

Selecting series

Select latest sample for series with a given metric name:

```
Select * from Kubernetes WHERE series=node_cpu_seconds_total
```

Select 5-minute range of samples for series with a given metric name:

```
Select * from Kubernetes WHERE series=node_cpu_seconds_total[5m]
```

Only series with given label values:

```
Select * from Kubernetes WHERE cpu="0" AND mode="idle" AND series=node_cpu_seconds_total
```

Complex label matchers:

```
Select * from Kubernetes WHERE cpu!="0" AND mode=~"user|system" AND series=node_cpu_seconds_total
```

=: Equality
!=: Non-equality
=~: Regex match
!~: Negative regex match

Math between series

Add all equally-labelled series from both sides:

```
Select * from Kubernetes where series=node_memory_MemFree_bytes + node_memory_Cached_bytes
```

Add series, matching only on the instance and job labels:

```
Select job, instance from Kubernetes where metricname=node_memory_MemFree_bytes AND
```

```
Select job, instance from Kubernetes where metricname=node_memory_Cached_bytes
```

Filtering series by value

Only keep series with a sample value greater than a given number:

```
Select from kubernetes where metricName=node_filesystem_avail_bytes > 101024*1024
```

Only keep series from the left-hand side whose sample values are larger than their right-hand-side matches:

```
Select * from kubernetes where go_goroutines > go_threads
```

Instead of filtering, return 0 or 1 for each compared series:

```
Select * from kubernetes where go_goroutines > bool go_threads
```

Match only on specific labels:

```
Select * from kubernetes where go_goroutines > bool go_threads group by (job, instance)
```

Available comparison operators: ==, !=, >, <, >=, <=

Set operations

Include any label sets that are either on the left or right side:

```
Select series, job from kubernetes where =up AND (job="prometheus" or job="node")
```

Match only on specific labels:

```
Select * from kubernetes where go_goroutines > bool go_threads group by (instance)
```

Aggregating over multiple series

Sum over all series:

```
Select * from Kubernetes where SUM by (node_filesystem_size_bytes)
```

Preserve the instance and job label dimensions:

```
Select * from Kubernetes where series=node_filesystem_size_bytes AMD SUM by sum by(job, instance)
```

Aggregate away the instance and job label dimensions:

```
Select * from Kubernetes where series=node_filesystem_size_bytes AMD SUM by sum without(job, instance)
```

Available aggregation operators: sum(), min(), max(), avg(), stddev(), stdvar(), count(), count_values(), group(), bottomk(), topk(), quantile()

