

---

**pg\_hint\_plan**

**NIPPON TELEGRAPH AND TELEPHONE CORPORATION**

**Jun 15, 2026**



# CONTENTS

<b>1 Synopsis</b>	<b>3</b>
<b>2 Description</b>	<b>5</b>
2.1 Basic Usage . . . . .	5
<b>3 The hint table</b>	<b>7</b>
3.1 Types of hints . . . . .	7
3.2 GUC parameters for <code>pg_hint_plan</code> . . . . .	10
<b>4 Installation</b>	<b>11</b>
4.1 Building binary module . . . . .	11
4.2 Installing from a binary package . . . . .	11
4.3 Loading <code>pg_hint_plan</code> . . . . .	11
<b>5 Uninstallation</b>	<b>13</b>
<b>6 Details in hinting</b>	<b>15</b>
6.1 Syntax and placement . . . . .	15
6.2 Using with PL/pgSQL . . . . .	15
6.3 Upper and lower case handling in object names . . . . .	16
6.4 Escaping special characters in object names . . . . .	16
6.5 Distinction between multiple occurrences of a table . . . . .	16
6.6 Underlying tables of views or rules . . . . .	17
6.7 Inheritance . . . . .	17
6.8 Hints in multistatements . . . . .	17
6.9 VALUES expressions . . . . .	17
6.10 Subqueries . . . . .	18
6.11 Using <code>IndexOnlyScan</code> hint . . . . .	18
6.12 About <code>NoIndexScan</code> . . . . .	18
6.13 Parallel hints and <code>UNION</code> . . . . .	18
6.14 Setting <code>pg_hint_plan</code> parameters by Set hints . . . . .	19
<b>7 Errors</b>	<b>21</b>
7.1 Syntax errors . . . . .	21
7.2 Incorrect Object definitions . . . . .	21
7.3 Redundant or conflicting hints . . . . .	21
7.4 Nested comments . . . . .	21
<b>8 Functional limitations</b>	<b>23</b>
8.1 Influence of planner GUC parameters . . . . .	23
8.2 Hints trying to enforce non-executable plans . . . . .	23

8.3	Queries in ECPG . . . . .	23
8.4	Query Identifiers . . . . .	23
<b>9</b>	<b>Requirements</b>	<b>25</b>
9.1	See also . . . . .	25
9.2	References . . . . .	25
<b>10</b>	<b>Hint list</b>	<b>27</b>

**Table of Contents**

- *pg\_hint\_plan 1.4*



## SYNOPSIS

`pg_hint_plan` makes it possible to tweak PostgreSQL execution plans using “hints” in SQL comments, as of `/*+ SeqScan(a) */`.

PostgreSQL uses a cost-based optimizer, which utilizes data statistics, not static rules. The planner (optimizer) estimates costs of each possible execution plans for a SQL statement then the execution plan with the lowest cost is executed. The planner does its best to select the best execution plan, but is not always perfect, since it doesn't take into account some of the data properties or correlations between columns.



## DESCRIPTION

### 2.1 Basic Usage

`pg_hint_plan` reads hinting phrases in a comment of special form given a SQL statement. A hint can be specified by prefixing it with the sequence `"/*+ "` and ending it with `"*/"`. Hint phrases consist of hint names and parameters enclosed by parentheses and delimited by whitespaces. Hint phrases can use newlines for readability.

In the example below, a hash join is selected as the join method while doing a sequential scan on `pgbench_accounts`:

```
=# /*+
    HashJoin(a b)
    SeqScan(a)
*/
EXPLAIN SELECT *
  FROM pgbench_branches b
  JOIN pgbench_accounts a ON b.bid = a.bid
  ORDER BY a.aid;
                                QUERY PLAN
-----
Sort (cost=31465.84..31715.84 rows=100000 width=197)
Sort Key: a.aid
-> Hash Join (cost=1.02..4016.02 rows=100000 width=197)
    Hash Cond: (a.bid = b.bid)
    -> Seq Scan on pgbench_accounts a (cost=0.00..2640.00 rows=100000
    width=97)
    -> Hash (cost=1.01..1.01 rows=1 width=100)
        -> Seq Scan on pgbench_branches b (cost=0.00..1.01 rows=1 width=100)
(7 rows)
```



## THE HINT TABLE

Hints can be specified in a comment, still this can be inconvenient in the case where queries cannot be edited. In the case, hints can be placed in a special table named "hint\_plan.hints". The table consists of the following columns:

column	description
id	Unique number to identify a row for a hint. This column is filled automatically by sequence.
norm_query_string	A string matching with the query to be hinted. Constants in the query are replaced by '?' as in the following example.
application_name	The value of application_name where sessions can apply a hint. The hint in the example below applies to sessions connected from psql. An empty string implies that all sessions will apply the hint.
hints	Hint phrase. This must be a series of hints excluding surrounding comment marks.

The following example shows how to operate with the hint table.

```
=# INSERT INTO hint_plan.hints(norm_query_string, application_name, hints)
  VALUES (
    'EXPLAIN (COSTS false) SELECT * FROM t1 WHERE t1.id = ?;',
    '',
    'SeqScan(t1)');
INSERT 0 1
=# UPDATE hint_plan.hints
  SET hints = 'IndexScan(t1)'
  WHERE id = 1;
UPDATE 1
=# DELETE FROM hint_plan.hints WHERE id = 1;
DELETE 1
```

The hint table is owned by the extension owner and has the same default privileges as of the time of its creation, during CREATE EXTENSION. Hints in the hint table are prioritized over hints in comments.

### 3.1 Types of hints

Hinting phrases are classified in multiple types based on what kind of object and how they can affect the planner. See *Hint list* for more details.

### 3.1.1 Hints for Scan methods

Scan method hints enforce specific scanning methods on the target table. `pg_hint_plan` recognizes the target table by alias names if any. These are for example `SeqScan` or `IndexScan`.

Scan hints work on ordinary tables, inheritance tables, UNLOGGED tables, temporary tables and system catalogs. External (foreign) tables, table functions, VALUES clause, CTEs, views and subqueries are not affected.

```

=# /*+
    SeqScan(t1)
    IndexScan(t2 t2_pkey)
*/
SELECT * FROM table1 t1 JOIN table2 t2 ON (t1.key = t2.key);

```

### 3.1.2 Hints for Join methods

Join method hints enforce the join methods of the joins involving the specified tables.

This can affect joins only on ordinary tables. Inheritance tables, UNLOGGED tables, temporary tables, external (foreign) tables, system catalogs, table functions, VALUES command results and CTEs are allowed to be in the parameter list. Joins on views and subqueries are not affected.

### 3.1.3 Hints for Joining order

This hint, named “Leading”, enforces the order of join on two or more tables. There are two methods of enforcing it. The first method enforces a specific order of joining but does not restrict the direction at each join level. The second method enforces the join direction additionally. See *hint list* for more details. For example:

```

=# /*+
    NestLoop(t1 t2)
    MergeJoin(t1 t2 t3)
    Leading(t1 t2 t3)
*/
SELECT * FROM table1 t1
JOIN table2 t2 ON (t1.key = t2.key)
JOIN table3 t3 ON (t2.key = t3.key);

```

### 3.1.4 Hints for Row number corrections

This hint, named “Rows”, changes the row number estimation of joins that comes from restrictions in the planner. For example:

```

=# /*+ Rows(a b #10) */ SELECT... ; Sets rows of join result to 10
=# /*+ Rows(a b +10) */ SELECT... ; Increments row number by 10
=# /*+ Rows(a b -10) */ SELECT... ; Subtracts 10 from the row number.
=# /*+ Rows(a b *10) */ SELECT... ; Makes the number 10 times larger.

```

### 3.1.5 Hints for parallel plans

This hint, named `Parallel`, enforces parallel execution configuration on scans. The third parameter specifies the strength of the enforcement. `soft` means that `pg_hint_plan` only changes `max_parallel_worker_per_gather` and leaves all the others to the planner to set. `hard` changes other planner parameters so as to forcibly apply the update. This can affect ordinary tables, inheritance parents, unlogged tables and system catalogs. External tables, table functions, `VALUES` clauses, CTEs, views and subqueries are not affected. Internal tables of a view can be specified by its real name or its alias as the target object. The following example shows that the query is enforced differently on each table:

```

=# EXPLAIN /*+ Parallel(c1 3 hard) Parallel(c2 5 hard) */
  SELECT c2.a FROM c1 JOIN c2 ON (c1.a = c2.a);
          QUERY PLAN
-----
Hash Join (cost=2.86..11406.38 rows=101 width=4)
  Hash Cond: (c1.a = c2.a)
    -> Gather (cost=0.00..7652.13 rows=1000101 width=4)
        Workers Planned: 3
            -> Parallel Seq Scan on c1 (cost=0.00..7652.13 rows=322613 width=4)
    -> Hash (cost=1.59..1.59 rows=101 width=4)
        -> Gather (cost=0.00..1.59 rows=101 width=4)
            Workers Planned: 5
                -> Parallel Seq Scan on c2 (cost=0.00..1.59 rows=59 width=4)

=# EXPLAIN /*+ Parallel(t1 5 hard) */ SELECT sum(a) FROM t1;
          QUERY PLAN
-----
Finalize Aggregate (cost=693.02..693.03 rows=1 width=8)
  -> Gather (cost=693.00..693.01 rows=5 width=8)
      Workers Planned: 5
          -> Partial Aggregate (cost=693.00..693.01 rows=1 width=8)
              -> Parallel Seq Scan on t1 (cost=0.00..643.00 rows=20000 width=4)

```

### 3.1.6 GUC parameters set during planning

Set hints change GUC parameters just while planning. GUC parameter shown in [Query Planning](#) can have the expected effects on planning unless an other hint conflicts with the planner method configuration parameters. When multiple hints change the same GUC, the last hint takes effect. *GUC parameters for `pg_hint_plan`* are also settable by this hint but it may not work as expected. See *Functional limitations* for details.

```

=# /*+ Set(random_page_cost 2.0) */
  SELECT * FROM table1 t1 WHERE key = 'value';
...

```

## 3.2 GUC parameters for pg\_hint\_plan

The following GUC parameters affect the behavior of pg\_hint\_plan:

Parameter name	Description	De- fault
pg_hint_plan. enable_hint	True enables pg_hint_plan.	on
pg_hint_plan. enable_hint_table	True enables hinting by table.	off
pg_hint_plan. parse_messages	Specifies the log level of hint parse error. Valid values are error, warning, notice, info, log, debug.	INFO
pg_hint_plan. debug_print	Controls debug print and verbosity. Valid values are off, on, detailed and verbose.	off
pg_hint_plan. message_level	Specifies message level of debug print. Valid values are error, warning, notice, info, log, debug.	INFO

## INSTALLATION

This section describes the installation steps.

### 4.1 Building binary module

Simply run `make` at the top of the source tree, then `make install` as an appropriate user. The `PATH` environment variable should be set properly to point to a PostgreSQL set of binaries:

```
$ tar xzvf pg_hint_plan-1.x.x.tar.gz
$ cd pg_hint_plan-1.x.x
$ make
$ su
$ make install
```

### 4.2 Installing from a binary package

On Debian and Ubuntu `pg_hint_plan` is available as a binary package from the `pgdg` (PostgreSQL Global Development Group) repository. Assuming you've already added the repository to `apt` sources, installing the package is as simple as:

```
sudo apt install postgresql-<postgres version>-pg-hint-plan
```

Please visit <https://www.postgresql.org/download/linux/> if you need help at adding the repository.

### 4.3 Loading `pg_hint_plan`

`pg_hint_plan` does not require `CREATE EXTENSION`. Loading it with a `LOAD` command will activate it and of course you can load it globally by setting `shared_preload_libraries` in `postgresql.conf`. Or you might be interested in `ALTER USER SET/ALTER DATABASE SET` for automatic loading in specific sessions.

```
postgres=# LOAD 'pg_hint_plan';
LOAD
```

Run `CREATE EXTENSION` and `SET pg_hint_plan.enable_hint_table TO on` if you are planning to use the hint table.



## UNINSTALLATION

`make uninstall` in the top directory of source tree will uninstall the installed files if you installed from the source tree and it is left available. Setting the environment variable `PATH` may be necessary.

```
$ cd pg_hint_plan-1.x.x
$ su
$ make uninstall
```



## DETAILS IN HINTING

### 6.1 Syntax and placement

`pg_hint_plan` reads hints from only the first block comment and stops parsing from any characters except alphabetical characters, digits, spaces, underscores, commas and parentheses. In the following example, `HashJoin(a b)` and `SeqScan(a)` are parsed as hints, but `IndexScan(a)` and `MergeJoin(a b)` are not:

```

=# /*+
    HashJoin(a b)
    SeqScan(a)
 */
/*+ IndexScan(a) */
EXPLAIN SELECT /*+ MergeJoin(a b) */ *
  FROM pgbench_branches b
  JOIN pgbench_accounts a ON b.bid = a.bid
  ORDER BY a.aid;
                                QUERY PLAN
-----
Sort (cost=31465.84..31715.84 rows=100000 width=197)
  Sort Key: a.aid
    -> Hash Join (cost=1.02..4016.02 rows=100000 width=197)
          Hash Cond: (a.bid = b.bid)
            -> Seq Scan on pgbench_accounts a (cost=0.00..2640.00 rows=100000 width=97)
            -> Hash (cost=1.01..1.01 rows=1 width=100)
                  -> Seq Scan on pgbench_branches b (cost=0.00..1.01 rows=1 width=100)
(7 rows)

```

### 6.2 Using with PL/pgSQL

`pg_hint_plan` works for queries in PL/pgSQL scripts with some restrictions.

- Hints affect only on the following kind of queries:
  - Queries that return one row (SELECT, INSERT, UPDATE and DELETE)
  - Queries that return multiple rows (RETURN QUERY)
  - Dynamic SQL statements (EXECUTE)
  - Cursor open (OPEN)
  - Loop over result of a query (FOR)

- A hint comment has to be placed after the first word in a query as preceding comments are not sent as a part of this query.

```

=# CREATE FUNCTION hints_func(integer) RETURNS integer AS $$
DECLARE
    id integer;
    cnt integer;
BEGIN
    SELECT /*+ NoIndexScan(a) */ aid
        INTO id FROM pgbench_accounts a WHERE aid = $1;
    SELECT /*+ SeqScan(a) */ count(*)
        INTO cnt FROM pgbench_accounts a;
    RETURN id + cnt;
END;
$$ LANGUAGE plpgsql;

```

### 6.3 Upper and lower case handling in object names

Unlike the way PostgreSQL handles object names, `pg_hint_plan` compares bare object names in hints against the database internal object names in a case-sensitive manner. Therefore, an object name `TBL` in a hint matches only “`TBL`” in the database and does not match any unquoted names like `TBL`, `tbl` or `Tbl`.

### 6.4 Escaping special characters in object names

The objects defined in a hint’s parameter can use double quotes if they include parentheses, double quotes and white spaces. The escaping rules are the same as PostgreSQL.

### 6.5 Distinction between multiple occurrences of a table

`pg_hint_plan` identifies the target object by using aliases if any. This behavior is useful to point to a specific occurrence among multiple occurrences of one table.

```

=# /*+ HashJoin(t1 t1) */
EXPLAIN SELECT * FROM s1.t1
JOIN public.t1 ON (s1.t1.id=public.t1.id);
INFO:  hint syntax error at or near "HashJoin(t1 t1)"
DETAIL:  Relation name "t1" is ambiguous.
...
=# /*+ HashJoin(pt st) */
EXPLAIN SELECT * FROM s1.t1 st
JOIN public.t1 pt ON (st.id=pt.id);
QUERY PLAN

-----
Hash Join (cost=64.00..1112.00 rows=28800 width=8)
Hash Cond: (st.id = pt.id)
-> Seq Scan on t1 st (cost=0.00..34.00 rows=2400 width=4)
-> Hash (cost=34.00..34.00 rows=2400 width=4)
    -> Seq Scan on t1 pt (cost=0.00..34.00 rows=2400 width=4)

```

## 6.6 Underlying tables of views or rules

Hints are not applicable on views, but they can affect the queries within the view if the object names match the names in the expanded query on the view. Assigning aliases to the tables in a view enables them to be manipulated from outside the view.

```

=# CREATE VIEW v1 AS SELECT * FROM t2;
=# EXPLAIN /*+ HashJoin(t1 v1) */
      SELECT * FROM t1 JOIN v1 ON (c1.a = v1.a);
      QUERY PLAN
-----
Hash Join (cost=3.27..18181.67 rows=101 width=8)
  Hash Cond: (t1.a = t2.a)
    -> Seq Scan on t1 (cost=0.00..14427.01 rows=1000101 width=4)
    -> Hash (cost=2.01..2.01 rows=101 width=4)
        -> Seq Scan on t2 (cost=0.00..2.01 rows=101 width=4)

```

## 6.7 Inheritance

Hints can only point to the parent of an inheritance tree and the hints affect all the tables in an inheritance tree. Hints pointing directly to inherited children have no effect.

## 6.8 Hints in multistatements

One multistatement can have exactly one hint comment and the hint affects all of the individual statements in the multistatement.

## 6.9 VALUES expressions

VALUES expressions in FROM clause are named as `*VALUES*` internally these can be hinted if it is the only `VALUES` of a query. Two or more `VALUES` expressions in a query cannot be distinguished by looking at an `EXPLAIN` result, resulting in ambiguous results:

```

=# /*+ MergeJoin(*VALUES*_1 *VALUES*) */
      EXPLAIN SELECT * FROM (VALUES (1, 1), (2, 2)) v (a, b)
      JOIN (VALUES (1, 5), (2, 8), (3, 4)) w (a, c) ON v.a = w.a;
INFO:  pg_hint_plan: hint syntax error at or near "MergeJoin(*VALUES*_1 *VALUES*) "
DETAIL:  Relation name "*VALUES*" is ambiguous.
      QUERY PLAN
-----
Hash Join (cost=0.05..0.12 rows=2 width=16)
  Hash Cond: ("*VALUES*_1".column1 = "*VALUES*".column1)
    -> Values Scan on "*VALUES*_1" (cost=0.00..0.04 rows=3 width=8)
    -> Hash (cost=0.03..0.03 rows=2 width=8)
        -> Values Scan on "*VALUES*" (cost=0.00..0.03 rows=2 width=8)

```

## 6.10 Subqueries

Subqueries context can be occasionally hinted using the name `ANY_subquery`:

```
IN (SELECT ... {LIMIT | OFFSET ...} ...)
= ANY (SELECT ... {LIMIT | OFFSET ...} ...)
= SOME (SELECT ... {LIMIT | OFFSET ...} ...)
```

For these syntaxes, the planner internally assigns the name to the subquery when planning joins on tables including it, so join hints are applicable on such joins using the implicit name. For example:

```
=# /*+HashJoin(a1 ANY_subquery)*/
EXPLAIN SELECT *
  FROM pgbench_accounts a1
 WHERE aid IN (SELECT bid FROM pgbench_accounts a2 LIMIT 10);
          QUERY PLAN
-----
Hash Semi Join (cost=0.49..2903.00 rows=1 width=97)
  Hash Cond: (a1.aid = a2.bid)
    -> Seq Scan on pgbench_accounts a1 (cost=0.00..2640.00 rows=100000 width=97)
    -> Hash (cost=0.36..0.36 rows=10 width=4)
          -> Limit (cost=0.00..0.26 rows=10 width=4)
                -> Seq Scan on pgbench_accounts a2 (cost=0.00..2640.00 rows=100000
↳width=4)
```

## 6.11 Using IndexOnlyScan hint

Index scan may be unexpectedly performed on another index when the index specified in `IndexOnlyScan` hint cannot perform an index only scan.

## 6.12 About NoIndexScan

A `NoIndexScan` hint implies `NoIndexOnlyScan`.

## 6.13 Parallel hints and UNION

A `UNION` can run in parallel only when all underlying subqueries are parallel-safe. Hence, enforcing parallel on any of the subqueries will let a parallel-executable `UNION` run in parallel. Meanwhile, a parallel hint with zero workers prevents a scan from being executed in parallel.

## 6.14 Setting pg\_hint\_plan parameters by Set hints

pg\_hint\_plan parameters influence their own behavior so some parameters will not work as one could expect:

- Hints to change `enable_hint`, `enable_hint_table` are ignored even though they are reported as “used hints” in debug logs.
- Setting `debug_print` and `message_level` in the middle of query processing.



## ERRORS

`pg_hint_plan` stops hint parsing on any error and will use the hints already parsed. Here are some typical errors.

### 7.1 Syntax errors

Any syntactical errors or wrong hint names are reported as a syntax error. These errors are reported in the server log with the message level specified by `pg_hint_plan.message_level` if `pg_hint_plan.debug_print` is on and above.

### 7.2 Incorrect Object definitions

Incorrect object definitions result in silently ignoring the hints. This kind of error is reported as a “Not Used Hint” in the server logs.

### 7.3 Redundant or conflicting hints

The last hint is considered when redundant hints are defined or hints conflict with each other. This kind of error is reported as a duplicated hints.

### 7.4 Nested comments

Hint comments cannot be recursive. If detected, hint parsing is immediately stopped and all the hints already parsed are ignored.



## FUNCTIONAL LIMITATIONS

### 8.1 Influence of planner GUC parameters

The planner does not try to consider joining order for FROM clause entries more than `from_collapse_limit`. `pg_hint_plan` cannot affect the joining order in this case.

### 8.2 Hints trying to enforce non-executable plans

Planner chooses any executable plans when the enforced plan cannot be executed:

- FULL OUTER JOIN to use nested loop.
- Use of indexes that do not have columns used in quals.
- TID scans for queries without ctid conditions.

### 8.3 Queries in ECPG

ECPG removes comments in queries written as embedded SQLs so hints cannot be passed to it. The only exception EXECUTE, that passes the query string to the server as-is. The hint table can be used in the case.

### 8.4 Query Identifiers

When `compute_query_id` is enabled, PostgreSQL generates a query ID, ignoring comments. Hence, queries with different hints, still written the same way, may compute the same query ID.



## REQUIREMENTS

pg\_hint\_plan 1.4 requires PostgreSQL 14.

PostgreSQL versions tested

- Version 14

OS versions tested

- CentOS 8.5

### 9.1 See also

### 9.2 References

- [EXPLAIN](#)
- [SET](#)
- [Server Config](#)
- [Parallel Plans](#)



**HINT LIST**

The available hints are listed below.

Group	Format	Description
Scan method	SeqScan(table)	Forces sequential scan on the table.
	TidScan(table)	Forces TID scan on the table.
	IndexScan(table[ index... ])	Forces index scan on the table. Restricts to specified indexes if any.
	IndexOnlyScan(table[ index... ])	Forces index-only scan on the table. Restricts to specified indexes if any. Index scan may be used if index-only scan is not available.
	BitmapScan(table[ index... ])	Forces bitmap scan on the table. Restricts to specified indexes if any.
	IndexScanRegex(table[ POSIX Regexp... ]) IndexOnlyScanRegex(table[ POSIX Regexp... ]) BitmapScanRegex(table[ POSIX Regexp... ])	Forces index scan, index-only scan (For PostgreSQL 9.2 and later) or bitmap scan on the table. Restricts to indexes that matches the specified POSIX regular expression pattern.
	NoSeqScan(table)	Forces to <i>not</i> do sequential scan on the table.
	NoTidScan(table)	Forces to <i>not</i> do TID scan on the table.
	NoIndexScan(table)	Forces to <i>not</i> do index scan and index-only scan on the table.
	NoIndexOnlyScan(table)	Forces to <i>not</i> do index only scan on the table.
	NoBitmapScan(table)	Forces to <i>not</i> do bitmap scan on the table.
Join method	NestLoop(table table[ table... ])	Forces nested loop for the joins on the tables specified.
	HashJoin(table table[ table... ])	Forces hash join for the joins on the tables specified.
	MergeJoin(table table[ table... ])	Forces merge join for the joins on the tables specified.
	NoNestLoop(table table[ table... ])	Forces to <i>not</i> do nested loop for the joins on the tables specified.
	NoHashJoin(table table[ table... ])	Forces to <i>not</i> do hash join for the joins on the tables specified.
	NoMergeJoin(table table[ table... ])	Forces to <i>not</i> do merge join for the joins on the tables specified.
Join order	Leading(table table[ table... ])	Forces join order as specified.
	Leading(<join pair>)	Forces join order and directions as specified. A join pair is a pair of tables and/or other join pairs enclosed by parentheses, which can make a nested structure.
Behavior control on Join	Memoize(table table[ table... ])	Allows the topmost join of a join among the specified tables to Memoize the inner result. Not enforced.
	NoMemoize(table table[ table... ])	Inhibits the topmost join of a join among the specified tables from Memoizing the inner result.
Row number correction	Rows(table table[ table... ] correction)	Corrects row number of a result of the joins on the tables specified. The available correction methods are absolute (#), addition (+), subtract(-) and multiplication (*). should be a string that strtod() can understand.