

Scenario 01 – Degraded Performance Across Multiple Services With Unknown Cause

The starting point for this scenario is as follows:

The Operations Center views a consolidated service health overview in the observability platform covering the services under its responsibility. Several services simultaneously show signs of performance degradation, including increased response times and reduced availability. The underlying cause is not immediately apparent and no obvious application related failure can be identified.

The Operations Center uses the observability platform to analyze the relationships between the affected services. The platform provides insight into shared underlying components and dependencies, making it clear that the observed issues are unlikely to be isolated incidents and are more likely caused by a common underlying problem.

By correlating metrics, logs, traces and events across applications, infrastructure and network components the Operations Center identifies the component most likely responsible for the issue. The platform's AI-assisted root cause analysis capabilities support this process by highlighting the most likely failing component based on observed correlations and by suggesting potential remediation actions based on historical incidents and recognized patterns. The Operations Center uses this information to validate the root cause and understand how it impacts the affected services.

The Operations Center assigns the incident directly to the appropriate specialist team within the observability platform, including all relevant context such as affected services, identified root cause and recommended actions. The specialist team receives the incident and can view a consolidated overview of the service they are responsible for, including service components and the impact of the identified issue.

At the same time, the incident is transferred to the IT service management platform where all relevant context and analysis from the observability platform is preserved.

The Operations Center and the specialist team collaborate on further investigation and resolution, either directly in the observability platform or through integrated collaboration tools. Relevant context, analysis and status updates follow the incident throughout its lifecycle without requiring manual transfer of information between systems.

Once the issue has been resolved, the Operations Center verifies that service performance has returned to normal and that all affected services have recovered.

Expected Solution Capabilities

- A consolidated service health view showing impacted services and business context.
- Automatic correlation of metrics, logs, traces and events across applications, infrastructure and network components.
- Identification of the most likely root cause and visualization of dependencies between affected components.
- AI assisted analysis that supports root cause identification and suggests possible remediation actions.

- Clear visibility into service impact, affected users and relationships between services and underlying components.
- Assignment of incidents to specialist teams while preserving all relevant context and investigation data.
- Integration with ITSM platforms, including automatic transfer of incident context and analysis results.
- Collaboration capabilities that allow multiple teams to work on the same incident without losing context.
- Service specific dashboards and views tailored to different roles and areas of responsibility.

Scenario 02 – Deviation From Normal Behaviour Detected Through Dynamic Thresholds

The starting point for this scenario is as follows:

The Operations Center and the responsible specialist team receive an alert from the observability platform indicating that a service is behaving abnormally. The service has operated reliably for an extended period, with predictable resource consumption and consistent log and metric patterns. No static threshold has been breached.

Further investigation reveals that resource utilization has gradually changed over time and that certain values now differ significantly from the service historical behavior. At the same time, patterns within the logs have changed, including an increase in specific error types and events that have previously remained stable.

Although none of the observed values exceed predefined thresholds the behaviour clearly deviates from the historical baseline of the service. The observability platform detects these anomalies using dynamic thresholds and analysis of historical trends and behaviour.

The platform highlights which values, signals and patterns differ from the established baseline and provides insight into how the behaviour has evolved over time. Based on this information, the specialist team can assess the potential risk and take corrective action before the issue develops into a user impacting incident.

Once remediation has been completed, the specialist team verifies that both resource utilization and log patterns have returned to normal operating conditions.

Expected Solution Capabilities

- Detection of anomalies without relying on static thresholds.
- Dynamic baselines established from historical service behaviour.
- Identification of abnormal patterns in both metrics, log data and traces.
- Visibility into what specifically deviates from normal behaviour and why it has been flagged.

- Visual comparison between normal and abnormal behaviour over time.
- Root cause analysis capabilities that help explain the detected deviation.
- Early warning capabilities that provide proactive indication of emerging issues before user impact occurs.
- Reduction of alert noise compared to traditional threshold based monitoring.
- Demonstration of anomaly detection capabilities for metrics, logs, traces, synthetic monitoring and Real User Monitoring (RUM) where supported by the platform.
- Verification that remediation actions have restored the service to its expected operating state.

Scenario 03 – Performance Degradation Following a Change

The starting point for this scenario is as follows:

A planned change is implemented on an application or component, either as a configuration change or as a new version deployed through a CI/CD pipeline. During the change window the application is automatically placed in maintenance mode to prevent unnecessary alerts from being generated.

After the change has been completed the Operations Center observes signs of performance degradation within the observability platform. The degradation coincides with the timing of the change. The platform also provides relevant information about the change, including the health of associated CI/CD pipelines, deployment activities and events related to the delivery process.

Using distributed tracing and correlation across metrics, logs, traces, events and other relevant observability data the Operations Center and the responsible specialist team analyze how the application behaves following the change. By comparing the state of the application before and after the change they identify which configuration change, deployment or application component introduced the bottleneck.

The platform provides visibility into the relationship between the change and the subsequent degradation and helps identify the underlying cause of the issue.

The incident is escalated to the responsible specialist team with full context, including details about the change, the affected service and the identified resource consumption or bottleneck. The specialist team either resolves the issue or rolls back the change based on the findings.

The Operations Center and the specialist team then verify that application performance returns to normal and that the service is operating as expected.

Expected Solution Capabilities

- Automatic correlation between service degradation and recent changes, deployments, configuration updates or approved changes originating from change management and CI/CD platforms.
- Visibility into CI/CD pipeline health, deployment activities, change records and related operational events.

- Comparison of service behaviour, configuration, topology and observability data before and after a change.
- Clear visibility into what changed, when the change occurred and how the change impacted the service.
- Correlation of metrics, logs, traces and events to determine the impact of a change.
- Ability to trace service degradation back to a specific change, deployment or configuration update.
- Distributed tracing capabilities that help identify where the degradation occurs.
- Identification of the underlying cause of the performance issue and the affected components.
- Clear visualization of the relationship between the change and the resulting degradation.
- Preservation of relevant context when escalating incidents to specialist teams.
- Support for validating the outcome of remediation actions or rollback activities.
- Verification that service performance has returned to its expected state after corrective actions have been completed.

Scenario 04 – Reducing Alert Noise and Identifying Underlying Problems

The starting point for this scenario is as follows:

The Operations Center and the responsible specialist team observe that an application or infrastructure component generates a large number of alerts over time, many of which are repetitive and do not represent actual service impact.

Using the observability platform they identify that several alerts occur repeatedly without resulting in corrective actions or escalation. This makes it difficult to distinguish genuinely important incidents from operational noise.

The specialist team uses the platform to filter, group and analyze events and alerts, gaining insight into which alerts repeatedly occur without providing operational value. Based on this analysis, they adjust monitoring configurations, alerting rules and thresholds to reduce unnecessary noise.

During the same investigation they discover that a particular event recurs under similar conditions over an extended period. While each individual occurrence appears to have limited impact, historical analysis reveals a recurring pattern that warrants further investigation.

The specialist team uses the observability platform to analyze the frequency, scope and historical context of the recurring event and determines that it represents an underlying problem rather than isolated incidents.

A problem record is created based on the information gathered from the observability platform and is transferred to the organization's problem management process with all relevant context attached.

Over time, the specialist team can verify that the number of unnecessary alerts decreases while recurring patterns are more effectively identified, tracked and managed as underlying problems.

Expected Solution Capabilities

- The ability to distinguish between operational noise and incidents that require action.
- Identification of recurring events and patterns across historical data.
- Correlation of alerts, events, metrics, logs and traces to improve understanding of recurring issues.
- Alert grouping, deduplication, suppression and noise reduction capabilities.
- Visibility into alert frequency, impact and historical trends over time.
- The ability for authorized users to modify alert rules, thresholds and monitoring configurations in a controlled and auditable manner.
- Support for identifying potential problems based on recurring operational patterns.
- Integration with problem management processes and external ITSM platforms.
- AI-assisted analysis that highlights recurring patterns, abnormal behaviour and potential underlying causes.
- Verification that monitoring improvements result in reduced alert noise while maintaining visibility into genuine service issues.

Scenario 05 – Topology, Auto-Discovery and Data Collection for New Services

The starting point for this scenario is as follows:

A new application is introduced into production and must be incorporated into the organization monitoring strategy. The application is instrumented using OpenTelemetry and sends metrics, logs and trace data to the observability platform.

The platform automatically discovers the new components, identifies the technologies involved and enables relevant monitoring capabilities based on the detected environment. At the same time the platform receives metrics, logs and traces through OpenTelemetry. If the platform uses proprietary collection mechanisms or agents in addition to OpenTelemetry all telemetry data is processed consistently and participates equally in analysis, correlation and visualization.

The Operations Center can immediately view the application within a consolidated and up-to-date topology view alongside existing services. The topology shows the application's components, the systems it communicates with and its dependencies on underlying

infrastructure and supporting services. Business context such as criticality, support information and responsible specialist teams is also available.

The platform provides visibility into the status of monitoring coverage and confirms that telemetry collection is functioning as expected for the newly discovered components.

When an issue occurs within the new application or one of its dependencies, the Operations Center can use the topology view to understand which components are affected, identify impacted services and use this information as the starting point for further investigation.

Expected Solution Capabilities

- Automatic discovery of newly deployed applications, services and infrastructure components.
- Integration with CMDB and service management systems to enrich discovered services and components with ownership, criticality, support information, lifecycle status and business context.
- Automatic identification of technologies, dependencies and relationships between components.
- Support for OpenTelemetry as a data source for metrics, logs and traces.
- Consistent handling of OpenTelemetry data alongside data collected through native platform collectors or agents.
- Automatic generation and maintenance of service maps and topology views.
- Visualization of dependencies between applications, infrastructure, databases, network components and external services.
- Enrichment of topology and service views with operational context such as criticality, ownership, support information and responsible teams.
- Visibility into monitoring coverage and confirmation that telemetry collection is functioning correctly for newly onboarded services.
- Automatic updates to topology and dependency views as services evolve over time.
- The ability to understand service impact and dependency relationships when investigating incidents affecting newly onboarded applications.

Scenario 06 – Monitoring and Correlation Across Environments and Data Sources

The starting point for this scenario is as follows:

The Operations Center receives alerts indicating degraded performance within a service that consists of components running both on premises and in a public cloud environment. Part of the telemetry from the cloud hosted components is not collected through the observability

platform's native collection mechanisms but originates from cloud native monitoring services, third-party observability tools or other internal and external data sources.

The service is experiencing performance degradation and the Operations Center uses the observability platform to obtain a consolidated view of the service regardless of where the components are running or where the telemetry originates. The platform processes and correlates metrics, events, logs and traces from external sources alongside data collected through its own collection mechanisms and presents all relevant information within the same operational context.

The platform visualizes the relationships between on premises components, cloud resources and dependent services, allowing the Operations Center to follow an issue across multiple environments within a single investigation workflow.

Through this analysis, the Operations Center identifies which part of the service chain is responsible for the degradation and escalates the incident to the appropriate specialist team with complete context covering the entire service path.

After the issue has been resolved, the Operations Center verifies that all relevant telemetry across environments and data sources has returned to normal operating conditions.

Expected Solution Capabilities

- A unified operational view of services spanning on premises, cloud, SaaS and hybrid environments.
- Correlation of metrics, events, logs and traces from multiple internal and external data sources.
- Consistent analysis, visualization, alerting and root cause investigation regardless of the origin of the telemetry.
- The ability to investigate incidents across multiple environments within a single workflow without switching between tools.
- Support for cloud native telemetry and integration with third-party monitoring and observability platforms.
- Correlation of externally sourced telemetry with data collected through the platform's own collectors and agents.
- End-to-end visibility across service dependencies spanning infrastructure, applications, cloud resources and external services.
- Clear visualization of relationships between services, components and environments.
- The ability to identify where an issue originates within a distributed service chain and understand downstream impact.
- Support for open standards such as OpenTelemetry and equivalent industry standard telemetry formats.
- Visibility into supported cloud platforms, monitoring solutions, logging platforms and other data sources available through native integrations.

- Verification that telemetry from all relevant sources contributes equally to analysis, correlation, dashboards, alerts and investigations.

Scenario 07 – Degraded User Experience Identified Through RUM, Synthetic Testing and Distributed Tracing

The starting point for this scenario is as follows:

The Operations Center receives alerts from the observability platform indicating degraded performance in a critical application. At the same time, users report slow response times during login and while performing key transactions.

Using Real User Monitoring (RUM) the Operations Center observes increased response times for actual user interactions and confirms that multiple users are experiencing degraded performance during the same period.

At the same time, synthetic tests that simulate critical business transactions show increasing response times, while some steps fail or take significantly longer than expected. The Operations Center can therefore confirm that the issue affects both real users and automated transaction monitoring.

Using distributed tracing, the Operations Center follows a transaction end to end through the application and its underlying components. The trace data reveals where delays occur within the service chain and identifies the component or service contributing most to the performance degradation.

The platform presents user experience data, application performance information and infrastructure context within a single view, making it possible to understand how technical issues translate into user impact.

The platform also provides visibility into how endpoint conditions, user connectivity and network performance affect the user experience, allowing the Operations Center to distinguish between application related problems and issues originating from the user workspace or network.

The incident is escalated with complete context, including affected transactions, relevant measurements and identified bottlenecks. The responsible specialist team investigates and resolves the issue.

Once remediation has been completed the Operations Center verifies that both real user experience metrics and synthetic tests indicate normal performance and that the service is operating as expected.

Expected Solution Capabilities

- Correlation of Real User Monitoring, synthetic testing and distributed tracing within a single investigation workflow.
- Visibility into how technical issues affect real users and critical business transactions.
- End-to-end tracing of user transactions across applications, services and infrastructure components.
- Identification of bottlenecks, delays and failures throughout the transaction flow.

- Visualization of user impact alongside application, infrastructure and network telemetry.
- The ability to distinguish between application issues and problems related to client devices, network connectivity or user workspaces.
- Support for proactive detection of performance degradation through synthetic testing before users are affected.
- Visibility into transaction performance, error rates, response times and user experience metrics.
- Correlation between user experience indicators and underlying application or infrastructure events.

Scenario 08 – Capacity Analysis and Future Capacity Planning

The starting point for this scenario is as follows:

A business-critical application is planned to support a larger number of users and there is a need to determine whether the existing infrastructure has sufficient capacity to support the expected growth.

Capacity managers and specialist teams use the observability platform to analyze historical resource utilization and performance data over time. They review trends across relevant infrastructure components such as servers, storage systems and network resources and identify how capacity has been utilized up to the present day.

By comparing different periods, they can see how resource consumption has evolved and identify components that are approaching their capacity limits. The platform also provides insight into which services, applications or workloads are driving resource consumption.

Based on this information the teams assess how the planned increase in demand is likely to affect the environment. The platform supports future planning by providing visibility into historical trends, growth patterns and projected capacity requirements.

Using these insights, capacity managers can identify the need for actions such as scaling, resource optimization or architectural changes before service performance is affected.

After the planned changes have been implemented, the teams use the observability platform to verify how resource utilization develops over time and whether the expected outcome has been achieved.

Expected Solution Capabilities

- Historical analysis of resource utilization across servers, storage, network infrastructure, databases and other monitored components.
- Identification of growth trends and long term capacity consumption patterns.
- Visibility into which services, applications or workloads contribute most to resource consumption.
- Capacity forecasting based on historical utilization and growth trends.

- The ability to identify components approaching capacity limits before service performance is impacted.
- Comparison of resource utilization across different time periods.
- Reporting capabilities that support capacity planning and management processes.
- Correlation between capacity utilization, performance metrics and operational events.
- Support for planning scaling activities and evaluating future resource requirements.
- Visibility into whether implemented capacity improvements have achieved the intended outcome.
- Consolidated dashboards and reports providing an overall view of current and projected capacity status.

Scenario 09 – Access Control, Data Handling and Traceability During an Incident

The starting point for this scenario is as follows:

The Operations Center identifies an incident within the observability platform that requires further investigation by the responsible specialist team.

The specialist team accesses the observability platform and is granted access only to the data associated with the service they are responsible for. They can view relevant logs, metrics, trace data and events related to that service but do not have access to data belonging to other services or applications.

As part of the investigation, external vendor support is involved. The organization grants limited access to the external resource, allowing visibility only into the data relevant to the specific incident. The platform ensures that information outside the approved scope remains inaccessible and that sensitive operational data is not exposed unnecessarily.

During the investigation, log data may contain sensitive information. The observability platform ensures that sensitive data is handled appropriately during ingestion by filtering, masking or removing sensitive content before it is stored and presented to the Operations Center, specialist teams or other authorized users.

A platform administrator can manage access permissions and retention settings in accordance with organizational policies and regulatory requirements. The platform ensures that metrics, logs, events and trace data are retained for an appropriate period to support troubleshooting, trend analysis and operational investigations while maintaining control of stored data.

Access permissions, retention settings and handling of sensitive information are managed directly within the platform without requiring assistance from the vendor.

Once the incident has been resolved, the Operations Center verifies that the investigation was conducted using only the data and access necessary to resolve the issue and that sensitive information remained protected throughout the process.

Expected Solution Capabilities

- Role-based access control that restricts visibility based on business need and area of responsibility.
- Granular access control for metrics, logs, events and trace data.
- The ability to restrict access at service, application, data source and telemetry level.
- Secure and controlled access for external vendors, consultants and third-party support resources.
- Protection of sensitive information through masking, filtering, redaction or similar mechanisms during data ingestion and storage.
- Independent administration of access policies, retention settings and governance controls without vendor involvement.
- Configurable retention policies for metrics, logs, traces and other telemetry data.
- Visibility into inherited permissions, role assignments and group memberships.
- Verification that users only have access to data relevant to their responsibilities and authorized scope.

Scenario 10 – Integration and Automation Through APIs

The starting point for this scenario is as follows:

The Operations Center and specialist teams need to integrate the observability platform with other enterprise systems and automate operational tasks that currently require manual effort.

A specialist team needs to retrieve data from the observability platform for use in custom workflows, integrations and operational processes. Using the platform's APIs, they can access relevant information such as service health, active alerts and events, monitoring coverage and historical telemetry data.

The retrieved information is used to build automated workflows and integrations with other systems without requiring manual exports or assistance from the vendor.

At the same time, the Operations Center needs to enrich the observability platform with information originating from external systems. The platform allows external events, configuration data, operational context and other relevant information to be ingested through APIs. Once ingested, this information is processed and analyzed alongside data collected directly by the platform.

Access to APIs is governed using the same authentication, authorization and security controls as the platform itself, ensuring consistent access management, traceability and governance.

The Operations Center verifies that data retrieved through APIs is consistent with the information presented in the platform and that the APIs are sufficiently documented to enable integrations, automation and ongoing maintenance by internal resources.

Expected Solution Capabilities

- Well-documented APIs that support both data retrieval and data ingestion.

- The ability to retrieve service status, alerts, events, topology information, observability data and historical telemetry through APIs.
- The ability to ingest external events, configuration data, business context and operational metadata through APIs.
- Support for automation of operational activities, monitoring configuration, onboarding processes and integration workflows.
- Consistent handling of externally sourced data alongside data collected directly by the platform.
- Secure API access using the same authentication and authorization mechanisms as the platform itself.
- Traceability and auditability of API activity and automated actions.
- Integration with ITSM, CMDB, automation platforms, CI/CD pipelines, collaboration tools and other enterprise systems.
- The ability to automate the creation, modification and maintenance of monitoring configurations through APIs.
- Consistency between data presented through APIs and data displayed within the user interface.
- API capabilities that can be utilized and maintained by internal resources without requiring vendor involvement.
- Clear documentation of API limitations, rate limits, authentication models and any functionality that is only available through the user interface.

Scenario 11 – Control of Telemetry Consumption, Licensing and AI Costs

The starting point for this scenario is as follows:

A platform administrator receives an alert from the observability platform indicating that telemetry data consumption has increased significantly compared to normal levels. At the same time forecasts show that the current growth trend may lead to increased licensing and operational costs if it continues.

The platform administrator investigates the anomaly using the platform's telemetry consumption views and identifies which data sources contribute most to the increase. Through historical data and trend analysis, the platform administrator determines when the anomaly began, how consumption has evolved over time and which areas deviate from the established baseline.

Further analysis reveals that a particular data source is generating substantially more telemetry than normal due to a configuration issue. The platform administrator verifies that existing control mechanisms have already limited data collection from the affected source according to

predefined policies and that the platform generated appropriate alerts so the issue could be investigated and corrected.

The platform administrator also gains insight into current license utilization and associated costs, including how these are influenced by telemetry consumption. The platform provides forecasts based on historical and current usage patterns, enabling future resource requirements, license consumption and costs to be evaluated and planned in advance.

Within the same operational view, the platform administrator can see how AI-related capabilities contribute to platform consumption and cost, including any impact on licensing models and future expenditure. Where required, limits, alerts and governance controls can be configured to ensure that AI functionality is used in a controlled and predictable manner.

Once corrective actions have been completed the platform administrator verifies that telemetry consumption has returned to expected levels, license utilization is under control and the platform continues to operate within defined technical and financial boundaries.

Expected Solution Capabilities

- Visibility into total telemetry consumption across the platform.
- Historical and current telemetry consumption analysis with trend visualization.
- Identification of the data sources, services, applications or technology domains contributing most to telemetry growth.
- Detection of abnormal telemetry consumption and unexpected growth patterns.
- Configurable alerts for abnormal telemetry usage and cost related anomalies.
- Built-in mechanisms to limit, control, throttle or otherwise manage excessive telemetry ingestion.
- Visibility into the relationship between telemetry consumption, license utilization and operational costs.
- Forecasting of future telemetry consumption, license requirements and associated costs based on historical trends.
- Breakdown of telemetry consumption and costs by service, application, technology domain, team or other organizational groupings.
- Visibility into AI-related consumption and the associated impact on licensing and operational costs.
- Governance controls for AI functionality, including limits, alerts, quotas and other cost-control mechanisms.
- The ability to identify and investigate the root cause of unexpected telemetry growth.
- Administration of telemetry controls, cost controls and AI governance settings without requiring vendor assistance.

Scenario 12 – Access Governance, Audit and Traceability

The starting point for this scenario is as follows:

A platform administrator receives an alert indicating that a specialist team has access to data outside its defined area of responsibility. The alert is generated by the observability platform's access governance and audit capabilities.

The platform administrator investigates the situation and identifies the roles, permissions and group memberships assigned to the affected users. The platform provides visibility into which services, applications and observability data are accessible as well as how permissions have been assigned through direct roles, inherited permissions or group memberships.

Using audit logs and traceability features the platform administrator determines when access was granted, who performed the change and whether additional privileged activities have been performed since the permissions were assigned. Historical information allows the administrator to verify whether the access was granted in accordance with established governance policies and legitimate business need.

The platform administrator performs an access review for the affected area and identifies users with excessive, unexpected or outdated permissions. The platform provides capabilities to review, validate and adjust access rights without requiring manual analysis across multiple systems.

The identified access deviations are corrected and the platform administrator verifies that access to services, applications and observability data is restricted to the appropriate users and groups. The platform supports granular access control at service, data source and telemetry level, allowing access to metrics, events, logs and trace data to be managed independently.

Once corrective actions have been completed, the platform administrator documents the findings, verifies that governance policies are being followed and confirms that all privileged activities are fully traceable through the platform's audit capabilities.

Expected Solution Capabilities

- Centralized administration of roles, groups and access permissions.
- Role based access control aligned with organizational responsibilities and business need.
- Granular access control at service, application, data source and telemetry level.
- Independent control of access to metrics, logs, events and trace data.
- Visibility into effective permissions, inherited permissions and group memberships.
- Audit logging of access changes, role assignments, permission modifications and privileged activities.
- Full traceability of who performed a change, what was changed and when the change occurred.
- Support for periodic access reviews and governance processes.
- Identification of excessive, unexpected or unused permissions.

- Historical visibility into permission assignments and access related activities.
- The ability to verify compliance with governance policies and the principle of least privilege.
- Administration of access controls and governance policies by internal resources without vendor involvement.