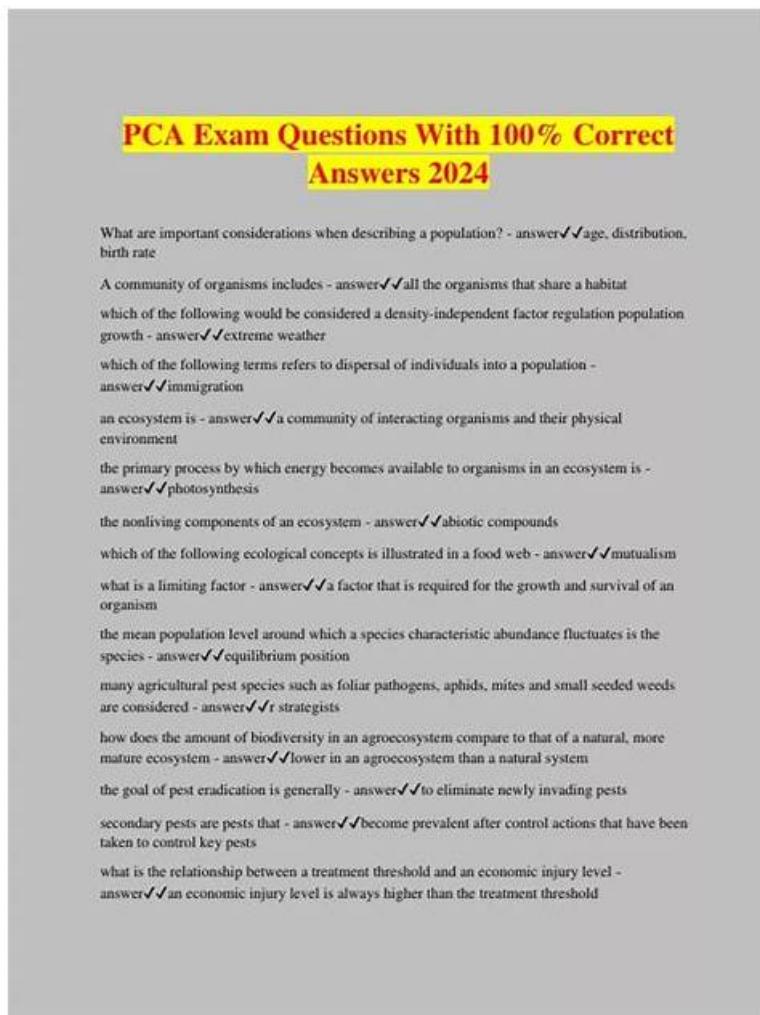


# PCA Valid Exam Camp Pdf - PCA Related Exams



What's more, part of that PracticeDump PCA dumps now are free: <https://drive.google.com/open?id=1WOUMn64q7bZ7PcEXUTg4NLiP68DGT2us>

Our PCA practice dumps compiled by the most professional experts can offer you with high quality and accuracy practice materials for your success. Up to now, we have more than tens of thousands of customers around the world supporting our PCA Exam Questions. If you are unfamiliar with our PCA study materials, please download the free demos for your reference, and to some unlearned exam candidates, you can master necessities by our PCA training guide quickly.

## Linux Foundation PCA Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"><li>Observability Concepts: This section of the exam measures the skills of Site Reliability Engineers and covers the essential principles of observability used in modern systems. It focuses on understanding metrics, logs, and tracing mechanisms such as spans, as well as the difference between push and pull data collection methods. Candidates also learn about service discovery processes and the fundamentals of defining and maintaining SLOs, SLAs, and SLIs to monitor performance and reliability.</li></ul>

Topic 2	<ul style="list-style-type: none"> <li>PromQL: This section of the exam measures the skills of Monitoring Specialists and focuses on Prometheus Query Language (PromQL) concepts. It covers data selection, calculating rates and derivatives, and performing aggregations across time and dimensions. Candidates also study the use of binary operators, histograms, and timestamp metrics to analyze monitoring data effectively, ensuring accurate interpretation of system performance and trends.</li> </ul>
Topic 3	<ul style="list-style-type: none"> <li>Prometheus Fundamentals: This domain evaluates the knowledge of DevOps Engineers and emphasizes the core architecture and components of Prometheus. It includes topics such as configuration and scraping techniques, limitations of the Prometheus system, data models and labels, and the exposition format used for data collection. The section ensures a solid grasp of how Prometheus functions as a monitoring and alerting toolkit within distributed environments.</li> </ul>
Topic 4	<ul style="list-style-type: none"> <li>Instrumentation and Exporters: This domain evaluates the abilities of Software Engineers and addresses the methods for integrating Prometheus into applications. It includes the use of client libraries, the process of instrumenting code, and the proper structuring and naming of metrics. The section also introduces exporters that allow Prometheus to collect metrics from various systems, ensuring efficient and standardized monitoring implementation.</li> </ul>
Topic 5	<ul style="list-style-type: none"> <li>Alerting and Dashboarding: This section of the exam assesses the competencies of Cloud Operations Engineers and focuses on monitoring visualization and alert management. It covers dashboarding basics, alerting rules configuration, and the use of Alertmanager to handle notifications. Candidates also learn the core principles of when, what, and why to trigger alerts, ensuring they can create reliable monitoring dashboards and proactive alerting systems to maintain system stability.</li> </ul>

**>> PCA Valid Exam Camp Pdf <<**

## PCA Related Exams - New PCA Test Test

Maybe this is the first time you choose our PCA practice materials, so it is understandable you may wander more useful information of our PCA exam dumps. Those free demos give you simple demonstration of our PCA study guide. It is unquestionable necessary for you to have an initial look of them before buying any. They are some brief introductions and basic information but also impressive. Just have a try and you will be interested in them!

## Linux Foundation Prometheus Certified Associate Exam Sample Questions (Q20-Q25):

### NEW QUESTION # 20

Which PromQL expression computes the rate of API Server requests across the different cloud providers from the following metrics?

```
apiserver_request_total{job="kube-apiserver", instance="192.168.1.220:6443", cloud="aws"} 1
apiserver_request_total{job="kube-apiserver", instance="192.168.1.121:6443", cloud="gcloud"} 5
```

- A. sum by (cloud)(rate(apiserver\_request\_total{job="kube-apiserver"})[5m]))
- B. rate(apiserver\_request\_total{job="kube-apiserver"})[5m]) by (cloud)
- C. sum by (cloud) (apiserver\_request\_total{job="kube-apiserver"})
- D. rate(sum by (cloud)(apiserver\_request\_total{job="kube-apiserver"}))[5m])

### Answer: A

Explanation:

The rate() function computes the per-second increase of a counter metric over a specified range, while sum by (label) aggregates those rates across dimensions - in this case, the cloud label.

The correct query is:

sum by (cloud)(rate(apiserver\_request\_total{job="kube-apiserver"})[5m])) This expression:

Calculates the rate of increase in API requests per second for each instance.

Groups and sums those rates by cloud, giving the total request rate per cloud provider.

Option A incorrectly places by (cloud) after rate(), which is not valid syntax.

Option B returns raw counter totals (not rates).

Option D incorrectly applies rate() after aggregation, which distorts the calculation since rate() must operate on individual time series before aggregation.

Reference:

Verified from Prometheus documentation - rate() Function, Aggregation Operators, and Querying Counters Across Labels sections.

## NEW QUESTION # 21

Which PromQL statement returns the sum of all values of the metric node\_memory\_MemAvailable\_bytes from 10 minutes ago?

- A. sum(node\_memory\_MemAvailable\_bytes) offset 10m
- B. sum(node\_memory\_MemAvailable\_bytes) setoff 10m
- C. sum(node\_memory\_MemAvailable\_bytes offset 10m)
- D. offset sum(node\_memory\_MemAvailable\_bytes[10m])

Answer: C

Explanation:

In PromQL, the offset modifier allows you to query metrics as they were at a past time relative to the current evaluation. To retrieve the value of node\_memory\_MemAvailable\_bytes as it was 10 minutes ago, you place the offset keyword inside the aggregation function's argument, not after it.

The correct query is:

sum(node\_memory\_MemAvailable\_bytes offset 10m)

This computes the total available memory across all instances, based on data from exactly 10 minutes in the past.

Placing offset after the aggregation (as in option B) is syntactically invalid because modifiers apply to instant and range vector selectors, not to complete expressions.

Reference:

Verified from Prometheus documentation - PromQL Evaluation Modifiers: offset, Aggregation Operators, and Temporal Query Examples.

## NEW QUESTION # 22

What are Inhibition rules?

- A. Inhibition rules inspect alerts when a matching set of alerts is firing.
- B. Inhibition rules inject a new set of alerts when a matching alert is firing.
- C. Inhibition rules mute a set of alerts when another matching alert is firing.
- D. Inhibition rules repeat a set of alerts when another matching alert is firing.

Answer: C

Explanation:

Inhibition rules in Prometheus's Alertmanager are used to suppress (mute) alerts that would otherwise be redundant when a higher-priority or related alert is already active. This feature helps avoid alert noise and ensures that operators focus on the root cause rather than multiple cascading symptoms.

For example, if a "DatacenterDown" alert is firing, inhibition rules can mute all "InstanceDown" alerts that share the same datacenter label, preventing redundant notifications. Inhibition is configured in the Alertmanager configuration file under the inhibit\_rules section. Each rule defines:

A source match (the alert that triggers inhibition),

A target match (the alert to mute), and

A match condition (labels that must be equal for inhibition to apply).

Only when the source alert is active are the target alerts silenced.

Reference:

Verified from Prometheus documentation - Alertmanager Configuration - Inhibition Rules, Alert Deduplication and Grouping, and Alert Routing Best Practices.

## NEW QUESTION # 23

Given the following Histogram metric data, how many requests took less than or equal to 0.1 seconds?

apiserver\_request\_duration\_seconds\_bucket{job="kube-apiserver", le="+Inf"} 3

```
apiserver_request_duration_seconds_bucket{job="kube-apiserver", le="0.05"} 0
apiserver_request_duration_seconds_bucket{job="kube-apiserver", le="0.1"} 1
apiserver_request_duration_seconds_bucket{job="kube-apiserver", le="1"} 3
apiserver_request_duration_seconds_count{job="kube-apiserver"} 3 apiserver_request_duration_seconds_sum{job="kube-apiserver"} 0.554003785
```

- A. 0
- B. 1
- C. 2
- D. 0.554003785

#### Answer: A

Explanation:

In Prometheus, histogram metrics use cumulative buckets to record the count of observations that fall within specific duration thresholds. Each bucket has a label le ("less than or equal to"), representing the upper bound of that bucket.

In the given metric, the bucket labeled le="0.1" has a value of 1, meaning exactly one request took less than or equal to 0.1 seconds. Buckets are cumulative, so:

le="0.05" → 0 requests  $\leq$  0.05 seconds

le="0.1" → 1 request  $\leq$  0.1 seconds

le="1" → 3 requests  $\leq$  1 second

le="+Inf" → all 3 requests total

The \_sum and \_count values represent total duration and request count respectively, but the number of requests below a given threshold is read directly from the bucket's le value.

Reference:

Verified from Prometheus documentation - Understanding Histograms and Summaries, Bucket Semantics, and Histogram Query Examples sections.

#### NEW QUESTION # 24

```
http_requests_total{verb="POST"} 30
http_requests_total{verb="GET"} 30
```

What is the issue with the metric family?

- A. Metric names are missing a prefix to indicate which application is exposing the query.
- B. Unit is missing in the http\_requests\_total metric name.
- C. The value represents two different things across the dimensions: code and verb.
- D. verb label content should be normalized to lowercase.

#### Answer: B

Explanation:

Prometheus metric naming best practices require that every metric name include a unit suffix that indicates the measurement type, where applicable. The unit should follow the base name, separated by an underscore, and must use base SI units (for example, \_seconds, \_bytes, \_total, etc.).

In the case of http\_requests\_total, while the metric correctly includes the \_total suffix-indicating it is a counter-it lacks a base unit of measurement (such as time, bytes, or duration). However, for event counters, \_total is itself considered the unit, representing "total occurrences" of an event. Thus, the naming would be acceptable in strict Prometheus terms, but if this metric were measuring something like duration, size, or latency, then including a specific unit would be mandatory.

However, since the question implies that the missing unit is the issue and not the label schema, the expected answer aligns with ensuring metric names convey measurable units when applicable.

Reference:

Prometheus documentation - Metric and Label Naming Conventions, Instrumentation Best Practices, and Metric Type Naming (Counters, Gauges, and Units) sections.

#### NEW QUESTION # 25

.....

Our experts offer help by diligently working on the content of PCA learning questions more and more accurate. Being an exam candidate in this area, we believe after passing the exam by the help of our PCA practice materials, you will only learn a lot from this

PCA Exam but can handle many problems emerging in a long run. You can much more benefited from our PCA study guide. Don't hesitate, it is worthy to purchase!

**PCA Related Exams:** [https://www.practicedump.com/PCA\\_actualtests.html](https://www.practicedump.com/PCA_actualtests.html)

DOWNLOAD the newest PracticeDump PCA PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1WOUMn64q7bZ7PcEXUTg4NLiP68DGT2us>