



SIMPLY
SECURE

G DATA TechPaper #0275

G DATA Network Monitoring



Contents

- Introduction 3**
- 1. The benefits of network monitoring..... 3**
 - 1.1. Availability 3
 - 1.2. Migration and expansion..... 3
 - 1.3. Compliance 3
 - 1.4. Security 4
 - 1.5. Cloud and virtualization 4
- 2. Choosing a network monitoring solution 4**
 - 2.1. Features 4
 - 2.2. Return On Investment 5
- 3. G DATA Network Monitoring..... 5**
 - 3.1. Deployment 5
 - 3.2. Configuration..... 7
 - 3.3. Analysis 8

Introduction

The interconnectivity of IT components such as workstations, servers, smart devices, printers and many other peripherals has risen explosively in the last decade. The availability of large numbers of networked devices leads to more and more complex enterprise network deployments and management strategies. This makes it difficult to keep track of all assets and maintaining the required level of know-how. Network monitoring helps IT staff ensure business continuity by tracking the status of a wide range of network assets, including hardware as well as software. Assisted by regular reports and configurable alarms, staff can provide maintenance and support, as well as proactively reduce the number of incidents and plan infrastructure deployments and extensions. Network monitoring makes asset management, performance optimization and maintenance more efficient and cost-effective for any business from SMB to global enterprise.

1. The benefits of network monitoring

1.1. Availability

The increasing number of network devices makes it harder for administrators to identify availability risks or performance bottlenecks. Network monitoring helps counter this trend. Not only does continuous monitoring enable administrators to spot performance issues as they occur, it also allows them to track trends and predict availability issues. Using historical trend data, weak points in the network can be optimized before the load affects performance or causes an outage. When users report problems with a database, CRM system or web shop that is not available, (historical) data and error logs are very useful.

1.2. Migration and expansion

Network monitoring can support infrastructure developments such as network migration and expansion. For example, through charting network topology, administrators can identify infrastructure components in need of improvement as well as making sure the network meets any deployment-specific prerequisites.

By tracking performance over a longer time, administrators can gain insight on performance levels. Measuring points can include application and infrastructure response time, utilization, throughput and capacity. These can then serve as baseline indicators when planning new infrastructure for migration and expansion scenarios, in order to make informed decisions about scalability and availability. This helps find the balance in capacity planning, making sure peak loads are processed appropriately while preventing infrastructure from going unused most of the time.

1.3. Compliance

Because network monitoring can be configured to log a wide range of data, it is very well suited for auditing and compliance purposes. Not only does it track usage data for devices in the network infrastructure, it can also monitor default configurations and log configuration changes. This



allows businesses to prepare their infrastructure for certification and make sure they do not drift out of compliance as infrastructure is expanded over time.

1.4. Security

Not only does network monitoring help administrators manage infrastructure: its features neatly tie in to security solutions. As one of the layers in the corporate network security concept, network monitoring can help detect signs of suspicious activities, such as an unusually high network load, which can indicate Denial of Service (DoS) attacks. Infected network devices may also exhibit atypical CPU load, service behavior or memory usage as well as runaway processes or traffic caused by malware infections. Combined with a patch management solution on endpoints, network monitoring helps administrators quickly and easily detect and mitigate vulnerabilities¹.

1.5. Cloud and virtualization

For cloud infrastructure management, virtual servers and other multi-tenant scenarios, network monitoring is essential to maintain a business model. Infrastructure needs to be built and serviced in order to host a considerable customer base; network monitoring helps estimate requirements and maintain performance for all applications and services on the network. The same methods help prepare the virtualization of physical servers by measuring their read/write access, network traffic, CPU usage and other performance-related statistics. Finally, network monitoring can also be used as a tool to monitor service level agreement (SLA) terms as well as for usage-based billing.

2. Choosing a network monitoring solution

Businesses of all sizes have already discovered the added value of integrating network monitoring into their management workflow. The market for network monitoring solutions is growing accordingly: the total revenue for 2012 was estimated at \$2.2 billion and is expected to grow to \$4.5 billion by 2017, with a compound annual growth rate of 15.2%². With such a multitude of solutions available, it is important to make an informed choice when preparing to purchase one.

2.1. Features

When looking at the features of network monitoring solutions, the first point to consider is the type of server architecture. Network monitoring is available as a locally hosted solution and as a cloud service. Locally hosted solutions require a significant amount of time to be invested into infrastructure deployment, management and maintenance. Cloud services, on the other hand, reduce the required amount of time and money by making use of hosted infrastructure, which also allows solutions to be scaled easily if demand increases.

The solution should be able to ingest data from all critical network devices. Compiling a network infrastructure inventory helps to find out which network monitoring protocols should be supported. SNMP support is the bare minimum, but any additional protocol support is a plus.

¹ For more information about patch management, see G DATA TechPaper #0271 Patch Management Best Practices.

² Frost & Sullivan: Network and Application Performance Management Market (2012).



Administrators should also think about what they are planning to do with the collected data. When analyzing trends, for example, the solution should be able to save data over a longer period of time and produce the appropriate trend analysis data and charts. When the goal is to set up an early warning system, the solution must allow administrators to set threshold values and to configure (real-time) alarms when data points exceed or fall below them.

In the end, network monitoring is about people as much as it is about technology. Any solution should be easy to manage, featuring a flat learning curve and a streamlined, automated workflow. All of the functionality should be available in a unified interface, offering administrators a clear dashboard allowing them to quickly see all important status information.

2.2. Return On Investment

In addition to a feature-based comparison, the required investment in network monitoring solutions is a major factor in decision making. The solution costs, both the initial implementation and the recurring maintenance, should be offset against the gains that can be realized. This Return On Investment (ROI) calculation can then be used to compare various solutions with each other.

The most obvious gain is preventing financial losses due to down time. For example, when running a web shop, availability is a crucial component in serving customers and ensuring revenue. Similarly, downtime of a crucial database server has the potential to disrupt the productivity of an entire office, causing immediate financial losses. By using network monitoring, the number of infrastructure outages can be reduced by taking corrective action before the network slows or goes down altogether. And even if a component requires emergency maintenance, alerts make sure that staff are immediately informed, ensuring a significant reduction of the personnel time required for taking care of infrastructure issues. Moreover, network monitoring alerts and trend data help reduce time to fix by providing staff with the relevant context. By reducing the amount time spent on firefighting, IT staff can work on other, more structural projects.

3. G DATA Network Monitoring

G DATA has integrated network monitoring in its existing portfolio of endpoint security solutions. This allows administrators to take advantage of synergy between the management components as well as between the client/agent components.

3.1. Deployment

Network Monitoring is available as an optional module for all G DATA business solutions from version 14 onwards and can be easily added to new and existing G DATA deployments. The architecture is cloud-based with support from the local G DATA ManagementServer. G DATA Security Client functions as agent, which collects data from local data points as well as other network assets. Those values are then reported to G DATA ManagementServer, which in turn synchronizes them with the cloud service G DATA ActionCenter. The cloud service aggregates and stores data, which can then be acted upon: either immediately, by the server sending out alerts, or by an administrator who analyses the data later on.

This setup means that administrators do not need to worry about the deploying and managing the monitoring server infrastructure. Administration and alerts can be carried out regardless of the current load of the customer's network. The cloud option is ideal for businesses that do not have budget or time to manage local network monitoring infrastructure.

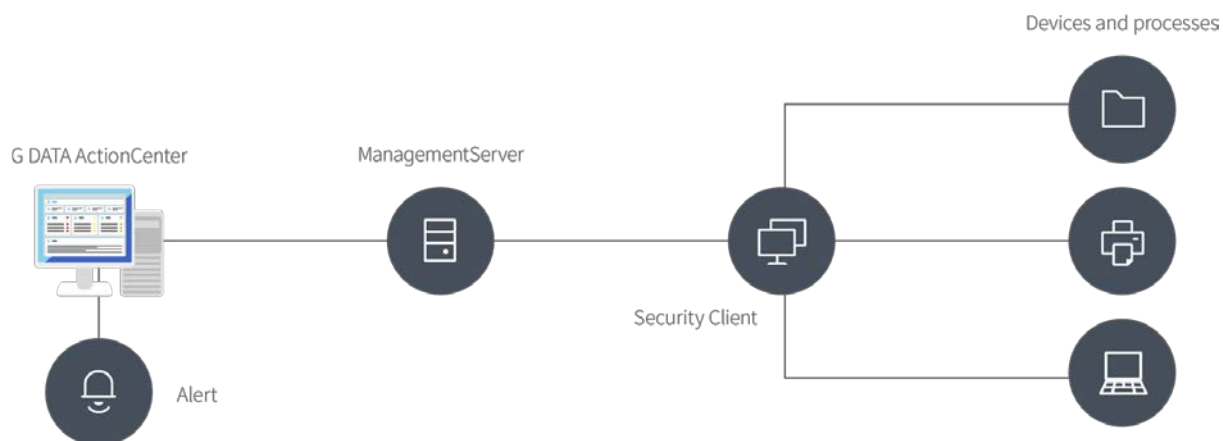


Figure 1: G DATA Network Monitoring architecture

By bringing together many different types of data, network monitoring allows for a comprehensive overview of data trends. It typically supports multiple protocols, covering virtually every network-connected asset – not just hardware, but also software and services. Examples include:

- Hardware
 - o Endpoints and servers
 - Hard drives
 - CPU
 - RAM
 - NAS
 - o Network infrastructure
 - Endpoint and server network interfaces
 - Routers
 - Switches
 - Access points
 - Firewalls
 - o Peripherals
 - (Network) printers
- Software
 - o Processes and services
 - Operating system
 - Applications
 - o Servers
 - Web
 - Database
 - Exchange
 - Domain Controller

This list is not exhaustive: many networked devices can potentially be monitored, even if there is no metric available yet. As metrics are deployed through a cloud-based service, the range of available metrics can be extended at any time. If a specific metric is required that does not exist yet, customers can file a development request with G DATA which, if feasible, will result in the metric being developed and provided in ActionCenter.

Network Monitoring relies on various protocols to gather data from assets. Because of its versatility, the Simple Network Management Protocol (SNMP) has become a popular choice for monitoring and managing many types of networked devices. It is characterized by a question-and-answer structure: a network monitoring server can use the protocol to request information from an SNMP-enabled asset about a specific property (e.g. CPU load, bandwidth usage). The device then responds with the corresponding value and the server saves the result. In addition to SNMP, two other important data acquisition sources are Performance Counters and the Windows Management Instrumentation (WMI) API, both available on the majority of agent-based Windows assets. Non-Windows-based assets can be contacted using ping or HTTP-based communication.

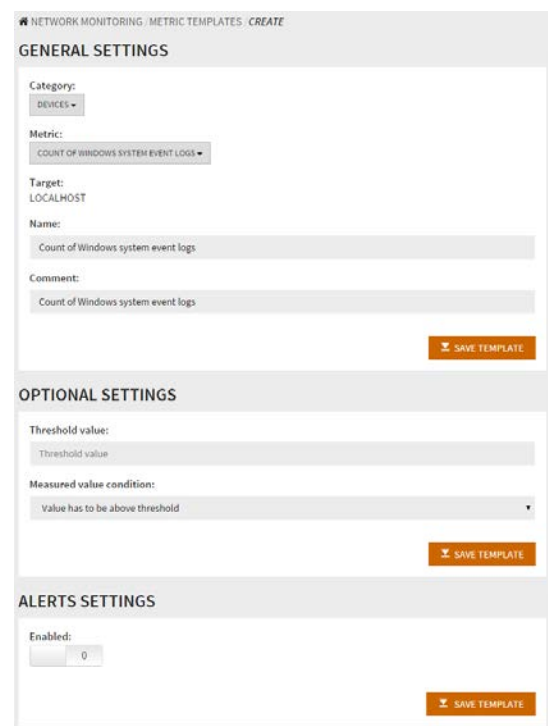
3.2. Configuration

Network monitoring is configured and managed through the web interface of G DATA ActionCenter at <https://ac.gdata.de>. In order to get started, one or more metric templates need to be created. A template contains a predefined type of monitoring and its configuration parameters. Examples include:

- Monitor a specific process on a Windows asset
- Monitor the availability of a server
- Monitor the toner level of a printer

Configuration parameters vary per template but can include a threshold value (to generate alerts if the value rises above or falls below the threshold) and one or more email addresses for alerts. For example, when measuring server availability, a threshold value of 1000 ms could be entered. The administrator will then receive an email message when it takes more than 1000 ms to reach the server.

After defining a template, a metric is created by assigning the template to one or more assets. The assets will periodically carry out the specific monitoring action and report the result to the associated ManagementServer. The server will synchronize the results to ActionCenter, which in turn carries out the actions that have been defined in the template (such as sending an alert).



The screenshot shows the 'CREATE' form for a metric template in the G DATA ActionCenter web interface. The form is organized into three main sections:

- GENERAL SETTINGS:** Includes fields for 'Category' (set to 'DEVICES'), 'Metric' (set to 'COUNT OF WINDOWS SYSTEM EVENT LOGS'), 'Target' (set to 'LOCALHOST'), 'Name' (set to 'Count of Windows system event logs'), and 'Comment' (set to 'Count of Windows system event logs'). A 'SAVE TEMPLATE' button is located at the bottom right of this section.
- OPTIONAL SETTINGS:** Includes a 'Threshold value' field (set to 'Threshold value'), a 'Measured value condition' dropdown (set to 'Value has to be above threshold'), and a 'SAVE TEMPLATE' button at the bottom right.
- ALERTS SETTINGS:** Includes an 'Enabled' checkbox (checked) and a 'SAVE TEMPLATE' button at the bottom right.

Figure 2: Template

3.3. Analysis

With one or more metrics created, administrators can use various ways to keep track of the reported data, each of which fits one or more of the typical use cases. For scenarios that rely on immediate reports, alarms are the recommended notification method. They allow for quick response times when an emergency happens. Alarms can be enabled in metric templates and are applied to all metrics that are based on the template. When enabling an alarm, it should be made sure that appropriate email groups have also been defined to make sure that crises can be swiftly dealt with. Alarm notifications can be sent to an email distribution list, such as an emergency response team as part of an IT department. It should be made sure that all alarm recipients can take action if they receive an alarm. If they should be able to carry out actions independently from the administrator, they can receive permissions to use ActionCenter themselves. At the very least, a workflow needs to be defined that makes sure that action can be taken quickly in case an emergency occurs.

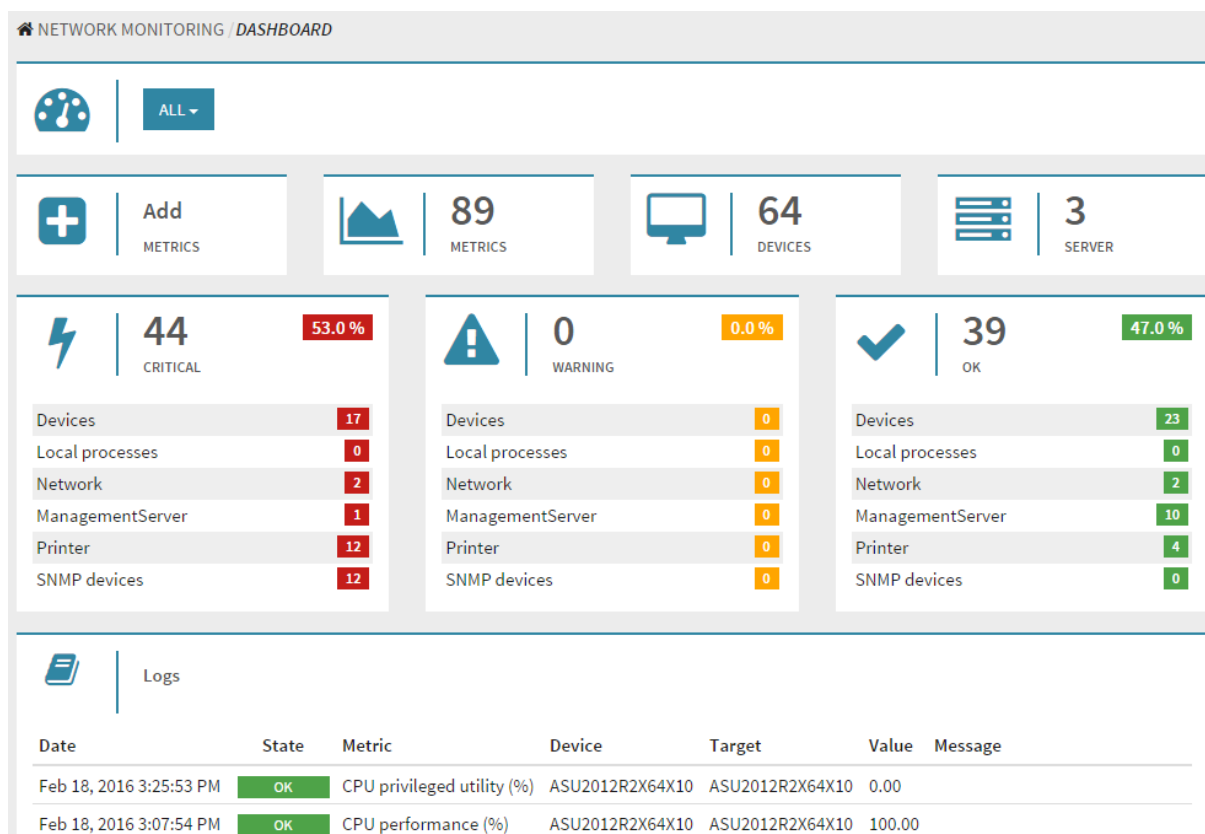


Figure 3: Dashboard

Administrators do not have to wait until alarms are sent. The Dashboard of G DATA ActionCenter shows essential information about the status of the network. Three status indicators display monitoring service statistics by priority (critical, warning and normal). This allows administrators to quickly see if any part of the network infrastructure needs to be examined. Individual services can be favorited in order to display them directly on the Dashboard, which is especially useful for important or often-used assets. The number of associated ManagementServers as well as the number of devices are also displayed, which helps keep an overview of the network.

If a more detailed analysis is required, the individual metric pages can be used. Each page shows a diagram, allowing administrators to spot trends even before they reach a critical level. The diagram can be configured to display values for a specific amount of time and can be used to pinpoint trends. When the RAM usage of a device displays an upward trend before suddenly decreasing, for example, this may indicate a memory usage problem of specific processes. Administrators can use the information to take immediate action, such as defining metrics for system processes or investigating problems locally on the device itself.

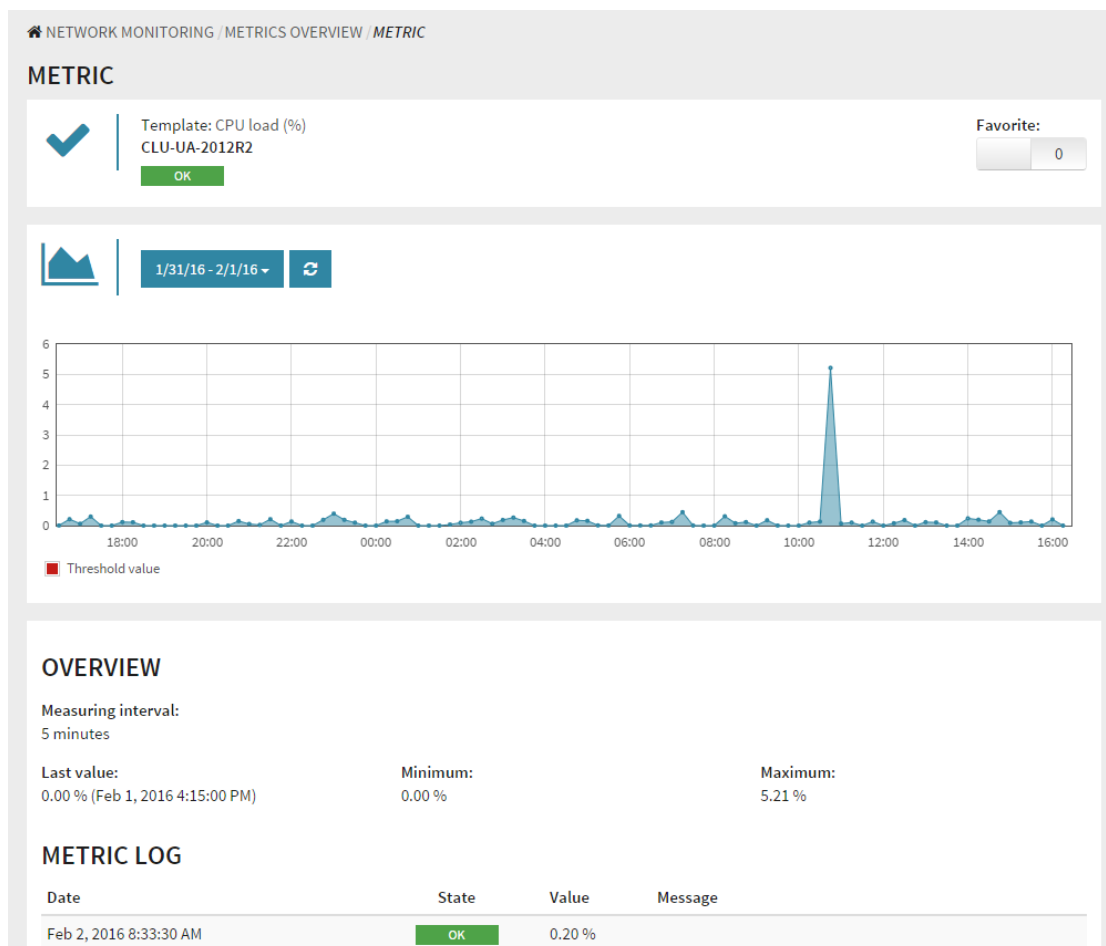


Figure 4: Metric

Using the available data for individual services, administrators can analyze trend data. Using historical monitoring data, for example, it is possible to identify peak and trough times for a network interface. These data points can be used to create a baseline of expected values and consequently to set alert thresholds. This is a process that should be optimized over time, as threshold values are not always easy to get right. Some services can still perform under load, leading to unnecessary alerts if thresholds are set too low. Others may stop functioning altogether when load increases, in which case warnings should be sent ahead of the indicator actually reaching the critical level. Which level is considered critical can be found out by tracking performance statistics over a longer time and correlating it with data about service availability and quality degradation. Administrators can also proactively set up performance tests that help find the breaking point for their infrastructure. The same tests can be run after a performance incident, making sure that any fixes that have been deployed are taking effect.