

```

+----+-----+-----+-----+-----+-----+-----+-----+
| fid | ST_AsText (g) |
+----+-----+-----+-----+-----+-----+-----+-----+
| 1 | LINESTRING (30250.4 15129.2,30248.8 15138.4,30238.2 15136. ... |
| 2 | LINESTRING (30220.2 15122.8,30217.2 15137.8,30207.6 15136, ... |
| 3 | LINESTRING (30179 15114.4,30176.6 15129.4,30167 15128,3016 ... |
| 4 | LINESTRING (30155.2 15121.4,30140.4 15118.6,30142 15109,30 ... |
| 5 | LINESTRING (30192.4 15085,30177.6 15082.2,30179.2 15072.4, ... |
| 6 | LINESTRING (30244 15087,30229 15086.2,30229.4 15076.4,3024 ... |
| 7 | LINESTRING (30200.6 15059.4,30185.6 15058.6,30186 15048.8, ... |
| 10 | LINESTRING (30179.6 15017.8,30181 15002.8,30190.8 15003.6, ... |
| 11 | LINESTRING (30154.2 15000.4,30168.6 15004.8,30166 15014.2, ... |
| 13 | LINESTRING (30105 15065.8,30108.4 15050.8,30118 15053,3011 ... |
| 21 | LINESTRING (30350.4 15828.8,30350.6 15845,30333.8 15845,30 ... |
| 22 | LINESTRING (30350.6 15871.4,30350.6 15887.8,30334 15887.8, ... |
| 23 | LINESTRING (30350.6 15914.2,30350.6 15930.4,30334 15930.4, ... |
| 24 | LINESTRING (30290.2 15823,30290.2 15839.4,30273.4 15839.4, ... |
| 25 | LINESTRING (30291.4 15866.2,30291.6 15882.4,30274.8 15882. ... |
| 26 | LINESTRING (30291.6 15918.2,30291.6 15934.4,30275 15934.4, ... |
| 154 | LINESTRING (30276.2 15143.8,30261.4 15141,30263 15131.4,30 ... |
| 155 | LINESTRING (30269.8 15084,30269.4 15093.4,30258.6 15093,30 ... |
| 157 | LINESTRING (30128.2 15011,30113.2 15010.2,30113.6 15000.4, ... |
| 249 | LINESTRING (30337.8 15938.6,30337.8 15946.8,30320.4 15946. ... |
+----+-----+-----+-----+-----+-----+-----+-----+
20 rows in set (0.46 sec)

```

12.6 The JSON Data Type

As of MySQL 5.7.8, MySQL supports a native [JSON](#) data type that enables efficient access to data in JSON (JavaScript Object Notation) documents. The [JSON](#) data type provides these advantages over storing JSON-format strings in a string column:

- Automatic validation of JSON documents stored in [JSON](#) columns. Invalid documents produce an error.
- Optimized storage format. JSON documents stored in [JSON](#) columns are converted to an internal format that permits quick read access to document elements. When the server later must read a JSON value stored in this binary format, the value need not be parsed from a text representation. The binary format is structured to enable the server to look up subobjects or nested values directly by key or array index without reading all values before or after them in the document.



Note

This discussion uses [JSON](#) in monospace to indicate specifically the JSON data type and “JSON” in regular font to indicate JSON data in general.

The size of JSON documents stored in [JSON](#) columns is limited to the value of the [max_allowed_packet](#) system variable. (While the server manipulates a JSON value internally in memory, it can be larger; the limit applies when the server stores it.)

[JSON](#) columns cannot have a default value.

[JSON](#) columns, like columns of other binary types, are not indexed directly; instead, you can create an index on a generated column that extracts a scalar value from the [JSON](#) column. See [Section 14.1.18.6, “Secondary Indexes and Generated Virtual Columns”](#), for a detailed example.

The MySQL optimizer also looks for compatible indexes on virtual columns that match JSON expressions.

MySQL Cluster NDB 7.5.2 and later supports [JSON](#) columns and MySQL JSON functions, including creation of an index on a column generated from a [JSON](#) column as a workaround for being unable to index a [JSON](#) column. A maximum of 3 [JSON](#) columns per [NDB](#) table is supported.

The following discussion covers these topics:

- [Creating JSON Values](#)
- [Normalization, Merging, and Autowrapping of JSON Values](#)
- [Searching and Modifying JSON Values](#)
- [Comparison and Ordering of JSON Values](#)
- [Aggregation of JSON Values](#)

Along with the `JSON` data type, a set of SQL functions is available to enable operations on JSON values, such as creation, manipulation, and searching. The follow discussion shows examples of these operations. For details about individual functions, see [Section 13.16, “JSON Functions”](#).

A set of spatial functions for operating on GeoJSON values is also available. See [Section 13.15.11, “Spatial GeoJSON Functions”](#).

Creating JSON Values

A JSON array contains a list of values separated by commas and enclosed within `[` and `]` characters:

```
["abc", 10, null, true, false]
```

A JSON object contains a set of key/value pairs separated by commas and enclosed within `{` and `}` characters:

```
{"k1": "value", "k2": 10}
```

As the examples illustrate, JSON arrays and objects can contain scalar values that are strings or numbers, the JSON null literal, or the JSON boolean true or false literals. Keys in JSON objects must be strings. Temporal (date, time, or datetime) scalar values are also permitted:

```
["12:18:29.000000", "2015-07-29", "2015-07-29 12:18:29.000000"]
```

Nesting is permitted within JSON array elements and JSON object key values:

```
[99, {"id": "HK500", "cost": 75.99}, ["hot", "cold"]]
{"k1": "value", "k2": [10, 20]}
```

You can also obtain JSON values from a number of functions supplied by MySQL for this purpose (see [Section 13.16.2, “Functions That Create JSON Values”](#)) as well as by casting values of other types to the `JSON` type using `CAST(value AS JSON)` (see [Converting between JSON and non-JSON values](#)). The next several paragraphs describe how MySQL handles JSON values provided as input.

In MySQL, JSON values are written as strings. MySQL parses any string used in a context that requires a JSON value, and produces an error if it is not valid as JSON. These contexts include inserting a value into a column that has the `JSON` data type and passing an argument to a function that expects a JSON value, as the following examples demonstrate:

- Attempting to insert a value into a `JSON` column succeeds if the value is a valid JSON value, but fails if it is not:

```
mysql> CREATE TABLE t1 (jdoc JSON);
Query OK, 0 rows affected (0.20 sec)
```

```
mysql> INSERT INTO t1 VALUES('{\"key1\": \"value1\", \"key2\": \"value2\"}');
Query OK, 1 row affected (0.01 sec)

mysql> INSERT INTO t1 VALUES('[1, 2,]');
ERROR 3140 (22032) at line 2: Invalid JSON text: \"Invalid value.\" at position 6 in value (or column) '[1, 2,
```

Positions for “at position *N*” in such error messages are 0-based, but should be considered rough indications of where the problem in a value actually occurs.

- The `JSON_TYPE()` function expects a JSON argument and attempts to parse it into a JSON value. It returns the value's JSON type if it is valid and produces an error otherwise:

```
mysql> SELECT JSON_TYPE('[\"a\", \"b\", 1]');
+-----+
| JSON_TYPE('[\"a\", \"b\", 1]') |
+-----+
| ARRAY |
+-----+

mysql> SELECT JSON_TYPE('hello');
+-----+
| JSON_TYPE('hello') |
+-----+
| STRING |
+-----+

mysql> SELECT JSON_TYPE('hello');
ERROR 3146 (22032): Invalid data type for JSON data in argument 1
to function json_type; a JSON string or JSON type is required.
```

MySQL handles strings used in JSON context using the `utf8mb4` character set and `utf8mb4_bin` collation. Strings in other character sets are converted to `utf8mb4` as necessary. (For strings in the `ascii` or `utf8` character sets, no conversion is needed because `ascii` and `utf8` are subsets of `utf8mb4`.)

As an alternative to writing JSON values using literal strings, functions exist for composing JSON values from component elements. `JSON_ARRAY()` takes a (possibly empty) list of values and returns a JSON array containing those values:

```
mysql> SELECT JSON_ARRAY('a', 1, NOW());
+-----+
| JSON_ARRAY('a', 1, NOW()) |
+-----+
| [\"a\", 1, \"2015-07-27 09:43:47.000000\"] |
+-----+
```

`JSON_OBJECT()` takes a (possibly empty) list of key/value pairs and returns a JSON object containing those pairs:

```
mysql> SELECT JSON_OBJECT('key1', 1, 'key2', 'abc');
+-----+
| JSON_OBJECT('key1', 1, 'key2', 'abc') |
+-----+
| {\"key1\": 1, \"key2\": \"abc\"} |
+-----+
```

`JSON_MERGE()` takes two or more JSON documents and returns the combined result:

```
mysql> SELECT JSON_MERGE('[\"a\", 1]', '{\"key\": \"value\"}');
+-----+
| JSON_MERGE('[\"a\", 1]', '{\"key\": \"value\"}') |
+-----+
```

```

+-----+
| ["a", 1, {"key": "value"}] |
+-----+

```

For information about the merging rules, see [Normalization, Merging, and Autowrapping of JSON Values](#).

JSON values can be assigned to user-defined variables:

```

mysql> SET @j = JSON_OBJECT('key', 'value');
mysql> SELECT @j;
+-----+
| @j |
+-----+
| {"key": "value"} |
+-----+

```

However, user-defined variables cannot be of `JSON` data type, so although `@j` in the preceding example looks like a JSON value and has the same character set and collation as a JSON value, it does *not* have the `JSON` data type. Instead, the result from `JSON_OBJECT()` is converted to a string when assigned to the variable.

Strings produced by converting JSON values have a character set of `utf8mb4` and a collation of `utf8mb4_bin`:

```

mysql> SELECT CHARSET(@j), COLLATION(@j);
+-----+-----+
| CHARSET(@j) | COLLATION(@j) |
+-----+-----+
| utf8mb4 | utf8mb4_bin |
+-----+-----+

```

Because `utf8mb4_bin` is a binary collation, comparison of JSON values is case sensitive.

```

mysql> SELECT JSON_ARRAY('x') = JSON_ARRAY('X');
+-----+
| JSON_ARRAY('x') = JSON_ARRAY('X') |
+-----+
| 0 |
+-----+

```

Case sensitivity also applies to the JSON `null`, `true`, and `false` literals, which always must be written in lowercase:

```

mysql> SELECT JSON_VALID('null'), JSON_VALID('Null'), JSON_VALID('NULL');
+-----+-----+-----+
| JSON_VALID('null') | JSON_VALID('Null') | JSON_VALID('NULL') |
+-----+-----+-----+
| 1 | 0 | 0 |
+-----+-----+-----+

mysql> SELECT CAST('null' AS JSON);
+-----+
| CAST('null' AS JSON) |
+-----+
| null |
+-----+
1 row in set (0.00 sec)

mysql> SELECT CAST('NULL' AS JSON);
ERROR 3141 (22032): Invalid JSON text in argument 1 to function cast_as_json:

```

```
"Invalid value." at position 0 in 'NULL'.
```

Case sensitivity of the JSON literals differs from that of the SQL `NULL`, `TRUE`, and `FALSE` literals, which can be written in any lettercase:

```
mysql> SELECT ISNULL(null), ISNULL(Null), ISNULL(NULL);
+-----+-----+-----+
| ISNULL(null) | ISNULL(Null) | ISNULL(NULL) |
+-----+-----+-----+
|          1 |          1 |          1 |
+-----+-----+-----+
```

Normalization, Merging, and Autowrapping of JSON Values

When a string is parsed and found to be a valid JSON document, it is also normalized: Members with keys that duplicate a key found earlier in the document are discarded (even if the values differ). The object value produced by the following `JSON_OBJECT()` call does not include the second `key1` element because that key name occurs earlier in the value:

```
mysql> SELECT JSON_OBJECT('key1', 1, 'key2', 'abc', 'key1', 'def');
+-----+-----+
| JSON_OBJECT('key1', 1, 'key2', 'abc', 'key1', 'def') |
+-----+-----+
| {"key1": 1, "key2": "abc"} |
+-----+-----+
```

The normalization performed by MySQL also sorts the keys of a JSON object (for the purpose of making lookups more efficient). The result of this ordering is subject to change and not guaranteed to be consistent across releases. In addition, extra whitespace between keys, values, or elements in the original document is discarded.

MySQL functions that produce JSON values (see [Section 13.16.2, “Functions That Create JSON Values”](#)) always return normalized values.

In contexts that combine multiple arrays, the arrays are merged into a single array by concatenating arrays named later to the end of the first array. In the following example, `JSON_MERGE()` merges its arguments into a single array:

```
mysql> SELECT JSON_MERGE('[1, 2]', '["a", "b"]', '[true, false]');
+-----+-----+
| JSON_MERGE('[1, 2]', '["a", "b"]', '[true, false]') |
+-----+-----+
| [1, 2, "a", "b", true, false] |
+-----+-----+
```

Multiple objects when merged produce a single object. If multiple objects have the same key, the value for that key in the resulting merged object is an array containing the key values:

```
mysql> SELECT JSON_MERGE('{ "a": 1, "b": 2 }', '{ "c": 3, "a": 4 }');
+-----+-----+
| JSON_MERGE('{ "a": 1, "b": 2 }', '{ "c": 3, "a": 4 }') |
+-----+-----+
| {"a": [1, 4], "b": 2, "c": 3} |
+-----+-----+
```

Nonarray values used in a context that requires an array value are autowrapped: The value is surrounded by `[` and `]` characters to convert it to an array. In the following statement, each argument is autowrapped as an array (`[1]`, `[2]`). These are then merged to produce a single result array:

```
mysql> SELECT JSON_MERGE('1', '2');
+-----+
| JSON_MERGE('1', '2') |
+-----+
| [1, 2]                |
+-----+
```

Array and object values are merged by autowrapping the object as an array and merging the two arrays:

```
mysql> SELECT JSON_MERGE('[10, 20]', '{"a": "x", "b": "y"}');
+-----+
| JSON_MERGE('[10, 20]', '{"a": "x", "b": "y"}') |
+-----+
| [10, 20, {"a": "x", "b": "y"}]                |
+-----+
```

Searching and Modifying JSON Values

A JSON path expression selects a value within a JSON document.

Path expressions are useful with functions that extract parts of or modify a JSON document, to specify where within that document to operate. For example, the following query extracts from a JSON document the value of the member with the `name` key:

```
mysql> SELECT JSON_EXTRACT('{ "id": 14, "name": "Aztalan" }', '$.name');
+-----+
| JSON_EXTRACT('{ "id": 14, "name": "Aztalan" }', '$.name') |
+-----+
| "Aztalan"                                                  |
+-----+
```

Path syntax uses a leading `$` character to represent the JSON document under consideration, optionally followed by selectors that indicate successively more specific parts of the document:

- A period followed by a key name names the member in an object with the given key. The key name must be specified within double quotation marks if the name without quotes is not legal within path expressions (for example, if it contains a space).
- `[N]` appended to a path that selects an array names the value at position `N` within the array. Array positions are integers beginning with zero.
- Paths can contain `*` or `**` wildcards:
 - `.[*]` evaluates to the values of all members in a JSON object.
 - `[*]` evaluates to the values of all elements in a JSON array.
 - `prefix**suffix` evaluates to all paths that begin with the named prefix and end with the named suffix.
- A path that does not exist in the document (evaluates to nonexistent data) evaluates to `NULL`.

Let `$` refer to this JSON array with three elements:

```
[3, {"a": [5, 6], "b": 10}, [99, 100]]
```

Then:

- `$.[0]` evaluates to `3`.
- `$.[1]` evaluates to `{"a": [5, 6], "b": 10}`.
- `$.[2]` evaluates to `[99, 100]`.
- `$.[3]` evaluates to `NULL` (it refers to the fourth array element, which does not exist).

Because `$.[1]` and `$.[2]` evaluate to nonscalar values, they can be used as the basis for more-specific path expressions that select nested values. Examples:

- `$.[1].a` evaluates to `[5, 6]`.
- `$.[1].a[1]` evaluates to `6`.
- `$.[1].b` evaluates to `10`.
- `$.[2][0]` evaluates to `99`.

As mentioned previously, path components that name keys must be quoted if the unquoted key name is not legal in path expressions. Let `$.` refer to this value:

```
{"a fish": "shark", "a bird": "sparrow"}
```

The keys both contain a space and must be quoted:

- `$. "a fish"` evaluates to `shark`.
- `$. "a bird"` evaluates to `sparrow`.

Paths that use wildcards evaluate to an array that can contain multiple values:

```
mysql> SELECT JSON_EXTRACT('{ "a": 1, "b": 2, "c": [3, 4, 5] }', '$.*');
+-----+
| JSON_EXTRACT('{ "a": 1, "b": 2, "c": [3, 4, 5] }', '$.*') |
+-----+
| [1, 2, [3, 4, 5]] |
+-----+
mysql> SELECT JSON_EXTRACT('{ "a": 1, "b": 2, "c": [3, 4, 5] }', '$.c[*]');
+-----+
| JSON_EXTRACT('{ "a": 1, "b": 2, "c": [3, 4, 5] }', '$.c[*]') |
+-----+
| [3, 4, 5] |
+-----+
```

In the following example, the path `$.**.b` evaluates to multiple paths (`$.a.b` and `$.c.b`) and produces an array of the matching path values:

```
mysql> SELECT JSON_EXTRACT('{ "a": { "b": 1 }, "c": { "b": 2 } }', '$**.b');
+-----+
| JSON_EXTRACT('{ "a": { "b": 1 }, "c": { "b": 2 } }', '$**.b') |
+-----+
| [1, 2] |
+-----+
```

In MySQL 5.7.9 and later, you can use `column->path` with a JSON column identifier and JSON path expression as a synonym for `JSON_EXTRACT(column, path)`. See [Section 13.16.3, “Functions That Search JSON Values”](#), for more information. See also [Section 14.1.18.6, “Secondary Indexes and Generated Virtual Columns”](#).

Some functions take an existing JSON document, modify it in some way, and return the resulting modified document. Path expressions indicate where in the document to make changes. For example, the `JSON_SET()`, `JSON_INSERT()`, and `JSON_REPLACE()` functions each take a JSON document, plus one or more path/value pairs that describe where to modify the document and the values to use. The functions differ in how they handle existing and nonexisting values within the document.

Consider this document:

```
mysql> SET @j = '{"a", {"b": [true, false]}, [10, 20]}';
```

`JSON_SET()` replaces values for paths that exist and adds values for paths that do not exist:

```
mysql> SELECT JSON_SET(@j, '$[1].b[0]', 1, '$[2][2]', 2);
+-----+
| JSON_SET(@j, '$[1].b[0]', 1, '$[2][2]', 2) |
+-----+
| [{"a", {"b": [1, false]}, [10, 20, 2]}] |
+-----+
```

In this case, the path `[$[1].b[0]` selects an existing value (`true`), which is replaced with the value following the path argument (`1`). The path `[$[2][2]` does not exist, so the corresponding value (`2`) is added to the value selected by `[$[2]`.

`JSON_INSERT()` adds new values but does not replace existing values:

```
mysql> SELECT JSON_INSERT(@j, '$[1].b[0]', 1, '$[2][2]', 2);
+-----+
| JSON_INSERT(@j, '$[1].b[0]', 1, '$[2][2]', 2) |
+-----+
| [{"a", {"b": [true, false]}, [10, 20, 2]}] |
+-----+
```

`JSON_REPLACE()` replaces existing values and ignores new values:

```
mysql> SELECT JSON_REPLACE(@j, '$[1].b[0]', 1, '$[2][2]', 2);
+-----+
| JSON_REPLACE(@j, '$[1].b[0]', 1, '$[2][2]', 2) |
+-----+
| [{"a", {"b": [1, false]}, [10, 20]}] |
+-----+
```

The path/value pairs are evaluated left to right. The document produced by evaluating one pair becomes the new value against which the next pair is evaluated.

`JSON_REMOVE()` takes a JSON document and one or more paths that specify values to be removed from the document. The return value is the original document minus the values selected by paths that exist within the document:

```
mysql> SELECT JSON_REMOVE(@j, '$[2]', '$[1].b[1]', '$[1].b[1]');
+-----+
| JSON_REMOVE(@j, '$[2]', '$[1].b[1]', '$[1].b[1]') |
+-----+
| [{"a", {"b": [true]}]} |
+-----+
```

The paths have these effects:

- `[$[2]` matches `[10, 20]` and removes it.

- The first instance of `$.b[1]` matches `false` in the `b` element and removes it.
- The second instance of `$.b[1]` matches nothing: That element has already been removed, the path no longer exists, and has no effect.

Comparison and Ordering of JSON Values

JSON values can be compared using the `=`, `<`, `<=`, `>`, `>=`, `<>`, `!=`, and `<=>` operators.

The following comparison operators and functions are not yet supported with JSON values:

- `BETWEEN`
- `IN()`
- `GREATEST()`
- `LEAST()`

A workaround for the comparison operators and functions just listed is to cast JSON values to a native MySQL numeric or string data type so they have a consistent non-JSON scalar type.

Comparison of JSON values takes place at two levels. The first level of comparison is based on the JSON types of the compared values. If the types differ, the comparison result is determined solely by which type has higher precedence. If the two values have the same JSON type, a second level of comparison occurs using type-specific rules.

The following list shows the precedences of JSON types, from highest precedence to the lowest. (The type names are those returned by the `JSON_TYPE()` function.) Types shown together on a line have the same precedence. Any value having a JSON type listed earlier in the list compares greater than any value having a JSON type listed later in the list.

```
BLOB
BIT
OPAQUE
DATETIME
TIME
DATE
BOOLEAN
ARRAY
OBJECT
STRING
INTEGER, DOUBLE
NULL
```

For JSON values of the same precedence, the comparison rules are type specific:

- `BLOB`

The first `N` bytes of the two values are compared, where `N` is the number of bytes in the shorter value. If the first `N` bytes of the two values are identical, the shorter value is ordered before the longer value.

- `BIT`

Same rules as for `BLOB`.

- `OPAQUE`

Same rules as for `BLOB`. `OPAQUE` values are values that are not classified as one of the other types.

- [DATETIME](#)

A value that represents an earlier point in time is ordered before a value that represents a later point in time. If two values originally come from the MySQL [DATETIME](#) and [TIMESTAMP](#) types, respectively, they are equal if they represent the same point in time.

- [TIME](#)

The smaller of two time values is ordered before the larger one.

- [DATE](#)

The earlier date is ordered before the more recent date.

- [ARRAY](#)

Two JSON arrays are equal if they have the same length and values in corresponding positions in the arrays are equal.

If the arrays are not equal, their order is determined by the elements in the first position where there is a difference. The array with the smaller value in that position is ordered first. If all values of the shorter array are equal to the corresponding values in the longer array, the shorter array is ordered first.

Example:

```
[ ] < [ "a" ] < [ "ab" ] < [ "ab", "cd", "ef" ] < [ "ab", "ef" ]
```

- [BOOLEAN](#)

The JSON false literal is less than the JSON true literal.

- [OBJECT](#)

Two JSON objects are equal if they have the same set of keys, and each key has the same value in both objects.

Example:

```
{ "a": 1, "b": 2 } = { "b": 2, "a": 1 }
```

The order of two objects that are not equal is unspecified but deterministic.

- [STRING](#)

Strings are ordered lexically on the first *N* bytes of the [utf8mb4](#) representation of the two strings being compared, where *N* is the length of the shorter string. If the first *N* bytes of the two strings are identical, the shorter string is considered smaller than the longer string.

Example:

```
"a" < "ab" < "b" < "bc"
```

This ordering is equivalent to the ordering of SQL strings with collation [utf8mb4_bin](#). Because [utf8mb4_bin](#) is a binary collation, comparison of JSON values is case sensitive:

```
"A" < "a"
```

- [INTEGER](#), [DOUBLE](#)

JSON values can contain exact-value numbers and approximate-value numbers. For a general discussion of these types of numbers, see [Section 10.1.2, “Number Literals”](#).

The rules for comparing native MySQL numeric types are discussed in [Section 13.2, “Type Conversion in Expression Evaluation”](#), but the rules for comparing numbers within JSON values differ somewhat:

- In a comparison between two columns that use the native MySQL [INT](#) and [DOUBLE](#) numeric types, respectively, it is known that all comparisons involve an integer and a double, so the integer is converted to double for all rows. That is, exact-value numbers are converted to approximate-value numbers.
- On the other hand, if the query compares two JSON columns containing numbers, it cannot be known in advance whether numbers will be integer or double. To provide the most consistent behavior across all rows, MySQL converts approximate-value numbers to exact-value numbers. The resulting ordering is consistent and does not lose precision for the exact-value numbers. For example, given the scalars 9223372036854775805, 9223372036854775806, 9223372036854775807 and 9.223372036854776e18, the order is such as this:

```
9223372036854775805 < 9223372036854775806 < 9223372036854775807
< 9.223372036854776e18 = 9223372036854776000 < 9223372036854776001
```

Were JSON comparisons to use the non-JSON numeric comparison rules, inconsistent ordering could occur. The usual MySQL comparison rules for numbers yield these orderings:

- Integer comparison:

```
9223372036854775805 < 9223372036854775806 < 9223372036854775807
```

(not defined for 9.223372036854776e18)

- Double comparison:

```
9223372036854775805 = 9223372036854775806 = 9223372036854775807 = 9.223372036854776e18
```

For comparison of any JSON value to SQL [NULL](#), the result is [UNKNOWN](#).

For comparison of JSON and non-JSON values, the non-JSON value is converted to JSON according to the rules in the following table, then the values compared as described previously.

Converting between JSON and non-JSON values. The following table provides a summary of the rules that MySQL follows when casting between JSON values and values of other types:

Table 12.1 JSON Conversion Rules

other type	CAST(other type AS JSON)	CAST(JSON AS other type)
JSON	No change	No change
utf8 character type (utf8mb4 , utf8 , ascii)	The string is parsed into a JSON value.	The JSON value is serialized into a utf8mb4 string.
Other character types	Other character encodings are implicitly converted to utf8mb4 and treated as described for utf8 character type.	The JSON value is serialized into a utf8mb4 string, then cast to the other

other type	CAST(other type AS JSON)	CAST(JSON AS other type)
		character encoding. The result may not be meaningful.
NULL	Results in a NULL value of type JSON.	Not applicable.
Geometry types	The geometry value is converted into a JSON document by calling ST_AsGeoJSON() .	Illegal operation. Workaround: Pass the result of CAST(json_val AS CHAR) to ST_GeomFromGeoJSON() .
All other types	Results in a JSON document consisting of a single scalar value.	Succeeds if the JSON document consists of a single scalar value of the target type and that scalar value can be cast to the target type. Otherwise, returns NULL and produces a warning.

[ORDER BY](#) and [GROUP BY](#) for JSON values works according to these principles:

- Ordering of scalar JSON values uses the same rules as in the preceding discussion.
- For ascending sorts, SQL [NULL](#) orders before all JSON values, including the JSON null literal; for descending sorts, SQL [NULL](#) orders after all JSON values, including the JSON null literal.
- Sort keys for JSON values are bound by the value of the [max_sort_length](#) system variable, so keys that differ only after the first [max_sort_length](#) bytes compare as equal.
- Sorting of nonscalar values is not currently supported and a warning occurs.

For sorting, it can be beneficial to cast a JSON scalar to some other native MySQL type. For example, if a column named `jdoc` contains JSON objects having a member consisting of an `id` key and a nonnegative value, use this expression to sort by `id` values:

```
ORDER BY CAST(JSON_EXTRACT(jdoc, '$.id') AS UNSIGNED)
```

If there happens to be a generated column defined to use the same expression as in the [ORDER BY](#), the MySQL optimizer recognizes that and considers using the index for the query execution plan. See [Section 9.3.9, “Optimizer Use of Generated Column Indexes”](#).

Aggregation of JSON Values

For aggregation of JSON values, SQL [NULL](#) values are ignored as for other data types. Non-[NULL](#) values are converted to a numeric type and aggregated, except for [MIN\(\)](#), [MAX\(\)](#), and [GROUP_CONCAT\(\)](#). The conversion to number should produce a meaningful result for JSON values that are numeric scalars, although (depending on the values) truncation and loss of precision may occur. Conversion to number of other JSON values may not produce a meaningful result.

12.7 Data Type Default Values

The [DEFAULT value](#) clause in a data type specification indicates a default value for a column. With one exception, the default value must be a constant; it cannot be a function or an expression. This means, for example, that you cannot set the default for a date column to be the value of a function such as [NOW\(\)](#) or [CURRENT_DATE](#). The exception is that you can specify [CURRENT_TIMESTAMP](#) as the default for [TIMESTAMP](#) and [DATETIME](#) columns. See [Section 12.3.5, “Automatic Initialization and Updating for TIMESTAMP and DATETIME”](#).

[BLOB](#), [TEXT](#), [GEOMETRY](#), and [JSON](#) columns cannot be assigned a default value.