the array should specify both "columns" and "select", and the "columns" MUST be non-overlapping sets.

4.1.6. Update Notification

The "update" notification is sent by the server to the client to report changes in tables that are being monitored following a "monitor" request as described above. The notification has the following members:

- o "method": "update"
- o "params": [<json-value>, <table-updates>]
- o "id": null

The <json-value> in "params" is the same as the value passed as the <json-value> in "params" for the corresponding "monitor" request. <table-updates> is an object that maps from a table name to a <table-update>. A <table-update> is an object that maps from the row's UUID (as a 36-byte string) to a <row-update> object. A <row-update> is an object with the following members:

- "old": <row> present for "delete" and "modify" updates
- "new": <row> present for "initial", "insert", and "modify" updates

Each table in which one or more rows has changed (or whose initial view is being presented) is represented in "updates". Each row that has changed (or whose initial view is being presented) is represented in its <table-update> as a member with its name taken from the row's _uuid member. The corresponding value is a <row-update>:

- o The "old" member is present for "delete" and "modify" updates. For "delete" updates, each monitored column is included. For "modify" updates, the prior value of each monitored column whose value has changed is included (monitored columns that have not changed are represented in "new").
- o The "new" member is present for "initial", "insert", and "modify" updates. For "initial" and "insert" updates, each monitored column is included. For "modify" updates, the new value of each monitored column is included.

4.1.7. Monitor_Cancel

The "monitor_cancel" request cancels a previously issued monitor request. The request object members are:

- o "method": "monitor_cancel"
- o "params": [<json-value>]
- o "id": <nonnull-json-value>

The <json-value> in "params" matches the <json-value> in "params" for the ongoing "monitor" request that is to be canceled. No more "update" messages will be sent for this table monitor. The response to this request has the following members:

- o "result": {}
- o "error": null
- o "id": the request "id" member

4.1.8. Lock Operations

Three RPC methods, "lock", "steal", and "unlock", allow a client to perform locking operations on the database. The database server supports an arbitrary number of locks, each of which is identified by a client-defined id. At any given time, each lock may have at most one owner. The RPC request objects have the following members:

- o "method": "lock", "steal", or "unlock"
- o "params": [<id>]
- o "id": <nonnull-json-value>

The response depends on the request, and has the following members:

- o "result": {"locked": boolean} for "lock"
- o "result": {"locked": true} for "steal"
- o "result": {} for "unlock"
- o "error": null
- o "id": same "id" as request

The three methods operate as follows:

- o "lock": The database will assign this client ownership of the lock as soon as it becomes available. When multiple clients request the same lock, they will receive it in first-come, first served order.
- o "steal": The database immediately assigns this client ownership of the lock. If there is an existing owner, it loses ownership.
- o "unlock": If the client owns the lock, releases it. If the client is waiting to obtain the lock, this cancels the request and stops waiting.

(Closing or otherwise disconnecting a database client connection unlocks all of its locks.)

For any given lock, the client MUST alternate "lock" or "steal" operations with "unlock" operations. That is, if the previous operation on a lock was "lock" or "steal", it MUST be followed by an "unlock" operation, and vice versa.

For a "lock" operation, the "locked" member in the response object is true if the lock has already been acquired, false if another client holds the lock and the client's request for it was queued. In the latter case, the client will be notified later with a "locked" message ([Section 4.1.9](#)) when acquisition succeeds.

These requests complete and send a response quickly, without waiting. The "locked" and "stolen" notifications (see below) report

asynchronous changes to ownership.

Note that the scope of a lock is a database server, not a database hosted by that server. A client may choose to implement a naming convention, such as "<db-name>__<lock-name>", which can effectively limit the scope of a lock to a particular database.

4.1.9. Locked Notification

The "locked" notification is provided to notify a client that it has been granted a lock it had previously request a lock with the "lock" method described above. The notification has the following members:

- o "method": "locked"
- o "params": [<id>]
- o "id": null

"Params" contains the name of the lock that was given in the "lock" request. The notified client now owns the lock named in "params".

The database server sends this notification after the reply to the corresponding "lock" request (but only if the "locked" member of the response was false), and before the reply to the client's subsequent "unlock" request.

4.1.10. Stolen Notification

The "stolen" notification is provided to notify a client, which had previously obtained a lock, that another client has stolen ownership of that lock. The notification has the following members:

- o "method": "stolen"
- o "params": [<id>]
- o "id": null

The notified client no longer owns the lock named in "params". The client **MUST** still issue an "unlock" request before performing any subsequent "lock" or "steal" operation on the lock.

If the client originally obtained the lock through a "lock" request, then it will automatically regain the lock later after the client that stole it releases it. (The database server will send the client a "locked" notification at that point to let it know.)

If the client originally obtained the lock through a "steal" request, the database server won't automatically reassign it ownership of the lock when it later becomes available. To regain ownership, the client must "unlock" and then "lock" or "steal" the lock again.

4.1.11. Echo

The "echo" method can be used by both clients and servers to verify the liveness of a database connection. It MUST be implemented by both clients and servers. The members of the request are:

- o "method": "echo"
- o "params": JSON array with any contents
- o "id": <json-value>

The response object has the following members:

- o "result": same as "params"
- o "error": null
- o "id": the request "id" member

5. Database Operations

This section describes the operations that may be specified in the "transact" method described in [Section 4.1.3](#).

5.1. Notation

We introduce the following notation for the discussion of operations.

<db-name>

An <id> that names a database. The valid <db-name>s can be obtained using a "list-db" request. The <db-name> is taken from the "name" member of <database-schema>.

<table>

An <id> that names a table.

<column>

An <id> that names a table column.

<row>

A JSON object that describes a table row or a subset of a table row. Each member is the name of a table column paired with the <value> of that column.

<value>

A JSON value that represents the value of a column in a table row, one of <atom>, a <set>, or a <map>.

<atom>

A JSON value that represents a scalar value for a column, one of <string>, <number>, <boolean>, <uuid>, <named-uuid>.

<set>

Either an <atom>, representing a set with exactly one element, or a 2-element JSON array that represents a database set value. The first element of the array must be the string "set" and the second element must be an array of zero or more <atom>s giving the values in the set. All of the <atom>s must have the same type.

<map>

A 2-element JSON array that represents a database map value. The first element of the array must be the string "map" and the second element must be an array of zero or more <pair>s giving the values in the map. All of the <pair>s must have the same key and value types.

(JSON objects are not used to represent <map> because JSON only allows string names in an object.)

<pair>

A 2-element JSON array that represents a pair within a database map. The first element is an <atom> that represents the key, the second element is an <atom> that represents the value.

<uuid>

A 2-element JSON array that represents a UUID. The first element of the array must be the string "uuid" and the second element must be a 36-character string giving the UUID in the format described by [RFC 4122](#). For example, the following <uuid> represents the UUID 550e8400-e29b-41d4-a716-446655440000:

```
["uuid", "550e8400-e29b-41d4-a716-446655440000"]
```

<named-uuid>

A 2-element JSON array that represents the UUID of a row inserted in an "insert" operation within the same transaction. The first element of the array must be the string "named-uuid" and the second element should be the <id> specified as the "uuid-name" for

an "insert" operation within the same transaction. For example, if an "insert" operation within this transaction specifies a "uuid-name" of "myrow", the following <named-uuid> represents the UUID created by that operation:

```
["named-uuid", "myrow"]
```

A <named-uuid> may be used anywhere a <uuid> is valid.

<condition>

A 3-element JSON array of the form [<column>, <function>, <value>] that represents a test on a column value. Except as otherwise specified below, <value> MUST have the same type as <column>. The meaning depends on the type of <column>:

integer or real

<function> must be "<", "<=", "==", "!=", ">=", ">", "includes", or "excludes".

The test is true if the column's value satisfies the relation <function> <value>, e.g. if the column has value 1 and <value> is 2, the test is true if <function> is "<", "<=" or "!=", but not otherwise.

"includes" is equivalent to "=="; "excludes" is equivalent to "!=".

boolean or string or uuid

<function> must be "!=", "==", "includes", or "excludes".

If <function> is "==" or "includes", the test is true if the column's value equals <value>. If <function> is "!=" or "excludes", the test is inverted.

set or map

<function> must be "!=", "==", "includes", or "excludes".

If <function> is "==", the test is true if the column's value contains exactly the same values (for sets) or pairs (for maps). If <function> is "!=", the test is inverted.

If <function> is "includes", the test is true if the column's value contains all of the values (for sets) or pairs (for maps) in <value>. The column's value may also contain other values or pairs.

If <function> is "excludes", the test is true if the column's value does not contain any of the values (for sets) or pairs

(for maps) in <value>. The column's value may contain other values or pairs not in <value>.

If <function> is "includes" or "excludes", then the required type of <value> is slightly relaxed, in that it may have fewer than the minimum number of elements specified by the column's type. If <function> is "excludes", then the required type is additionally relaxed in that <value> may have more than the maximum number of elements specified by the column's type. <function> One of "<", "<=", "==", "!=", ">=", ">", "includes", "excludes".

<function>

One of "<", "<=", "==", "!=", ">=", ">", "includes", "excludes".

<mutation>

A 3-element JSON array of the form [<column>, <mutator>, <value>] that represents a change to a column value. Except as otherwise specified below, <value> must have the same type as <column>. The meaning depends on the type of <column>:

integer or real

<mutator> must be "+=", "-=", "*=", "/=" or (integer only) "%=". The value of <column> is changed to the sum, difference, product, quotient, or remainder, respectively, of <column> and <value>.

Constraints on <column> are ignored when parsing <value>.

boolean or string or uuid

No valid <mutator>s are currently defined for these types.

set

Any <mutator> valid for the set's element type may be applied to the set, in which case the mutation is applied to each member of the set individually. <value> must be a scalar value of the same type as the set's element type, except that constraints are ignored.

If <mutator> is "insert", then each of the values in the set in <value> is added to <column> if it is not already present. The required type of <value> is slightly relaxed, in that it may have fewer than the minimum number of elements specified by the column's type.

If <mutator> is "delete", then each of the values in the set in <value> is removed from <column> if it is present there. The required type is slightly relaxed in that <value> may have more

or less than the maximum number of elements specified by the column's type.

map

<mutator> must be "insert" or "delete".

If <mutator> is "insert", then each of the key-value pairs in the map in <value> is added to <column> only if its key is not already present. The required type of <value> is slightly relaxed, in that it may have fewer than the minimum number of elements specified by the column's type.

If <mutator> is "delete", then <value> may have the same type as <column> (a map type) or it may be a set whose element type is the same as <column>'s key type:

- + If <value> is a map, the mutation deletes each key-value pair in <column> whose key and value equal one of the key-value pairs in <value>.
- + If <value> is a set, the mutation deletes each key-value pair in <column> whose key equals one of the values in <value>.

For "delete", <value> may have any number of elements, regardless of restrictions on the number of elements in <column>.

<mutator>

One of "+=", "-=", "*=", "/=", "%=", "insert", "delete".

5.2. Operations

The operations that may be performed as part of a "transact" RPC request (see [Section 4.1.3](#)) are described in the following sections. Each of these operations is a JSON object that may be included as one of the elements of the "params" array that is one of the elements of the "transact" request. The details of each object, its semantics, results, and possible errors are described below.

5.2.1. Insert

The "insert" object contains the following members:

"op": "insert"	required
"table": <table>	required
"row": <row>	required
"uuid-name": <id>	optional

The corresponding result object contains the following member:

```
"uuid": <uuid>
```

The operation inserts "row" into "table". If "row" does not specify values for all the columns in "table", those columns receive default values. The default value for a column depends on its type. The default for a column whose <type> specifies a "min" of 0 is an empty set or empty map. Otherwise, the default is a single value or a single key-value pair, whose value(s) depend on its <atomic-type>:

- o "integer" or "real": 0
- o "boolean": false
- o "string": "" (the empty string)
- o "uuid": 00000000-0000-0000-0000-000000000000

The new row receives a new, randomly generated UUID. If "uuid-name" is supplied, then it is an error if <id> is not unique among the "uuid-name"s supplied on all the "insert" operations within this transaction. The UUID for the new row is returned as the "uuid" member of the result.

The errors that may be returned are as follows:

```
"error": "duplicate uuid-name"
```

The same "uuid-name" appears on another "insert" operation within this transaction.

```
"error": "constraint violation"
```

One of the values in "row" does not satisfy the immediate constraints for its column's <base-type>. This error will occur for columns that are not explicitly set by "row" if the default value does not satisfy the column's constraints.

5.2.2. Select

The "select" object contains the following members:

```
"op": "select"           required
"table": <table>         required
"where": [<condition>*]   required
"columns": [<column>*]    optional
```

The corresponding result object contains the following member:


```
"rows": [<row>*]
```

The operation searches "table" for rows that match all the conditions specified in "where". If "where" is an empty array, every row in "table" is selected.

The "rows" member of the result is an array of objects. Each object corresponds to a matching row, with each column specified in "columns" as a member, the column's name as the member name and its value as the member value. If "columns" is not specified, all the table's columns are included. If two rows of the result have the same values for all included columns, only one copy of that row is included in "rows". Specifying "_uuid" within "columns" will avoid dropping duplicates, since every row has a unique UUID.

The ordering of rows within "rows" is unspecified.

5.2.3. Update

The "update" object contains the following members:

```
"op": "update"           required
"table": <table>         required
"where": [<condition>*]   required
"row": <row>              required
```

The corresponding result object contains the following member:

```
"count": <integer>
```

The operation updates rows in a table. It searches "table" for rows that match all the conditions specified in "where". For each matching row, it changes the value of each column specified in "row" to the value for that column specified in "row". The "_uuid" and "_version" columns of a table may not be directly updated with this operation. Columns designated read-only in the schema also may not be updated.

The "count" member of the result specifies the number of rows that matched.

The error that may be returned is:

```
"error": "constraint violation"
  One of the values in "row" does not satisfy the immediate
  constraints for its column's <base-type>.
```


5.2.4. Mutate

The "mutate" object contains the following members:

```
"op": "mutate"           required
"table": <table>         required
"where": [<condition>*]   required
"mutations": [<mutation>*] required
```

The corresponding result object contains the following member:

```
"count": <integer>
```

The operation mutates rows in a table. It searches "table" for rows that match all the conditions specified in "where". For each matching row, it mutates its columns as specified by each <mutation> in "mutations", in the order specified.

The "_uuid" and "_version" columns of a table may not be directly modified with this operation. Columns designated read-only in the schema also may not be updated.

The "count" member of the result specifies the number of rows that matched.

The errors that may be returned are:

```
"error": "domain error"
```

The result of the mutation is not mathematically defined, e.g. division by zero.

```
"error": "range error"
```

The result of the mutation is not representable within the database's format, e.g. an integer result outside the range INT64_MIN...INT64_MAX or a real result outside the range -DBL_MAX...DBL_MAX.

```
"error": "constraint violation"
```

The mutation caused the column's value to violate a constraint, e.g. it caused a column to have more or fewer values than are allowed, an arithmetic operation caused a set or map to have duplicate elements, or it violated a constraint specified by a column's <base-type>.

5.2.5. Delete

The "delete" object contains the following members:


```
"op": "delete"           required
"table": <table>         required
"where": [<condition>*]   required
```

The corresponding result object contains the following member:

```
"count": <integer>
```

The operation deletes all the rows from "table" that match all the conditions specified in "where". The "count" member of the result specifies the number of deleted rows.

[5.2.6.](#) Wait

The "wait" object contains the following members:

```
"op": "wait"             required
"timeout": <integer>      optional
"table": <table>         required
"where": [<condition>*]   required
"columns": [<column>*]    required
"until": "==" or "!="     required
"rows": [<row>*]          required
```

There is no corresponding result object.

The operation waits until a condition becomes true.

If "until" is "==", it checks whether the query on "table" specified by "where" and "columns", which is evaluated in the same way as specified for "select", returns the result set specified by "rows". If it does, then the operation completes successfully. Otherwise, the entire transaction rolls back. It is automatically restarted later, after a change in the database makes it possible for the operation to succeed. The client will not receive a response until the operation permanently succeeds or fails.

If "until" is "!=", the sense of the test is negated. That is, as long as the query on "table" specified by "where" and "columns" returns "rows", the transaction will be rolled back and restarted later.

If "timeout" is specified, then the transaction aborts after the specified number of milliseconds. The transaction is guaranteed to be attempted at least once before it aborts. A "timeout" of 0 will abort the transaction on the first mismatch.

The errors that may be returned are:

"error": "not supported"

One or more of the columns in this table do not support triggers.
This error will not occur if "timeout" is 0.

"error": "timed out"

The "timeout" was reached before the transaction was able to complete.

[5.2.7.](#) **Commit**

The "commit" object contains the following members:

"op": "commit"	required
"durable": <boolean>	required

There is no corresponding result object.

If "durable" is specified as true, then the transaction, if it commits, will be stored durably (to disk) before the reply is sent to the client.

The error that may be returned is:

"error": "not supported"

When "durable" is true, this database implementation does not support durable commits.

[5.2.8.](#) **Abort**

The "abort" object contains the following member:

"op": "abort"	required
---------------	----------

There is no corresponding result object (the operation never succeeds).

The operation aborts the entire transaction with an error. This may be useful for testing.

The error that will be returned is:

"error": " aborted"

This operation always fails with this error.

[5.2.9.](#) **Comment**

The "comment" object contains the following members:

"op": "comment"	required
"comment": <string>	required

There is no corresponding result object.

The operation provides information to a database administrator on the purpose of a transaction. The OVSDB server, for example, adds comments in transactions that modify the database to the database journal.

5.2.10. Assert

The assert object contains the following members:

"op": "assert"	required
"lock": <id>	required

Result object has no members.

The assert operation causes the transaction to be aborted if the client does not own the lock named <string>.

The error that may be returned is:

```
"error": "not owner"
  The client does not own the named lock.
```

6. IANA Considerations

This document requests the assignment of a TCP port in the User Port range. The value that has been used in the past is 6632.

7. Security Considerations

The main security issue that needs to be addressed for the OVSDB protocol is the authentication, integrity, and privacy of communications between a client and server implementing this protocol. To provide such protection, an OVSDB connection SHOULD be secured using Transport Layer Security (TLS) [[RFC5246](#)]. The precise details of how clients and servers authenticate each other is highly dependent on the operating environment. It is often the case that OVSDB clients and servers operate in a tightly controlled environment, e.g., on machines in a single data center where they communicate on a isolated management network.

8. Acknowledgements

Thanks to Jeremy Stribling and Justin Pettit for their helpful input to this document.

9. References

9.1. Normative References

- [DCE] "DCE: Remote Procedure Call", Open Group CAE Specification C309, ISBN 1-85912-041-5, August 1994.
- [JSON-RPC] "JSON-RPC Specification, Version 1.0", <<http://json-rpc.org/wiki/specification>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4627] Crockford, D., "The application/json Media Type for JavaScript Object Notation (JSON)", [RFC 4627](#), July 2006.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), August 2008.

9.2. Informative References

- [DB-SCHEMA] "Open vSwitch Database Schema", <<http://openvswitch.org/ovs-vswitchd.conf.db.5.pdf>>.
- [OF-SPEC] "OpenFlow Switch Specification, version 1.3", <<https://www.opennetworking.org/images/stories/downloads/openflow/openflow-spec-v1.3.0.pdf>>.
- [OVS] "Open vSwitch", <<http://openvswitch.org/>>.

Authors' Addresses

Ben Pfaff
Nicira, Inc.
3460 W. Bayshore Rd.
Palo Alto, CA 94303
USA

Email: blp@nicira.com

Bruce Davie (editor)
Nicira, Inc.
3460 W. Bayshore Rd.
Palo Alto, CA 94303
USA

Email: bsd@nicira.com