

Exam : C1000-189

**Title : IBM Certified Instana
Observability v1.0.277
Administrator - Professional**

Version : DEMO

1.Which data source on the analytics page shows traces?

- A. Infrastructure
- B. Applications
- C. Logs
- D. Websites

Answer: B

Explanation:

Instana's Analytics page provides a consolidated environment for users to query and visualize operational data across their stack. According to the official IBM Instana Observability documentation, traces—comprising the end-to-end journey of requests across services—are found specifically under the Applications data source. The Applications section gives interactive access to traces, requests, response times, call hierarchies, and distributed dependencies. This is possible because Instana's agent and tracers automatically instrument applications to capture and send detailed trace data. The documentation states, "The Applications analytics section allows you to interactively work with service traces and requests, providing distributed tracing visibility." This allows users to drill down, identify bottlenecks, and analyze errors at the service interaction and code execution level.

Infrastructure data source focuses on system-level metrics (CPU, memory, disk), Logs cover textual/semi-structured log output, and Websites relate to synthetic and real-user measurements— but only Applications feature distributed tracing as per the IBM Instana Observability product documentation. Thus, for incident response, root-cause analysis, and performance breakdowns, always consult the Applications data source for trace-level data.

Reference: IBM Instana Observability Documentation, Analytics Overview.

2.At which level can AWS agent polling intervals for CloudWatch API be configured?

- A. Resource group
- B. Region
- C. Account
- D. Service

Answer: B

Explanation:

AWS monitoring through Instana involves integration with the CloudWatch API to retrieve platform and service metrics. The official IBM Instana Observability documentation affirms that polling intervals for CloudWatch can be set at the Region level. This means an administrator configures how frequently Instana's agent queries CloudWatch within each specified region independently. This level of granularity provides flexibility: for example, mission-critical regions may be monitored more frequently, while others are polled less often to reduce API costs or remain within AWS rate limits. The documentation specifies: "Instana Agents for AWS can be configured with a polling interval for CloudWatch that is set per Region to customize granularity and resource consumption." Polling cannot be set at the account, resource group, or individual service level in default configuration. Instana's region-based polling helps balance data accuracy and overhead, especially in global or multi-region deployments. If needed, changes are applied through YAML configuration or UI during AWS agent integration setup.

Reference: IBM Instana Observability Documentation, AWS Monitoring, Agent Configuration.

3.When installing the Instana host agent on Kubernetes, which option is valid?

- A. Homebrew
- B. Binary
- C. Operator
- D. RPM

Answer: C

Explanation:

The Instana Operator is the officially recommended and supported method for deploying the Instana host agent on Kubernetes clusters. The IBM Instana Observability documentation states, "The recommended method to install the Instana agent on Kubernetes clusters is via the Instana Operator, which uses Custom Resources to simplify lifecycle management." The Operator pattern in Kubernetes automates not just installation, but also upgrades, configuration, and management of agents across the entire cluster. This ensures security and reliability because the Operator reacts to cluster changes and can self-heal agent deployments. Other install options such as Homebrew, direct binary, or RPM are for traditional VM or bare-metal hosts—not for orchestrated container environments like Kubernetes. Only with the Operator does Instana support automated scaling, configuration through CRDs, and native Kubernetes best practices. Helm charts are also often involved in configuring the Operator, further streamlining agents' deployment in public, private, or hybrid cloud clusters.

Reference: IBM Instana Observability Documentation, Kubernetes Installation, Operator Lifecycle Management.

4.Which HTTP header is automatically collected?

- A. x-client-id
- B. Instana-probe
- C. Instana-id
- D. X-Instana-Service

Answer: D

Explanation:

Instana traces and analyzes every request. Services and endpoints are automatically discovered, and relationships between services, endpoints, and your infrastructure are autocorrelated and stored in our Dynamic Graph.

Based on the data that is collected from tracers and sensors, KPIs are calculated for calls, latency, and erroneous calls. KPIs help you discover the health of every individual service and then the health of your entire infrastructure.

Services are a part of application monitoring and provide a logical view of your system. Services are derived from infrastructure entities such as hosts, containers, and processes. Incoming calls are correlated to infrastructure entities and enriched with infrastructure data; for example, the Kubernetes pod label or SpringBoot application name. After this infrastructure-linking processing step, a service mapping step maps the enriched calls to generate a service name per call based on a set of rules. Instana comes with an extensive set of predefined rules to generate the best possible service name for you automatically. To fine-tune the service mapping, you can create your own custom rules, see [customize service mapping](#).

5.Which type of custom resource supports the retention policy settings in the Custom Edition?

- A. StorageConf

- B. CoreSpec
- C. UnitProp
- D. ConfigYaml

Answer: B

Explanation:

According to the official IBM Instana Observability documentation (v1.0.304), retention policy settings in Custom Edition are NOT configured in a custom resource called "StorageConf." Instead, they are configured as properties within the CoreSpec of the Core custom resource. The documentation explicitly states: "Overwriting the default retention settings is optional and should only be done consciously. These retention setting values are configured as properties in the CoreSpec."

The actual configuration looks like this:

text

kind: Core

metadata:

name: instana-core

namespace: instana-core

spec:

properties:

- name: retention.metrics.rollup5

value: "86400"

- name: config.appdata.shortterm.retention.days

value: "7"

- name: config.synthetics.retention.days

value: "60"

The retention policies for infrastructure metrics, application data, and synthetic monitoring are all configured as properties within the Core spec, not in a separate "StorageConf" custom resource. "StorageConf" refers to storage configurations for raw spans (S3, GCS, Azure), not retention policies. Reference: IBM Instana Observability Documentation (v1.0.304) — Installing the Instana Backend, Overwriting Data Retention Defaults section in CoreSpec.