
Monitoring and Troubleshooting

Monitoring Tools and Incident Response Workflows

A Comprehensive Guide to IT Infrastructure
Monitoring and Incident Management



— What is Monitoring?

Monitoring is the continuous process of **collecting, analyzing, and tracking information** about systems (like servers, applications, or networks) to ensure they are **working properly**.



Why Monitoring Matters in Modern IT Infrastructure



Modern IT infrastructure demands **proactive monitoring to prevent costly downtime** and ensure business

continuity as organizations rely heavily on applications for revenue streams




Infrastructure monitoring provides the **oxygen to your IT infrastructure** by collecting data needed for a complete picture of availability, performance, and resource efficiency



Proactive monitoring enables organizations to **respond to issues before they impact users**, preventing loss of time and money



The shift to cloud-native technologies requires **continuous monitoring across cloud environments**, operating systems, servers, and virtualized systems



Understanding Infrastructure Monitoring

 Infrastructure monitoring **collects data from operating systems, hypervisors, containers, databases**, network devices, applications, logs, and metrics

 The monitoring process involves **four key steps: data collection, data analysis** using statistical methods and machine learning, alerting, and remediation

 Two main approaches exist: **agent-based monitoring** with software agents on each system, and **agentless monitoring** using SSH, SNMP, and WMI protocols

 Effective monitoring delivers critical capabilities including **optimizing user experience, detecting performance degradations, pinpointing root causes**, and triggering automated remediation



Types of Monitoring Tools



Network Monitoring

Tracks network health, performance, traffic patterns, and identifies bottlenecks affecting network performance across distributed environments



Server Monitoring

Provides real-time insights into server health including CPU usage, memory utilization, disk space, and capacity planning metrics



Application Monitoring

Focuses on software performance, response times, transaction volumes, error rates, and end-user experience metrics



Infrastructure Monitoring

Assesses overall system health across hardware, network components, providing comprehensive visibility into IT infrastructure



Security Monitoring

Includes Endpoint Detection and Response (EDR) and Network Detection and Response (NDR) tools that detect potential security threats and ensure compliance with regulations and standards

Essential Observability Data Sources

High-quality monitoring insights require collecting diverse data types from multiple sources

Metrics

Quantitative data as counts or measures aggregated over time, essential for creating visualizations and identifying performance patterns

Event Logs

Generated by every system and service, offering crucial insights into system behavior and aiding in troubleshooting operational issues

Distributed Traces

Record transaction journeys through infrastructure, revealing how various environment components interact with one another

Metadata

Topology details, namespaces, and priority data that help understand the significance and impact of events across infrastructure components

UX Data

Page load times and latency metrics provide real-time insights into how users experience applications and services

Open-source Telemetry

Tools like OpenTelemetry, Prometheus, and StatsD along with cloud integrations like CloudWatch enable comprehensive observability

Monitoring Best Practices



Leverage Automation

Implement tools with automation capabilities for complete observability and transition to AIOps



Comprehensive Alerts

Create specific alerts to reduce false positives while ensuring redundancy



Prioritize Alerts

Organize notifications to ensure critical alerts are never missed



Role-Specific Dashboards

Develop custom dashboards for ITOps, security teams, and business leaders



Review Metrics Regularly

Review tracked metrics at regular intervals to avoid blind spots



Conduct Test Runs

Schedule test runs before relying on monitoring systems daily

Common Monitoring Use Cases - Part 1



Detecting Network Issues

Monitoring network traffic and metrics helps identify bottlenecks affecting performance, enabling teams to resolve issues before impacting end users



Ensuring Security and Compliance

Continuous monitoring detects potential security threats and verifies infrastructure compliance with relevant regulations and standards



Tracking Server Health

Real-time monitoring of CPU usage, memory utilization, and disk space identifies capacity issues before they affect application performance

Common Monitoring Use Cases

Part 2: Capacity Planning and Application Performance



Capacity Planning

Performance data analysis identifies areas requiring additional resources, enabling informed decisions about resource allocation for maximum efficiency



Application Performance Monitoring

Tracking response times, transaction volumes, and error rates enables early detection and resolution of application performance issues

Introduction to Incident Response

Incident Response is the **organized approach** to detecting, managing, and recovering from **unexpected events (incidents)** that disrupt normal operations such as system outages, data breaches, or security attacks.

Incident response addresses


⚠️ **IT threats including cyberattacks, security breaches, server downtime**, and service outages through systematic frameworks

🔧 The incident response lifecycle provides **step-by-step frameworks for identifying and reacting** to service outages or security threats

🔄 Effective incident response is **not linear but cyclical**, emphasizing continuous learning and improvement from each incident

🛡️ Organizations cannot prevent all incidents, making incident response capability **necessary for rapidly detecting incidents, minimizing loss**, and restoring IT services

🔧 Modern incident response integrates **monitoring, alerting, communication, and collaboration tools** to enable rapid detection and resolution



NIST Incident Response Lifecycle - Phase 1 & 2

1

Phase 1: Preparation

-  **Establish tools and resources:** Set up the right monitoring, alerting, and incident management tools needed for effective response
-  **Train the incident response team:** Develop incident response playbooks, define roles and responsibilities, and establish communication channels before incidents occur
-  **Implement preventive activities:** Based on risk assessments, deploy measures to reduce the number of incidents that can occur

Preparation is the foundation that enables rapid and effective incident response

2

Phase 2: Detection and Analysis

-  **Accurate detection and assessment:** NIST identifies this as the most difficult part of incident response for many organizations
-  **Multiple detection sources:** Detection typically starts with monitoring and alerting tools, though incidents may also be reported by customers or team members
-  **Assess impact and severity:** Analyze incident impact, apply severity levels, and determine appropriate response actions and stakeholder communications

Early and accurate detection minimizes incident impact and enables faster response

NIST Incident Response Lifecycle - Phase 3 & 4

3

Containment, Eradication, and Recovery

-  **Containment:** Isolate affected systems to keep incident impact as small as possible and prevent further spread
-  **Eradication:** Remove the root cause of the incident, including malware, unauthorized access, or system vulnerabilities
-  **Recovery:** Restore services to normal operations while monitoring for any signs of residual issues
-  Requires coordinated actions from incident response teams, clear role delegation, and continuous stakeholder communication

4

Post-Event Activity

-  Analyze the incident and response efforts to understand what happened and how effectively the team responded
-  Conduct **blameless postmortems** and root cause analysis using techniques like the 5 Whys
-  Limit chances of incident recurrence by implementing preventive measures and system improvements
-  Transform each incident into an **opportunity for organizational growth** through continuous improvement cycles

Practical Incident Response Workflow

1

Detect the Incident

Detection through monitoring tools or customer/team member reports

2

Set Up Communication

Establish dedicated Slack channels and video conference bridges

3

Assess Impact

Assign severity levels to trigger appropriate automated actions

4

Communicate with Stakeholders

Quickly and accurately inform internal and external stakeholders

5

Escalate to Responders

Page additional teams and bring them to the incident ticket

6

Delegate Roles

Assign specific incident response roles using established playbooks

7

Resolve and Learn

End business impact, transition to cleanup tasks, and conduct thorough postmortem. Maintain a robust incident timeline throughout the response process to capture crucial data for post-incident analysis and future prevention

Key Incident Response Metrics

Measuring incident response effectiveness requires tracking specific performance metrics



Severity Levels

Classify incident impact on business operations, guiding response prioritization and resource allocation decisions



Mean Time to Detect (MTTD)

Measures how quickly incidents are identified, with faster detection enabling earlier response and reduced impact



Mean Time to Resolve (MTTR)

Tracks the average time from incident detection to resolution, indicating incident response efficiency



Mean Time to Failure (MTTF)


Provides insights into system reliability by measuring the average time between failures



Service Level Agreements (SLA), Objectives (SLO), and Indicators (SLI)

Establish performance targets and measure actual performance against commitments. Tracking these metrics enables continuous improvement through data-driven insights into incident response performance and infrastructure reliability

Windows Monitoring Tools Overview



Performance Monitor (PerfMon)

Built-in tool that collects and analyzes performance data using customizable counters, creates data collector sets, and generates detailed reports

Event Viewer

Provides access to event logs for local and remote machines, categorized into Application, Security, System, and Setup logs for troubleshooting

Resource Monitor


Offers real-time monitoring of CPU, memory, disk, and network usage with detailed process-level information

Task Manager

Provides quick access to performance metrics, running processes, startup programs, and resource utilization data

Practical Exercise - Part 1: Performance Monitor

Hands-on lab to monitor system performance using Windows Performance Monitor



Performance Monitor Exercise (15 minutes)

Open PerfMon: Press Win+R, type "perfmon", press Enter

Add Counters: Click green "+", add Memory > Available MBytes and PhysicalDisk > % Disk Time

Create Data Collector:

Expand Data Collector Sets > Right-click User Defined > New > Data Collector Set

Configure:

Name it "System Performance", select "Create manually", add CPU, Memory, Disk counters

Set Interval: Set sample interval to 15 seconds, click Finish

Start Collection:

Right-click your set and select "Start", run for 5 minutes while using applications


View Results:

Stop collector, navigate to Reports > User Defined > System Performance

→ TRY IT OUT

Event Viewer Analysis and Expected Outcomes

Part 2: Event Viewer Analysis



Part 2: Event Viewer Analysis (10 min)

Open Event Viewer: Press Win+R, type "eventvwr", press Enter

Review System Logs:

Navigate to Windows Logs > System, note Warning or Error entries

Examine Details:

Click on an Error event to view Event ID, Source, and Description

Check Applications:

Navigate to Windows Logs > Application, identify errors or crashes

Filter Logs:

Use "Filter Current Log" to show Critical/Error/Warning for past 24 hours

Document Findings:

Record at least three events with Event ID, Source, and potential impact

✓ **Expected Outcomes:**

Monitor real-time system performance, create custom data collectors, use Event Viewer to identify system issues, and correlate performance metrics with system events

Understand how to monitor real-time system performance using Performance Monitor

Learn to create custom data collector sets for ongoing performance tracking

Gain proficiency in using Event Viewer to identify and diagnose system issues

Develop skills in correlating performance metrics with system events for effective troubleshooting

Summary and Key Takeaways

Effective monitoring and incident response are essential for maintaining reliable IT infrastructure



Infrastructure monitoring provides **proactive visibility into system health**, enabling organizations to detect and resolve issues before they impact users



Comprehensive monitoring requires collecting **diverse data types including metrics, logs, traces, and user experience data** from multiple sources



The NIST incident response lifecycle provides a proven framework: **Preparation, Detection and Analysis, Containment/Eradication/Recovery, and Post-Event Activity**



Successful incident response combines **monitoring tools, clear communication channels, defined roles, severity classification**, and continuous improvement through postmortems



Windows built-in tools like **Performance Monitor and Event Viewer** provide powerful capabilities for monitoring and troubleshooting on Windows systems



Implementing monitoring best practices including **automation, comprehensive alerting, role-specific dashboards, and regular metric reviews** maximizes monitoring program value



The incident response process is **cyclical, not linear**—each incident provides learning opportunities that improve future detection, response, and prevention capabilities

Continuous Monitoring + Effective Response = Reliable Infrastructure