

MORTAR

PIG

Cheat Sheet

We love Apache Pig for data processing—it's easy to learn, it works with all kinds of data, and it plays well with Python, Java, and other popular languages. And, of course, Pig runs on Hadoop, so it's built for high-scale data science.

Whether you're just getting started with Pig or you've already written a variety of Pig scripts, this compact reference gathers in one place many of the tools you'll need to make the most of your data using Pig 0.12.

Contents

Data Types	2
Diagnostic Operators	
Relational Operators	3
Syntax Tips	4
How to Use UDFs	5
Mathematical Functions	6
Eval Functions	7
String Functions	8
Date/Time Functions	9
Bag and Tuple Functions	10
Load/Store Functions	11
SQL -> Pig	12

Additional Resources

Official Pig website

pig.apache.org

Programming Pig, by Alan Gates

[chimera.labs.oreilly.com/
books/1234000001811](http://chimera.labs.oreilly.com/books/1234000001811)

Mortar's Pig page

www.mortardata.com/pig-resources




Data Types

Pig is written in Java, so Pig data types correspond to underlying Java data types. When using a UDF (see page 5), the Java types need to be translated into the target language.

	TYPES	DESCRIPTION	EXAMPLE	PYTHON/JYTHON UDF TYPE
SIMPLE	int	Signed 32-bit integer	10	int
	long	Signed 64-bit integer	10L	long
	float	32-bit floating point	10.5f	float
	double	64-bit floating point	10.5	float
	chararray	Character array (string) in Unicode UTF-8 format	hello world	str or unicode
	bytearray	Byte array (blob)		bytearray
	boolean	boolean	true/false	bool
	datetime	org.joda.time.DateTime	1970-01-01T00:00:00.000+00:00	datetime
	bigint	Java BigInteger	100000000000	long
	bigdecimal	Java BigDecimal	42.4242424242424	float
COMPLEX	tuple	An ordered set of fields.	(17, random)	tuple
	bag	A collection of tuples.	{(17,random), (42,life)}	list
	map	A set of key value pairs.	[foo#bar,baz#quux]	dict

Diagnostic Operators

- DESCRIBE** Returns the schema of a relation.
- DUMP** Dumps or displays results to screen.
- EXPLAIN** Displays execution plans.
- ILLUSTRATE** Displays a step-by-step execution of a sequence of statements.



Relational Operators

These operators are the heart of Pig set operations. The fundamental ways to manipulate data appear here, including GROUP, FILTER, and JOIN.

NAME	DESCRIPTION
<code>COGROUP alias BY (col1, col2)</code> ↗	COGROUP is the same as GROUP. For readability, programmers usually use GROUP when only one relation is involved and COGROUP when multiple relations are involved.
<code>CROSS alias1, alias2</code> ↗	Computes the cross product of two or more relations.
<code>CUBE alias BY CUBE(exp1, exp2, ...)</code> ↗	Cube operation computes aggregates for all possible combinations of specified group by dimensions. The number of group by combinations generated by cube for n dimensions will be 2^n .
<code>CUBE alias BY ROLLUP(exp1, exp2, ...)</code> ↗	Rollup operation computes multiple levels of aggregates based on hierarchical ordering of specified group by dimensions. Rollup is useful when there is hierarchical ordering on the dimensions. The number of group by combinations generated by rollup for n dimensions will be n+1.
<code>DISTINCT alias</code> ↗	Removes duplicate tuples in a relation.
<code>FILTER alias BY expression</code> ↗	Selects tuples from a relation based on some condition.
<code>...FLATTEN(tuple)</code> <code>...FLATTEN(bag)</code> ↗	Un-nests a bag or a tuple.
<code>FOREACH ... GENERATE</code> ↗	Performs transformations on each row in the data.
<code>GROUP alias ALL</code> <code>GROUP alias BY expression</code> <code>GROUP alias BY (exp1, exp2, ...)</code> ↗	Groups the data in one or multiple relations.
<code>JOIN smaller_alias BY expression, larger_alias BY expression</code> ↗	Performs inner, equijoin of two or more relations based on common field values.
<code>JOIN smaller_alias BY expression [LEFT RIGHT OUTER], larger_alias BY expression</code> ↗	Performs an outer join of two or more relations based on common field values.
<code>JOIN big_alias BY expression, small_alias BY expression USING 'replicated'</code> ↗	Efficient join when one or more relations are small enough to fit in main memory.
<code>LIMIT alias n</code> ↗	Limits the number of output tuples.
<code>LOAD 'data' [USING function] [AS schema]</code> ↗	Loads data from the file system.



Relational Operators (cont...)

NAME	DESCRIPTION
<code>ORDER</code> alias <code>BY</code> col [<code>ASC</code> <code>DESC</code>] ↗	Sorts a relation based on one or more fields.
<code>RANK</code> alias <code>BY</code> col [<code>ASC</code> <code>DESC</code>] ↗	Returns each tuple with the rank within a relation.
<code>SAMPLE</code> alias size ↗	Selects a random sample of data based on the specified sample size.
<code>SPLIT</code> alias <code>INTO</code> alias1 <code>IF</code> expression, alias2 <code>IF</code> expression... ↗	Partitions a relation into two or more relations.
<code>STORE</code> alias <code>INTO</code> 'directory' [<code>USING</code> function] ↗	Stores or saves results to the file system.
<code>UNION</code> alias1, alias2 ↗	Computes the union of two or more relations.

Syntax Tips

The fields in a Pig relation have an explicit order (unlike SQL columns). As a result, there are syntax shortcuts that rely on that field order.

Assume:

```
my_data = LOAD 'file.tsv' AS (field1:int, field2:chararray, field3:float, field4:int);
```

STATEMENT	OUTPUT
<code>FOREACH</code> my_data <code>GENERATE</code> \$0, \$2;	(field1, field3)
<code>FOREACH</code> my_data <code>GENERATE</code> field2..field4;	(field2, field3, field4)
<code>FOREACH</code> my_data <code>GENERATE</code> field2..;	(field2, field3, field4)
<code>FOREACH</code> my_data <code>GENERATE</code> *;	(field1, field2, field3, field4)



How to Use UDFs

Much of Pig's power comes from the fact that it can be extended using other languages. User-defined functions (UDFs) can be written in Java, Python, Jython, Ruby, and javascript.

Python

PIG

```
REGISTER 'udfs.py' USING streaming_python AS udfs;
```

PYTHON

```
@outputSchema("user_products:bag{t:(product_id:long)}")
def deserialize_user_products(product_ids):
    return [ (product_id, ) for product_id in product_ids.split(',') ]
```

Jython

PIG

```
REGISTER 'udfs.py' USING jython AS udfs;
```

JYTHON

```
@outputSchema("user_products:bag{t:(product_id:long)}")
def deserialize_user_products(product_ids):
    return [ (product_id, ) for product_id in product_ids.split(',') ]
```

Java

PIG

```
REGISTER 'udf-project.jar';
DEFINE My_Function my.function.path.MyFunction();
```

JAVA

For an introduction to java UDFs and Loaders, see www.mortardata.com/java-pig

Download the [Mortar Pig Java Template](#) for a template project with working pom.xml files.

Ruby

PIG

```
REGISTER 'test.rb' USING jruby AS myfuncs;
```

RUBY

```
require 'pigudf'
class Myudfs < PigUdf
  outputSchema num:int
  def square num
    return nil if num.nil?
    num**2
  end
end
```

Mathematical Functions

These functions are very similar to what you can find in `java.lang.Math`—basic mathematical functions and the trigonometric repertoire.

NAME (SIGNATURE)	RETURN TYPE	DESCRIPTION
<code>ABS(int a)</code> , <code>ABS(long a)</code> , <code>ABS(float a)</code> , <code>ABS(double a)</code> ↗	int, long, float, double	Returns the absolute value of an expression.
<code>ACOS(double a)</code> ↗	double	Returns the arc cosine of an expression.
<code>ASIN(double a)</code> ↗	double	Returns the arc sine of an expression.
<code>ATAN(double a)</code> ↗	double	Returns the arc tangent of an expression.
<code>CBRT(double a)</code> ↗	double	Returns the cube root of an expression.
<code>CEIL(double a)</code> ↗	double	Returns the value of an expression rounded up to the nearest integer.
<code>COS(double a)</code> ↗	double	Returns the cosine of an expression.
<code>COSH(double a)</code> ↗	double	Returns the hyperbolic cosine of an expression.
<code>EXP(double a)</code> ↗	double	Returns Euler's number e raised to the power of x.
<code>FLOOR(double a)</code> ↗	double	Returns the value of an expression rounded down to the nearest integer.
<code>LOG(double a)</code> ↗	double	Returns the natural logarithm (base e) of an expression.
<code>LOG10(double a)</code> ↗	double	Returns the base 10 logarithm of an expression.
<code>RANDOM()</code> ↗	double	Returns a pseudo random number greater than or equal to 0.0 and less than 1.0.
<code>ROUND(float a)</code> , <code>ROUND(double a)</code> ↗	int, long	Returns the value of an expression rounded to an integer.
<code>SIN(double a)</code> ↗	double	Returns the sine of an expression.
<code>SINH(double a)</code> ↗	double	Returns the hyperbolic sine of an expression.
<code>SQRT(double a)</code> ↗	double	Returns the positive square root of an expression.
<code>TAN(double a)</code> ↗	double	Returns the tangent of an expression.
<code>TANH(double a)</code> ↗	double	Returns the hyperbolic tangent of an expression.



Eval Functions

This group contains aggregate functions such as COUNT and SUM, along with useful utility methods such as isEmpty.

NAME (SIGNATURE)	RETURN TYPE	DESCRIPTION
AVG(col) ↗	double	Computes the average of the numeric values in a single column of a bag.
CONCAT(String expression1, String expression2) CONCAT(byte[] expression1, byte[] expression2) ↗	String, byte[]	Concatenates two expressions of identical type.
COUNT(DataBag bag) ↗	long	Computes the number of elements in a bag. Does not include null values.
COUNT_STAR(DataBag bag) ↗	long	Computes the number of elements in a bag, including null values.
DIFF(DataBag bag1, DataBag bag2) ↗	DataBag	Compares two bags. Any tuples that are in one bag but not the other are returned in a bag.
IsEmpty(DataBag bag), IsEmpty(Map map) ↗	boolean	Checks if a bag or map is empty.
MAX(col) ↗	int, long, float, double	Computes the maximum of the numeric values or chararrays in a single-column bag.
MIN(col) ↗	int, long, float, double	Computes the minimum of the numeric values or chararrays in a single-column bag.
DEFINE pluck PluckTuple(expression1) pluck(expression2) ↗	Tuple	Allows the user to specify a string prefix, and then filter for the columns in a relation that begin with that prefix.
SIZE(expression) ↗	long	Computes the number of elements based on any Pig data type.
SUBTRACT(DataBag bag1, DataBag bag2) ↗	DataBag	Returns bag composed of bag1 elements not in bag2.
SUM(col) ↗	double, long	Computes the sum of the numeric values in a single-column bag.
TOKENIZE(String expression [, 'field_delimiter']) ↗	DataBag	Splits a string and outputs a bag of words. If field_delimiter is null or not passed, the following will be used as delimiters: space [], double quote ["], comma [,], parenthesis [()], star [*]

String Functions

Providing convenient ways to handle text fields, these functions mirror commonly used String functions in other languages.

NAME (SIGNATURE)	RETURN TYPE	DESCRIPTION
<code>ENDSWITH(String string, String testAgainst)</code> ↗	boolean	Tests inputs to determine if the first argument ends with the string in the second.
<code>EqualsIgnoreCase(String string1, String string2)</code> ↗	boolean	Compares two strings ignoring case considerations.
<code>INDEXOF(String string, String 'character', int startIndex)</code> ↗	int	Returns the index of the first occurrence of a character in a string, searching forward from a start index.
<code>LAST_INDEX_OF(String string, String 'character')</code> ↗	int	Returns the index of the last occurrence of a character in a string, searching backward from the end of the string.
<code>LCFIRST(String expression)</code> ↗	String	Converts the first character in a string to lower case.
<code>LOWER(String expression)</code> ↗	String	Converts all characters in a string to lower case.
<code>LTRIM(String expression)</code> ↗	String	Returns a copy of a string with only leading white space removed.
<code>REGEX_EXTRACT(String string, String regex, int index)</code> ↗	String	Performs regular expression matching and extracts the matched group defined by an index parameter.
<code>REGEX_EXTRACT_ALL (String string, String regex)</code> ↗	Tuple	Performs regular expression matching and extracts all matched groups.
<code>REPLACE(String string, String 'regExp', String 'newChar')</code> ↗	String	Replaces existing characters in a string with new characters.
<code>RTRIM(String expression)</code> ↗	String	Returns a copy of a string with only trailing white space removed.
<code>STARTSWITH(String string, String testAgainst)</code> ↗	boolean	Tests inputs to determine if the first argument starts with the string in the second.
<code>STRSPLIT(String string, String regex, int limit)</code> ↗	Tuple	Splits a string around matches of a given regular expression.
<code>SUBSTRING(String string, int startIndex, int stopIndex)</code> ↗	String	Returns a substring from a given string.
<code>TRIM(String expression)</code> ↗	String	Returns a copy of a string with leading and trailing white space removed.
<code>UCFIRST(String expression)</code> ↗	String	Returns a string with the first character converted to upper case.
<code>UPPER(String expression)</code> ↗	String	Returns a string converted to upper case.



DateTime Functions

Pig uses the Joda-Time DateTime class for date and time handling.

NAME (SIGNATURE)	RETURN TYPE	DESCRIPTION
<code>AddDuration(DateTime datetime, String duration)</code> ↗	DateTime	Returns the result of a DateTime object plus an ISO 8601 duration string.
<code>CurrentTime()</code> ↗	DateTime	Returns the DateTime object of the current time.
<code>DaysBetween(DateTime datetime1, DateTime datetime2)</code> ↗	long	Returns the number of days between two DateTime objects.
<code>GetDay(DateTime datetime)</code> ↗	int	Returns the day of a month from a DateTime object.
<code>GetHour(DateTime datetime)</code> ↗	int	Returns the hour of a day from a DateTime object.
<code>GetMilliSecond(DateTime datetime)</code> ↗	int	Returns the millisecond of a second from a DateTime object.
<code>GetMinute(DateTime datetime)</code> ↗	int	Returns the minute of an hour from a DateTime object.
<code>GetMonth(DateTime datetime)</code> ↗	int	Returns the month of a year from a DateTime object.
<code>GetSecond(DateTime datetime)</code> ↗	int	Returns the second of a minute from a DateTime object.
<code>GetWeek(DateTime datetime)</code> ↗	int	Returns the week of a week year from a DateTime object.
<code>GetWeekYear(DateTime datetime)</code> ↗	int	Returns the week year from a DateTime object.
<code>GetYear(DateTime datetime)</code> ↗	int	Returns the year from a DateTime object.
<code>HoursBetween(DateTime datetime1, DateTime datetime2)</code> ↗	long	Returns the number of hours between two DateTime objects.
<code>MillisecondsBetween(datetime1, datetime2)</code> ↗	long	Returns the number of milliseconds between two DateTime objects.
<code>MinutesBetween(DateTime datetime1, DateTime datetime2)</code> ↗	long	Returns the number of minutes between two DateTime objects.
<code>MonthsBetween(DateTime datetime1, DateTime datetime2)</code> ↗	long	Returns the number of months between two DateTime objects.
<code>SecondsBetween(DateTime datetime1, DateTime datetime2)</code> ↗	long	Returns the number of seconds between two DateTime objects.
<code>SubtractDuration(DateTime datetime, String duration)</code> ↗	DateTime	Returns the result of a DateTime object minus an ISO 8601 duration string.

DateTime Functions (cont...)

NAME (SIGNATURE)	RETURN TYPE	DESCRIPTION
<code>ToDate(long milliseconds)</code> ↗	<code>DateTime</code>	Returns a <code>DateTime</code> object according to parameters.
<code>ToDate(String isostring)</code> ↗	<code>DateTime</code>	Returns a <code>DateTime</code> object according to parameters.
<code>ToDate(String userstring, String format)</code> ↗	<code>DateTime</code>	Returns a <code>DateTime</code> object according to parameters.
<code>ToDate(String userstring, String format, String timezone)</code> ↗	<code>DateTime</code>	Returns a <code>DateTime</code> object according to parameters.
<code>ToMilliseconds(DateTime datetime)</code> ↗	<code>long</code>	Returns the number of milliseconds elapsed since January 1, 1970, 00:00:00.000 GMT for a <code>DateTime</code> object.
<code>ToString(DateTime datetime)</code> ↗	<code>String</code>	Converts the <code>DateTime</code> object to the ISO string.
<code>ToString(DateTime datetime, String format)</code> ↗	<code>String</code>	Converts the <code>DateTime</code> object to the customized string.
<code>ToUnixTime(DateTime datetime)</code> ↗	<code>long</code>	Returns the Unix Time as long for a <code>DateTime</code> object. Unix Time is the number of seconds elapsed since January 1, 1970, 00:00:00.000 GMT.
<code>WeeksBetween(DateTime datetime1, DateTime datetime2)</code> ↗	<code>long</code>	Returns the number of weeks between two <code>DateTime</code> objects.
<code>YearsBetween(DateTime datetime1, DateTime datetime2)</code> ↗	<code>long</code>	Returns the number of years between two <code>DateTime</code> objects.

Bag and Tuple Functions

One of the things that makes Pig powerful is its use of complex data types. These functions exist to manipulate data stored in bags, tuples, and maps.

NAME (SIGNATURE)	RETURN TYPE	DESCRIPTION
<code>TOTUPLE(expression [, expression ...])</code> ↗	<code>Tuple</code>	Converts one or more expressions to type tuple.
<code>TOBAG(expression [, expression ...])</code> ↗	<code>DataBag</code>	Converts one or more expressions to type bag.
<code>TOMAP(key-expression, value-expression [, key-expression, value-expression ...])</code> ↗	<code>Map</code>	Converts key/value expression pairs into a map.
<code>TOP(int topN, column, relation)</code> ↗	<code>DataBag</code>	Returns the top-n tuples from a bag of tuples.



Load/Store Functions

Crucial to data manipulation in Pig, load and store functions bring data into Pig and push it out again in a multitude of formats.

NAME	FILE FORMAT	LOAD/STORE	
<pre>org.apache.pig.piggybank.storage.avro.AvroStorage() org.apache.pig.piggybank.storage.avro.AvroStorage('no_schema_check', '\$SCHEMA_FILE', '\$PATH');</pre>	Avro	Load, Store	
<pre>org.apache.pig.piggybank.storage.apachelog.CommonLogLoader()</pre>	Common Log Format, Combined Log Format	Load	
<pre>org.apache.pig.piggybank.storage.CSVExcelStorage() org.apache.pig.piggybank.storage.CSVExcelStorage('\$DELIMITER', 'YES_MULTILINE', 'NOCHANGE', 'SKIP_INPUT_HEADER')</pre>	CSV	Load, Store	
<pre>com.mortardata.pig.storage.DynamoDBStorage('\$DYNAMODB_TABLE', '\$DYNAMODB_AWS_ACCESS_KEY_ID', '\$DYNAMODB_AWS_SECRET_ACCESS_KEY')</pre>	DynamoDB	Store	
<pre>see PostgreSQL</pre>	Greenplum	Store	
<pre>org.apache.pig.piggybank.storage.FixedWidthLoader('-5, 7-10, 10-18') org.apache.pig.piggybank.storage.FixedWidthLoader('-5, 7-10, 10-18', 'SKIP_HEADER', '\$SCHEMA')</pre>	Fixed-Width	Load, Store	
<pre>org.apache.pig.piggybank.storage.HiveColumnarLoader('\$SCHEMA')</pre>	Hive	Load	
<pre>org.apache.pig.piggybank.storage.JsonLoader() org.apache.pig.piggybank.storage.JsonLoader('\$SCHEMA')</pre>	JSON	Load, Store	
<pre>com.mongodb.hadoop.pig.MongoLoader() com.mongodb.hadoop.pig.MongoLoader('\$SCHEMA')</pre>	MongoDB	Load, Store	
<pre>com.mortardata.pig.PapertrailLoader()</pre>	Papertrail Logs	Load	
<pre>PigStorage() PigStorage('\$DELIMITER')</pre>	character-delimited (CSV, TSV)	Load, Store	
<pre>org.apache.pig.piggybank.storage.DBStorage('org.postgresql.Driver', 'jdbc:postgresql://\$POSTGRES_HOST/\$POSTGRES_DATABASE', '\$POSTGRES_USER', '\$POSTGRES_PASS', 'INSERT INTO my_table(my_col_1,my_col_2,my_col_3) VALUES (?, ?, ?)')</pre>	PostgreSQL	Store	
<pre>TextLoader()</pre>	Text/Unformatted	Load	
<pre>org.apache.pig.piggybank.storage.StreamingXMLLoader('\$DOCUMENT')</pre>	XML	Load	

SQL -> Pig

Because many people come to Pig from a relational database background, we've included a handy translation from SQL concepts to their Pig equivalents.

SQL FUNCTION	SQL	PIG
SELECT	<code>SELECT column_name, column_name FROM table_name;</code>	<code>FOREACH alias GENERATE column_name, column_name;</code>
SELECT *	<code>SELECT * FROM table_name;</code>	<code>FOREACH alias GENERATE *;</code>
DISTINCT	<code>SELECT DISTINCT column_name, column_name FROM table_name;</code>	<code>DISTINCT(FOREACH alias GENERATE column_name, column_name);</code>
WHERE	<code>SELECT column_name, column_name FROM table_name WHERE column_name operator value;</code>	<code>FOREACH (FILTER alias BY column_name operator value) GENERATE column_name, column_name;</code>
AND/OR	<code>... WHERE (column_name operator value1 AND column_name operator value2) OR column_name operator value3;</code>	<code>FILTER alias BY (column_name operator value1 AND column_name operator value2) OR column_name operator value3;</code>
ORDER BY	<code>... ORDER BY column_name ASC DESC, column_name ASC DESC;</code>	<code>ORDER alias BY column_name ASC DESC, column_name ASC DESC;</code>
TOP/LIMIT	<code>SELECT TOP number column_name FROM table_name ORDER BY column_name ASC DESC;</code> <code>SELECT column_name FROM table_name ORDER BY column_name ASC DESC LIMIT number;</code>	<code>FOREACH (GROUP alias BY column_name) GENERATE LIMIT alias number;</code> <code>TOP(number, column_index, alias);</code>
GROUP BY	<code>SELECT function(column_name) FROM table GROUP BY column_name;</code>	<code>FOREACH (GROUP alias BY column_name) GENERATE function(alias.column_name);</code>
LIKE	<code>... WHERE column_name LIKE pattern;</code>	<code>FILTER alias BY REGEX_EXTRACT(column_name, pattern, 1) IS NOT NULL;</code>
IN	<code>... WHERE column_name IN (value1,value2,...);</code>	<code>FILTER alias BY column_name IN (value1, value2,...);</code>
JOIN	<code>SELECT column_name(s) FROM table1 JOIN table2 ON table1.column_name=table2.column_name;</code>	<code>FOREACH (JOIN alias1 BY column_name, alias2 BY column_name) GENERATE column_name(s);</code>



SQL -> Pig (cont...)

SQL FUNCTION	SQL	PIG
LEFT/RIGHT/FULL OUTER JOIN	<pre>SELECT column_name(s) FROM table1 LEFT RIGHT FULL OUTER JOIN table2 ON table1.column_name=table2.column_name;</pre>	<pre>FOREACH (JOIN alias1 BY column_name LEFT RIGHT FULL, alias2 BY column_name) GENERATE column_name(s);</pre>
UNION ALL	<pre>SELECT column_name(s) FROM table1 UNION ALL SELECT column_name(s) FROM table2;</pre>	<pre>UNION alias1, alias2;</pre>
AVG	<pre>SELECT AVG(column_name) FROM table_name;</pre>	<pre>FOREACH (GROUP alias ALL) GENERATE AVG(alias.column_name);</pre>
COUNT	<pre>SELECT COUNT(column_name) FROM table_name;</pre>	<pre>FOREACH (GROUP alias ALL) GENERATE COUNT(alias);</pre>
COUNT DISTINCT	<pre>SELECT COUNT(DISTINCT column_name) FROM table_name;</pre>	<pre>FOREACH alias { unique_column = DISTINCT column_name; GENERATE COUNT(unique_column); };</pre>
MAX	<pre>SELECT MAX(column_name) FROM table_name;</pre>	<pre>FOREACH (GROUP alias ALL) GENERATE MAX(alias.column_name);</pre>
MIN	<pre>SELECT MIN(column_name) FROM table_name;</pre>	<pre>FOREACH (GROUP alias ALL) GENERATE MIN(alias.column_name);</pre>
SUM	<pre>SELECT SUM(column_name) FROM table_name;</pre>	<pre>FOREACH (GROUP alias ALL) GENERATE SUM(alias.column_name);</pre>
HAVING	<pre>... HAVING aggregate_function(column_name) operator value;</pre>	<pre>FILTER alias BY aggregate_function(column_name) operator value;</pre>
UCASE/UPPER	<pre>SELECT UCASE(column_name) FROM table_name;</pre>	<pre>FOREACH alias GENERATE UPPER(column_name);</pre>
LCASE/LOWER	<pre>SELECT LCASE(column_name) FROM table_name;</pre>	<pre>FOREACH alias GENERATE LOWER(column_name);</pre>
SUBSTRING	<pre>SELECT SUBSTRING(column_name,start,length) AS some_name FROM table_name;</pre>	<pre>FOREACH alias GENERATE SUBSTRING(column_name, start, start+length) as some_name;</pre>
LEN	<pre>SELECT LEN(column_name) FROM table_name;</pre>	<pre>FOREACH alias GENERATE SIZE(column_name);</pre>
ROUND	<pre>SELECT ROUND(column_name,0) FROM table_name;</pre>	<pre>FOREACH alias GENERATE ROUND(column_name);</pre>