

InfluxDB 2.0

Paul Dix
@pauldix
paul@influxdata.com

Biggest Change Since 0.9

Clean Migration Path

Compatibility Layer



- MIT Licensed
- Multi-tenanted
- Telegraf, InfluxDB, Chronograf, Kapacitor rolled into 1
- OSS single server



- MIT Licensed
- Multi-tenanted
- Telegraf, InfluxDB, Chronograf, Kapacitor rolled into 1
- OSS single server
- Cloud usage based pricing
- Dedicated Cloud
- Enterprise on-premise

TICK is dead

Long Live InfluxDB 2.0

(and Telegraf)



flux

Consistent Documented API

Collection, Write/Query, Streaming & Batch Processing, Dashboards

[Code](#)[Issues 361](#)[Pull requests 16](#)[Projects 1](#)[Wiki](#)[Insights](#)[Settings](#)

Branch: master ▾

[platform / http / swagger.yml](#)[Find file](#) [Copy path](#) **goller** Merge pull request #1275 from influxdata/feature/query-plan

3e54ef9 5 days ago

14 contributors 

4895 lines (4895 sloc) | 134 KB

[Raw](#)[Blame](#)[History](#)

```
1  openapi: "3.0.0"
2  info:
3    title: Influx API Service
4    version: 0.1.0
5  servers:
6    - url: /api/v2
7  paths:
8    /signin:
9      post:
10       summary: Exchange basic auth credentials for session
11       security:
12         - basicAuth: []
13     responses:
14       '204':
15         description: successfully authenticated
16       default:
17         description: unsuccessful authentication
18         content:
```

Officially Supported Client Libraries

Go, Node.js, Ruby, Python, PHP, Java, C#, C, Kotlin

Visualization Libraries

Multi-tenant roles

- Operator
- Organization Administrator
- User

Data Model

- Organizations
 - Buckets (retention)
 - Time series data
 - Tasks
 - Runs
 - Logs
 - Dashboards
- Users
 - Tokens
 - Authorizations
- Protos (templates)
- Scrapers
- Telegrafs
- Labels

All-in-one but separable

Demo

<https://influxdata.com/download>

Status

- Alpha 1 released 4 weeks ago
- New alpha build every week
- Alphas deliver features
- Beta once feature complete
- Beta releases for performance and stability

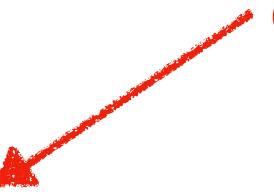
Thank you

Paul Dix
@pauldix
paul@influxdata.com

Flux Language Primer

```
// get all data from the telegraf db
from(bucket:"telegraf/autogen")
// filter that by the last hour
|> range(start:-1h)
// filter further by series with a specific measurement and field
|> filter(fn: (r) => r._measurement == "cpu" and r._field == "usage_system")
```

Comments



```
// get all data from the telegraf db
from(bucket:"telegraf/autogen")
  // filter that by the last hour
  |> range(start:-1h)
  // filter further by series with a specific measurement and field
  |> filter(fn: (r) => r._measurement == "cpu" and r._field == "usage_system")
```

Named Arguments

```
// get All data from the telegraf db
from(bucket:"telegraf/autogen")
  // filter that by the last hour
  |> range(start:-1h)
  // filter further by series with a specific measurement and field
  |> filter(fn: (r) => r._measurement == "cpu" and r._field == "usage_system")
```

String Literals

```
// get all data from the telegraf db
from(bucket:"telegraf/autogen")
  // filter that by the last hour
  |> range(start:-1h)
  // filter further by series with a specific measurement and field
  |> filter(fn: (r) => r._measurement == "cpu" and r._field == "usage_system")
```

Buckets, not DBs



```
// get data from the telegraf db
from(bucket:"telegraf/autogen")
// filter that by the last hour
|> range(start:-1h)
// filter further by series with a specific measurement and field
|> filter(fn: (r) => r._measurement == "cpu" and r._field == "usage_system")
```

```
// get all data from the telegraf db
from(bucket:"telegraf/autogen")
// filter that by the last hour
|> range(start:-1h) ← Duration Literal
// filter further by series with a specific measurement and field
|> filter(fn: (r) => r._measurement == "cpu" and r._field == "usage_system")
```

```
// get all data from the telegraf db
from(bucket:"telegraf/autogen")
// filter that by the last hour
|> range(start:2018-11-07T00:00:00Z)← Time Literal
// filter further by series with a specific measurement and field
|> filter(fn: (r) => r._measurement == "cpu" and r._field == "usage_system")
```

```
// get all data from the telegraf db
from(bucket:"telegraf/autogen")
  // filter that by the last hour
  |> range(start:-1h)
  // filter further by series with a specific measurement and field
  |> filter(fn: (r) => r._measurement == "cpu" and r._field == "usage_system")
```



Pipe forward operator

```
// get all data from the telegraf db
from(bucket:"telegraf/autogen")
// filter that by the last hour
|> range(start:-1h)
// filter further by series with a specific measurement and field
|> filter(fn: (r) => r._measurement == "cpu" and r._field == "usage_system")
```



Anonymous Function

```
// get all data from the telegraf db
from(bucket:"telegraf/autogen")
  // filter that by the last hour
  |> range(start:-1h)
  // filter further by series with a specific measurement and field
  |> filter(fn: (r) => (r._measurement == "cpu" or r._measurement == "cpu")
            and r.host == "serverA")
```



Predicate Function

```
// variables
some_int = 23
```

```
// variables
some_int = 23
some_float = 23.2
```

```
// variables
some_int = 23
some_float = 23.2
some_string = "cpu"
```

```
// variables
some_int = 23
some_float = 23.2
some_string = "cpu"
some_duration = 1h
```

```
// variables
some_int = 23
some_float = 23.2
some_string = "cpu"
some_duration = 1h
some_time = 2018-10-10T19:00:00
```

```
// variables
some_int = 23
some_float = 23.2
some_string = "cpu"
some_duration = 1h
some_time = 2018-10-10T19:00:00
some_array = [1, 6, 20, 22]
```

```
// variables
some_int = 23
some_float = 23.2
some_string = "cpu"
some_duration = 1h
some_time = 2018-10-10T19:00:00
some_array = [1, 6, 20, 22]
some_object = {foo: "hello" bar: 22}
```

Data Model & Working with Tables

Example Series

```
_measurement=mem,host=A,region=west,_field=free  
_measurement=mem,host=B,region=west,_field=free  
_measurement=cpu,host=A,region=west,_field=usage_system  
_measurement=cpu,host=A,region=west,_field=usage_user
```

Example Series

```
_measurement=mem,host=A,region=west,_field=free  
_measurement=mem,host=B,region=west,_field=free  
_measurement=cpu,host=A,region=west,_field=usage_system  
_measurement=cpu,host=A,region=west,_field=usage_user
```



Measurement

Example Series

```
_measurement=mem,host=A,region=west,_field=free  
_measurement=mem,host=B,region=west,_field=free  
_measurement=cpu,host=A,region=west,_field=usage_system  
_measurement=cpu,host=A,region=west,_field=usage_user
```



Field

<u>measurement</u>	<u>host</u>	<u>region</u>	<u>field</u>	<u>time</u>	<u>value</u>
mem	A	west	free	2018-06-14T09:15:00	10
mem	A	west	free	2018-06-14T09:14:50	10



Table

<u>measurement</u>	<u>host</u>	<u>region</u>	<u>field</u>	<u>time</u>	<u>value</u>
mem	A	west	free	2018-06-14T09:15:00	10
mem	A	west	free	2018-06-14T09:14:50	10



Column

<u>measurement</u>	<u>host</u>	<u>region</u>	<u>field</u>	<u>time</u>	<u>value</u>
mem	A	west	free	2018-06-14T09:15:00	10
mem	A	west	free	2018-06-14T09:14:50	10



Record

<code>_measurement</code>	<code>host</code>	<code>region</code>	<code>_field</code>	<code>_time</code>	<code>_value</code>
mem	A	west	free	2018-06-14T09:15:00	10
mem	A	west	free	2018-06-14T09:14:50	10

`_measurement=mem,host=A,region=west,_field=free`



Group Key

_measurement	host	region	_field	_time	_value
mem	A	west	free	2018-06-14T09:15:00	10
mem	A	west	free	2018-06-14T09:14:50	10



_measurement=mem,host=A,region=west,_field=free

**Every record has
the same value!**

Table Per Series

_measurement	host	region	_field	_time	_value
mem	A	west	free	2018-06-14T09:15:00	10
mem	A	west	free	2018-06-14T09:14:50	11
_measurement	host	region	_field	_time	_value
mem	B	west	free	2018-06-14T09:15:00	20
mem	B	west	free	2018-06-14T09:14:50	22
_measurement	host	region	_field	_time	_value
cpu	A	west	usage_user	2018-06-14T09:15:00	45
cpu	A	west	usage_user	2018-06-14T09:14:50	49
_measurement	host	region	_field	_time	_value
cpu	A	west	usage_system	2018-06-14T09:15:00	35
cpu	A	west	usage_system	2018-06-14T09:14:50	38

input tables -> function -> output tables

input tables -> function -> output tables

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:50, stop:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r._field == "free")
|> sum()
```

input tables -> function -> output tables

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:50, stop:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r._field == "free")
|> sum()
```



What to sum on?

input tables -> function -> output tables

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:50, stop:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r._field == "free")
|> sum(columns: ["_value"])
```



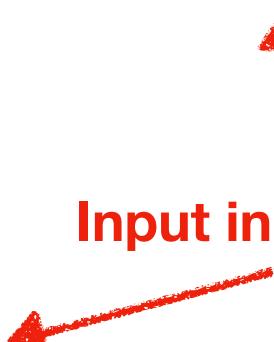
Default columns argument

input tables -> function -> output tables

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:50, stop:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r._field == "free")
|> sum()
```

_meas	host	region	_field	_time	_value
mem	A	west	free	2018-06-	10
mem	A	west	free	2018-06-	11
_meas	host	region	_field	_time	_value
mem	B	west	free	2018-06-	20
mem	B	west	free	2018-06-	22

Input in table form



input tables -> function -> output tables

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:50, stop:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r._field == "free")
|> sum()
```

_meas	host	region	_field	_time	_value
mem	A	west	free	2018-06-	10
mem	A	west	free	2018-06-	11
_meas	host	region	_field	_time	_value
mem	B	west	free	2018-06-	20
mem	B	west	free	2018-06-	22



sum()

input tables -> function -> output tables

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:50, stop:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r._field == "free")
|> sum()
```

_meas	host	region	_field	_time	_value
mem	A	west	free	2018-06-14T09:14:50	10
mem	A	west	free	2018-06-14T09:15:01	11

sum()

_meas	host	region	_field	_time	_value
mem	B	west	free	2018-06-14T09:14:50	20
mem	B	west	free	2018-06-14T09:15:01	22

_meas	host	region	_field	_time	_value
mem	A	west	free	2018-06-14T09:15:01	21

_meas	host	region	_field	_time	_value
mem	B	west	free	2018-06-14T09:15:01	42

N to N table mapping

(1 to 1 mapping)

N to M table mapping

window

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:30, end:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r._field == "free")
|> window(every:20s)
```



30s of data (4 samples)

window

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:30, end:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r._field == "free")
|> window(every:20s)
```



split into 20s windows

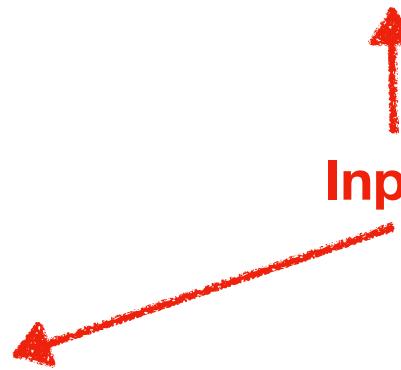
window

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:30, end:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r._field == "free")
|> window(every:20s)
```

meas	host	region	_field	_time	_valu
mem	A	west	free	...14:30	10
mem	A	west	free	...14:40	11
mem	A	west	free	...14:50	12
mem	A	west	free	...15:00	13

meas	host	region	field	time	valu
mem	B	west	free	...14:30	20
mem	B	west	free	...14:40	22
mem	B	west	free	...14:50	23
mem	B	west	free	...15:00	24

Input



window

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:30, end:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r._field == "free")
|> window(every:20s)
```

meas	host	region	_field	_time	_valu
mem	A	west	free	...14:30	10
mem	A	west	free	...14:40	11
mem	A	west	free	...14:50	12
mem	A	west	free	...15:00	13

meas	host	region	field	time	valu
mem	B	west	free	...14:30	20
mem	B	west	free	...14:40	22
mem	B	west	free	...14:50	23
mem	B	west	free	...15:00	24

→ **window(
every:20s)**

window

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:30, end:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
          r._field == "free")
|> window(every:20s)
```

meas	host	region	_field	_time	_valu
mem	A	west	free	...14:30	10
mem	A	west	free	...14:40	11
mem	A	west	free	...14:50	12
mem	A	west	free	...15:00	13

meas	host	region	field	time	valu
mem	B	west	free	...14:30	20
mem	B	west	free	...14:40	22
mem	B	west	free	...14:50	23
mem	B	west	free	...15:00	24

→ window(
every:20s) →

_meas	host	region	_field	_time	_valu
mem	A	west	free	...14:30	10
mem	A	west	free	...14:40	11

_meas	host	region	_field	_time	_valu
mem	A	west	free	...14:50	12
mem	A	west	free	...15:00	13

_meas	host	region	_field	_time	_valu
mem	B	west	free	...14:30	20
mem	B	west	free	...14:40	22

_meas	host	region	_field	_time	_valu
mem	B	west	free	...14:50	23
mem	B	west	free	...15:00	24

window

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:30, end:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
          r._field == "free")
|> window(every:20s)
```

N to M tables

meas	host	region	_field	_time	_value
mem	A	west	free	...14:30	10
mem	A	west	free	...14:40	11
mem	A	west	free	...14:50	12
mem	A	west	free	...15:00	13

meas	host	region	field	time	value
mem	B	west	free	...14:30	20
mem	B	west	free	...14:40	22
mem	B	west	free	...14:50	23
mem	B	west	free	...15:00	24

window(
every:20s)

meas	host	region	_field	_time	_value
mem	A	west	free	...14:30	10
mem	A	west	free	...14:40	11

meas	host	region	_field	_time	_value
mem	A	west	free	...14:50	12
mem	A	west	free	...15:00	13

meas	host	region	_field	_time	_value
mem	B	west	free	...14:30	20
mem	B	west	free	...14:40	22

meas	host	region	_field	_time	_value
mem	B	west	free	...14:50	23

Window based on time

_start and _stop columns

group

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:30, end:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
          r._field == "free")
|> group(keys:[ "region" ])
```

group

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:30, end:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r. field == "free")
|> group(keys:[ "region" ])
```



new group key

group

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:30, end:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
           r._field == "free")
|> group(keys:[ "region" ])
```

meas	host	region	_field	_time	valu
mem	A	west	free	...14:30	10
mem	A	west	free	...14:40	11
mem	A	west	free	...14:50	12
mem	A	west	free	...15:00	13

meas	host	region	field	time	valu
mem	B	west	free	...14:30	20
mem	B	west	free	...14:40	22
mem	B	west	free	...14:50	23
mem	B	west	free	...15:00	24

group

```
// example query
from(db:"telegraf")
|> range(start:2018-06-14T09:14:30, end:2018-06-14T09:15:01)
|> filter(fn: r => r._measurement == "mem" and
          r._field == "free")
|> group(keys:["region"])
```

N to M tables

M == cardinality(group keys)

meas	host	region	_field	_time	_value
mem	A	west	free	...14:30	10
mem	A	west	free	...14:40	11
mem	A	west	free	...14:50	12
mem	A	west	free	...15:00	13

meas	host	region	field	time	value
mem	B	west	free	...14:30	20
mem	B	west	free	...14:40	22
mem	B	west	free	...14:50	23
mem	B	west	free	...15:00	24

group(
keys:
["region"])

meas	host	region	_field	_time	_value
mem	A	west	free	...14:30	10
mem	B	west	free	...14:30	20
mem	A	west	free	...14:40	11
mem	B	west	free	...14:40	21
mem	A	west	free	...14:50	12
mem	B	west	free	...14:50	22
mem	B	west	free	...15:00	13
mem	B	west	free	...15:00	23

Group based on columns