Statseeker.

The Top 5 Network Monitoring KPIs

Essential metrics for optimal network performance.

What are network KPIs?

Network key performance indicators (KPIs) are benchmarks by which optimal network performance is determined.

Tracking performance against KPIs helps network managers make proactive decisions to ensure agreed service levels are met. KPIs offer measurable data that helps the network team make informed decisions about infrastructure upgrades, performance improvements and resource allocation, based on actual needs and usage trends.

Which KPIs do you use to measure network performance?

Key performance indicators aren't just for evidencing performance levels to others; they're an invaluable tool for you, enabling you to measure and assess the quality and capability of your network.

When troubleshooting network degradation or outages, KPIs are essential for determining the root cause of network degradation, such as packet loss, saturation, bandwidth hogs, or interface and network device outages.

With the right monitoring solution, you can quickly identify and address these issues, ensuring the network runs smoothly and service levels are consistently met. Statseeker removes a lot of the manual work needed to find and compile this data. This means the network team spends less time searching for problems and more time taking decisions and making network improvements.

You need to understand which network metrics provide the essential information required for measuring network performance. While there is not a defined set of universal standards, the following five categories tend to be the most commonly used:

Device health

CPU, memory utilization and component temperatures are the foundation for proactive device health monitoring. You should be able to obtain and report on these metrics (or similar) from every device on your network.

Device health metrics can help identify hardware or firmware issues, environmental changes, and monitor equipment performance. As CPU, memory and temperature metrics are the most common device health KPIs, setting alerts on these metrics means problems can be addressed before network users notice or report an issue.

Statseeker tracks device health metrics over time and creates a baseline of regular behavior. Even the smallest deviation from the baseline can be identified using threshold rules and alerts and caught before it leads to a bigger issue.

For example, consistent air conditioning temperature readings in a data center will form part of a regular pattern of behaviour for devices. However, an overworked device that causes a rise in temperature readings can be identified as an anomaly, your team alerted, and resolved in real time.

Statseeker dashboard showing Device Overview

2

Device availability

Ideally, you want to be able to see the availability of every device on your network, in as near to real time as you can. If you only become aware that a router or a managed switch is not working when someone calls you to complain, that is too late.

Statseeker helps customers spot device outages in under a minute of the outage occurring.

- Devices are ping polled every 15 seconds to determine outages every 45 seconds.
- Port availability is SNMP polled every minute.

Outage event data is retained for as long as you need it, which means you can dig into the time devices are down, going back weeks, months, or more.

The more you know, the better informed you are about which devices experience the most outages. A device which suffers a five-minute outage every day, for example, causes four business days of impact to your team and business over the course of a year.

Outage data can be reported on by day, week, month, or year. Out-of-the-box Statseeker reports include Device Availability and Port SLA.

And as Statseeker can generate alerts as soon as a device goes down, your network team will always be the first to know about it.

Statseeker dashboard showing Device Availability Over Time

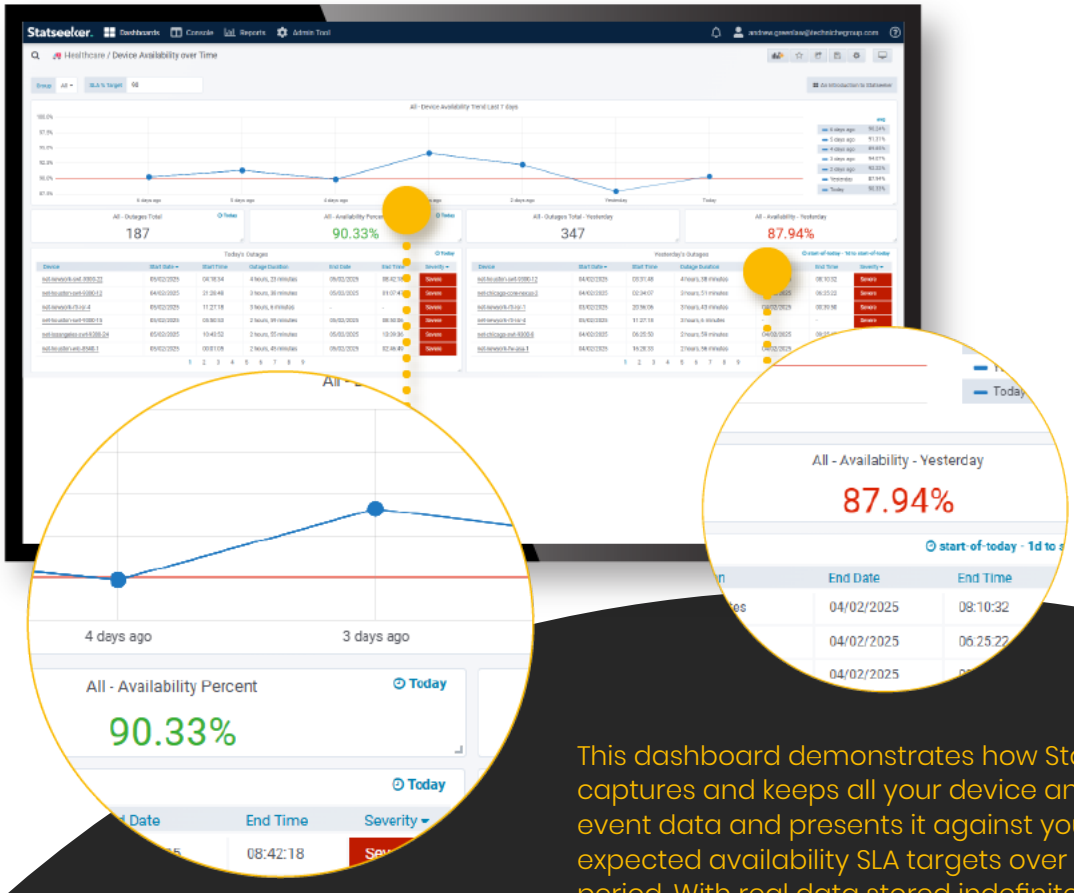

This dashboard demonstrates how Statseeker captures and keeps all your device and port event data and presents it against your expected availability SLA targets over a specified period. With real data stored indefinitely, your outages can be tracked day by day, week by week, and month by month.

3 Latency and packet loss

Today, every industry depends heavily on high quality voice and video communications. The performance of these services can be significantly affected by challenges such as latency, jitter, or packet loss.

Measuring network latency and packet loss together gives you an early warning sign and an indication of network problems. A trending increase in latency or packet loss usually affects the user experience and services delivered across the network. It's time to investigate!

Statseeker provides best-in-class historical data, which is never averaged, and which can be used to highlight changes to normal patterns.

Statseeker stores latency data to three decimal places and can expose sub millisecond RTT latency. Being able to report to this level of granular detail is essential for environments where network latency is operation-critical, such as hospitals or high-frequency trading companies.

Statseeker can also make use of Observability Appliances that are deployed for specific devices. A single device can be compared from two separate locations to view latency differences that impact the experience of that device and its users.

Statseeker anomaly detection enables you to determine whether reported data deviates from the baseline norm by analyzing the past six months of historical data, providing a clear and immediate indication of potential issues.

This feature is available in dashboards, custom reports and integrates seamlessly with existing threshold and alerting capabilities, enhancing your ability to identify and respond to anomalies efficiently.

Statseeker dashboard showing Ping RTT Comparison

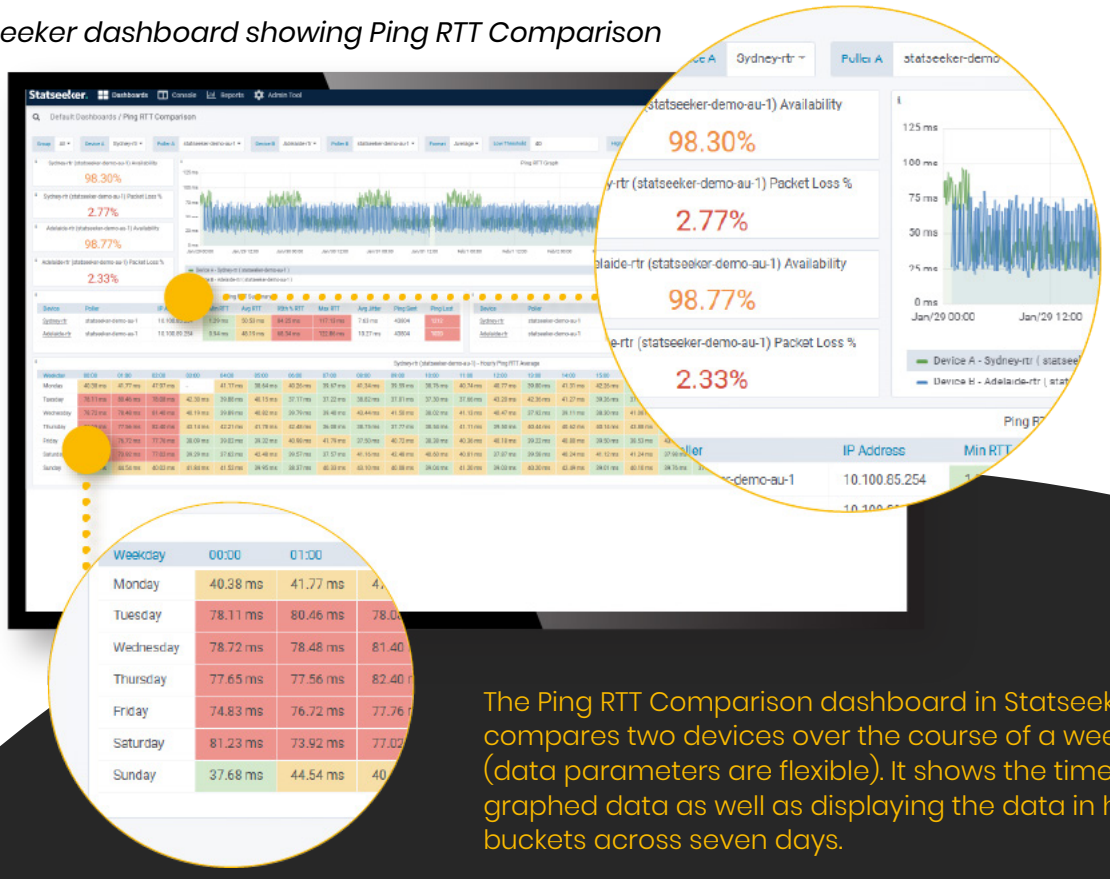

The Ping RTT Comparison dashboard in Statseeker compares two devices over the course of a week (data parameters are flexible). It shows the timeseries graphed data as well as displaying the data in hourly buckets across seven days.

4 Network interface

Using SNMP polling you can see common KPIs, such as volume of network traffic. You can also see errors and discards per interface, inbound and outbound, and you can poll an interface for availability or inactivity.

Measuring the availability and utilization of interfaces on your network gives you insights into whether these critical links are delivering the expected service level. Any downtime on an interface during core service hours could translate into revenue loss and a reduction in user productivity.

With Statseeker's built-in baseline and predictive analytic tools, you can use bandwidth data to assess the capacity of the network and predict the need for capacity upgrades.

Statseeker polls network interfaces every minute for traffic statistics, errors and discards that combine to signal the health and performance of a link. This highlights any potential degradation or saturation that is happening in real time, has been identified consistently, or is trending in a way that raises concern.

If a link is being upgraded or needs further inspection and troubleshooting, Statseeker's one second poller can be initiated for a live view of the link traffic.

Statseeker offers unrivalled bandwidth statistics, meaning that you do not have to select which links or ports to monitor; it polls every port and link in your network. Additionally, fast access to unaveraged historical data, allowing you to revisit any point in time and report on the data exactly as it occurred, makes Statseeker an exceptionally powerful tool.

Statseeker dashboard showing Interface Overview

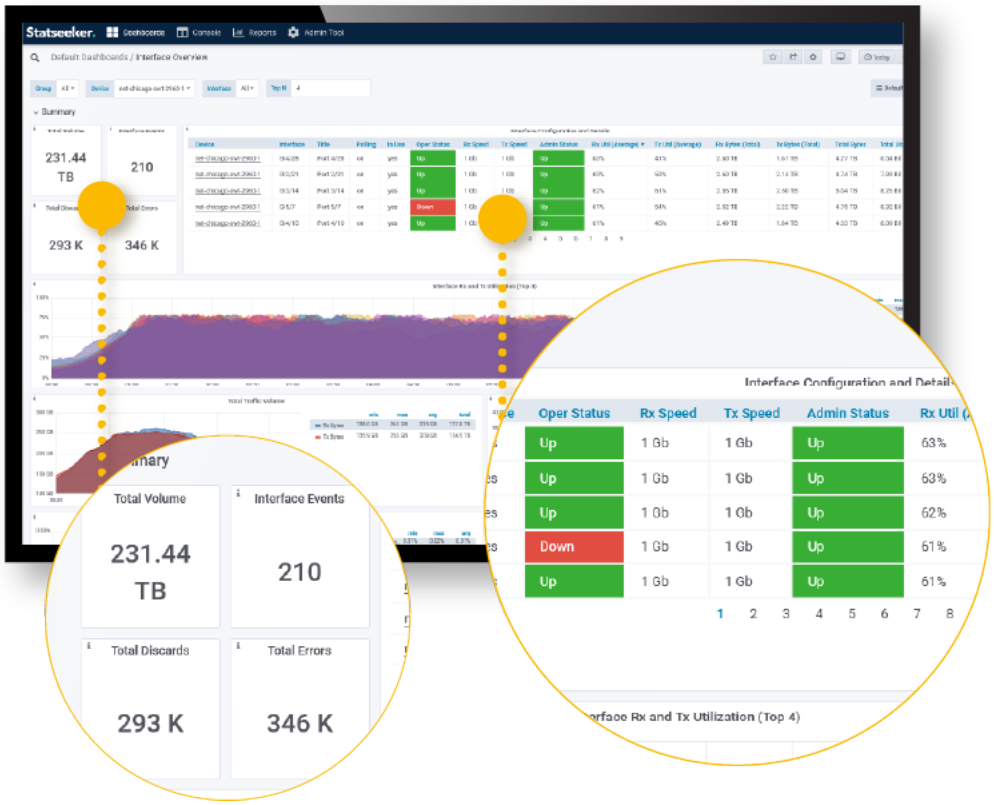

5

Device-specific metrics

Most SNMP-enabled network devices provide common key performance metrics. However, certain devices provide additional metrics for network functions which are equally important for monitoring user experience as the universal data.

For example, additional metrics can help determine whether UPS units are charging correctly, ensure load balancers are effectively distributing connections across different pools, detect unusual spikes in VPN bandwidth that may be impacting user experience or monitor firewall activity to identify potential security threats and ensure proper traffic filtering.

Being able to access and visualize this critical performance data in your NOC is key to ensuring the networking team can resolve incidents before they become outages.

With Statseeker, users can monitor device-specific metrics through a single dashboard. Statseeker brings together metrics from multiple device types through its extensive range of custom device types (CDTs), without the need to examine multiple tools to get to the right information. New CDTs are constantly added to support new hardware and vendors in the network industry.

Additionally, the Meraki for Statseeker module enhances the standard Cisco Meraki dashboard experience and provides a consolidated data view of your Meraki and non-Meraki infrastructure. Statseeker also integrates seamlessly with Cisco ACI deployed networks, collecting health and configuration data, and enabling you to monitor legacy and SDN networks through a single pane of glass.

Statseeker dashboard showing both Cisco ACI and Cisco Meraki metrics

Statseeker.

Statseeker was developed by network engineers who understood the importance of proactive monitoring. Statseeker now has active deployments in over 22 countries, and many Fortune 100 firms consider Statseeker to be an integral part of their day-to-day operations.

Start a free 30-day trial now:

statseeker.com/free-trial