
# HOW WILL the Internet of Things Disrupt Your Performance Monitoring Strategy?

Gartner has called it the most hyped technology of 2014. IDC forecasts it will represent a $7.1 trillion worldwide market in 2020. Janus Bryzek (known as "the father of sensors") has referred to it as, "the largest growth in the history of humans."

*What is it?*

The Internet of Things (IoT). Or the Internet of Everything. Or the Internet of Your Things. It all depends on the vendor in the conversation.

The IoT is a connected world of billions of IP-enabled sensors, machines, and other non-human stuff. Most often, these devices communicate between themselves and with other control systems. These machine-to-machine communications don't follow the typical unicast traffic models. Their exchanges consist of small amounts of event-driven data.

The IoT promises greater levels of visibility, agility, and control over our personal and professional lives. It's about connected cities and connected farms. Connected homes and connected businesses. Connected humans and even connected brains.

Amidst the vendor hype, staggering forecasts, and various naming conventions, a simple truth remains: The IoT has been present in enterprise and service provider infrastructures for some time, and will experience profound growth in the near term.

But, what makes the IoT relevant to network and IT teams? For starters, it poses concerns over security. A recent study by HP found that 70 percent of the most commonly used IoT devices contain security vulnerabilities. But this is only one of the challenges the IoT brings. If you're responsible for monitoring the performance of your infrastructure, here are five other key impacts to consider:

- Massive amounts of data
- New devices and protocols
- Bursty traffic
- Heavier reliance on the cloud
- Increased IPv6 deployment

By ignoring the inevitable trend toward an all-IP connected world, network and IT teams run the risk of creating visibility gaps over their infrastructure performance. These gaps put SLAs, customers, end user experience, and revenue at risk. This paper outlines the ways the IoT may disrupt your current performance monitoring process. It also offers recommendations for sensible strategies to counter those impacts.

# MASSIVE AMOUNTS OF DATA.

More than anything, the IoT represents an issue of scale. Connected devices, sensors, wearables, automobiles, and other non-human stuff will produce massive amounts of data. When viewed on a device-by-device basis, the volume of data produced doesn't warrant concern. Yet, when considered in aggregate, this data poses a significant threat to our networks and the systems that monitor them.

For example, one company has begun implanting wireless Internet-connected sensors in the ears of cattle. These sensors help farmers monitor the cattle's health and prevent the spread of disease from contaminated meat or milk. It's estimated that each cow transmits 200MB of data every year. That seems miniscule, but when you multiply that by the 1.5 billion cattle in the world, you get a different picture.

Let's take this concept from the farm to the enterprise, where a Fortune 100 company currently tracks energy consumption in their data centers. By monitoring the IP-enabled power strips that support servers, they can detect energy inefficiencies. Every day, they collect and baseline more than 25 billion performance metrics!

**A 2014 research project by IDG estimated that the amount of data managed by organizations will increase 76% within the next 12-18 months.**

Over the coming years, enterprise and service provider networks will experience a tremendous uptick in the amount of data they will handle. If your performance monitoring platform can't intuitively and cost-effectively scale with this increase in data, you risk creating a visibility gap over your infrastructure.

When the limitations of your performance monitoring platform force you to choose what you will and won't monitor, the following problems arise:

- **You monitor some things, but not others.** You're unable to predict how failure of any given device or object impacts application and service delivery across your complex architecture.
- **You monitor things with less frequency.** You're blind to activity that happens at sub-minute levels.
- **You roll up and average performance metrics over time.** You're left with inaccurate data for historical reporting and proper capacity forecasts.

The sensible approach is to build upon a performance monitoring platform engineered for speed at scale. That means abandoning products built around a centralized database architecture that will eventually fold under the weight of massive data and leave you with a product that fails to provide near real-time information about the health of your infrastructure. Instead, consider a monitoring platform based on a distributed computing model. By keeping your performance data distributed across your network, you're better equipped to handle the challenge of massive data generated by the IoT.

## NEW DEVICES AND PROTOCOLS.

When we think of IT equipment, we think of servers, routers, switches, load balancers, firewalls, and the like. We rarely think of IoT devices, sensors, and wearables as IT equipment. But they are IT equipment. Like it or not, responsibility for things that connect to your network end up in the lap of IT.

The IoT will be a major driver of data center investment going forward. In addition to the connected devices and sensors themselves, organizations will invest in larger pipes, data warehouses, more efficient compute power, and hybrid data center/cloud solutions. Cisco believes that an infrastructure optimized for IoT demands a new approach with a "strong centralized data center and robust edge."

Many of these connected IoT entities may already exist in your data center:

- Environmental management systems
- Smart racks
- Security sensors, cameras, and locks
- Power generators and distribution systems
- Uninterruptible Power Suppliest
- Building control systems

Depending on who you listen to, anywhere from 26 billion to 200 billion connected devices will be dropped onto our worldwide networks in the upcoming years. The installed base of active wireless connected devices alone already exceeds 16 billion. Analysts expect that number to more than double by 2020. It's believed that 75 percent of this growth will come from non-hub devices such as sensor nodes and accessories.

In addition to the challenge posed by the volume of data these devices generate, network and IT teams will also have to monitor the performance of the devices themselves. This will force them to consider protocols other than SNMP. Not only is it not an IoT-friendly protocol due to security concerns, vendors have begun to drift away from SNMP support in their gear, especially in the carrier space.

When gathering performance metrics from IoT devices, look for a monitoring solution that takes a data agnostic approach to collection. In other words, the monitoring platform should be able to collect any time series data, regardless of source. Also, the vendor should have a proven track record for ingesting data from other third party platforms and element management systems.

With many standards bodies (such as the AllSeen Alliance, Industrial Internet Consortium, and Open Internet Consortium) currently competing for the IoT development and connectivity frameworks, there's no telling which – or how many – standards may exist a few years from now. You need to prepare for the possibilities by building on a data agnostic monitoring platform.

## BURSTY TRAFFIC.

Imagine a world of tens of billions of devices and sensors sending quick bursts of data across our networks at both regular and sporadic intervals. This is the IoT world. It's no longer acceptable to monitor network traffic at five minute intervals. The nature of IoT traffic demands closer scrutiny if you're to understand the actual activity transpiring at any given time.

Even one minute polling cycles will not suffice. At times, you'll need to poll down to the second. The reason? Performance monitoring tools average total traffic over polling cycles. At five-minute time spans, brief spikes that often disrupt application and service performance get flattened. This is especially relevant when troubleshooting latency-sensitive services like telephony and video.

*SevOne allows you to poll your network down to the second for extremely granular views of infrastructure performance.*

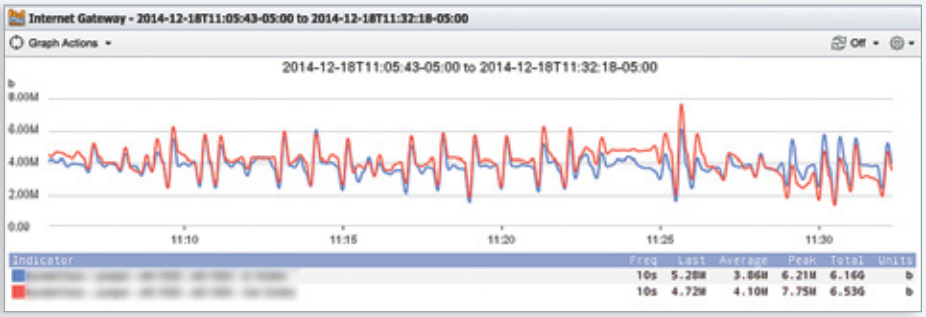


The difference between sub-minute and five minute polling cycles is akin to the difference between high definition and standard definition television. The detail you see at sub-minute levels paints a much clearer picture of what's happening in your environment. It can be impossible to troubleshoot a performance issue when looking at five minute – or even one minute – snapshots of your infrastructure.

One SevOne customer – a wireless provider—was able to show that spikes in network traffic, when viewed at 1.5 second intervals, can be more than 250 percent higher (on average) than what you see at one minute intervals.

If you haven't already, consider a monitoring platform capable of high frequency polling down to the second. While you might not always ratchet up your polling cycles to such granularity, you'll need to do so when investigating performance issues in the IoT world.

## HEAVIER RELIANCE ON THE CLOUD.

Gartner forecasts that half of large enterprises will have cloud deployments by the end of 2017. They also predict that cloud computing platforms will constitute the bulk of IT spend by 2016. Not only has the cloud become mainstream, it's poised to back the staggering demands of the IoT.

To support the billions of connected devices expected to arrive by 2020, the world would need to deploy about 340 application servers per day (or 120,000 servers per year). For companies designing and manufacturing IoT systems that require backend services, the cloud surfaces as an obvious solution.

Yet, the ongoing migration of services to the cloud poses a crucial question: How do you monitor what you no longer own? If you're utilizing Infrastructure as a Sevice (IaaS), Platform as a Service (PaaS), or Software as a Service (SaaS), there are components you no longer own. However, you're still responsible for their performance as far as they impact the applications and services your organization delivers.

*Cloud service models offer customers varying levels of control over assets and services. This presents performance monitoring challenges.*

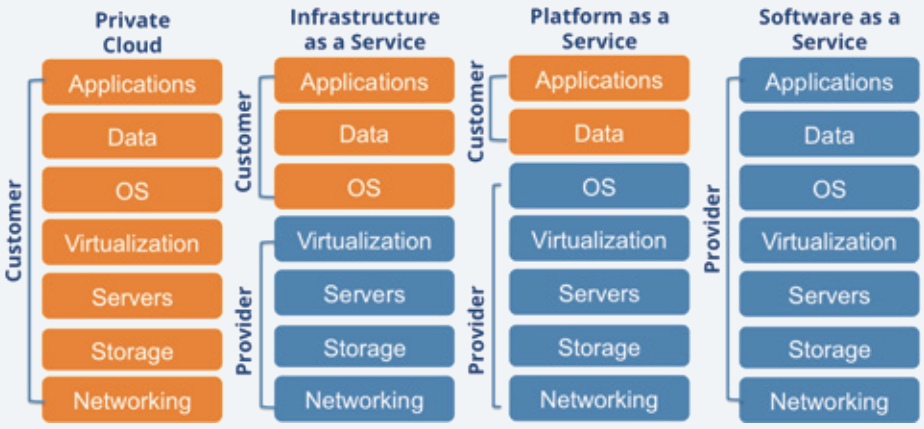


How do you determine if your performance monitoring platform is capable of reporting on the performance of cloud services? The National Institute of Standards and Technology's (NIST) working definition of cloud computing offers a foundation for discussion. They cite five essential characteristics of cloud computing:

- **Broad Network Access** – Capabilities are available over the network and accessed across a varying set of platforms or end user devices, such as mobile phones, tablets, laptops, etc.
- **Resource Pooling** – The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, whether we're talking about storage, processing, memory, or network bandwidth.
- **On-Demand Self-Service** – A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.
- **Rapid Elasticity** – You have capabilities that can be rapidly provisioned and released to scale in a manner that's commensurate with demand.
- **Measured Service** – Clouds leverage a metering capability at some level, so resource usage can be monitored, controlled, and reported, which provides transparency for both the provider and consumer of the cloud service.

Use the 5 NIST characteristics above to better understand how a performance monitoring solution addresses the challenges presented by cloud environments and services.

For example, how well does the solution handle Wi-Fi monitoring (broad network access)? Can it report on usage-based billing (measured service)? Questions like these reveal more than simply asking a vendor, "Does your product monitor performance in the cloud?"

Of course, you're also going to look for the ability to ingest performance metrics from mainstream cloud platforms such as Amazon AWS and the ability to report on hybrid infrastructure from a single performance dashboard.

(For more information, refer to SevOne's whitepaper: *A Sensible Approach to Monitoring Cloud Services*)

# INCREASED IPV6 DEPLOYMENTS.

Next to the US Federal Government, the IoT will perhaps be the largest driver of IPv6 adoption. Then again, IPv6 may be driving IoT development, since with IPv6, it's possible to use a global network to develop one's own network of smart things and interconnect them with the rest of the world.

Aside from the chicken and egg debate, it's clear the IoT is not an IPv4 world. IPv4 supports 32-bit addresses, which equates to about 4.3 billion addresses. This number has been largely exhausted by the world's connected devices. The need for tens of billions of new addresses for IoT connected devices demands a highly scalable address scheme. IPv6 has us covered, with 3.4 × 1038 possible addresses. It's an almost limitless number that can amply handle all conceivable IoT devices.

With IPv6 deployment comes the need to monitor IPv6 devices and networks. But it's not enough for a performance monitoring solution to claim itself as "IPv6 capable." The platform must support reporting of both IPv4 and IPv6 data from the same report or dashboard. It must also be capable of monitoring native IPv6 environments and not simply IPv6 data transmitted across an IPv4 connection.

When evaluating performance monitoring platforms, be sure the solution is IPv6 compliant and not just IPv6 capable, as this may cause headaches down the road. It limits your ability to troubleshoot and obtain an end-to-end view of your infrastructure health.

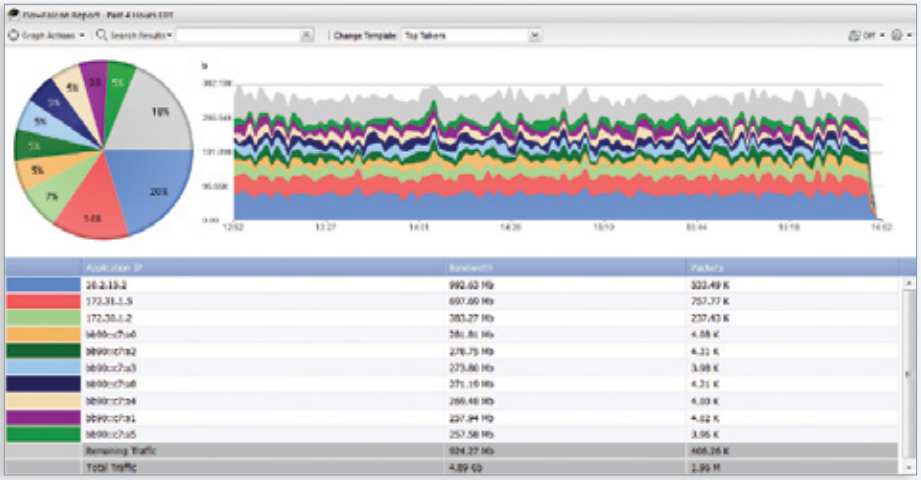

*SevOne supports NetFlow reporting across both IPv4 and IPv6 traffic within a single dashboard.*

# SUMMARY.

Many have referred to the IoT as the "third wave" of the Internet evolution. It takes us beyond the second wave of mobile access that connects approximately two billion people today. The IoT introduces a new visible world with tens of billions of connected sensors and devices.

According to a recent study by IDC, the worldwide IoT install base will see a compound annual growth rate of 17.5 percent from 2013 to 2020. It's not a question of if, but how soon the IoT will have a material impact on the global enterprise and service provider market.

Clearly, the IoT will disrupt how you currently think about your infrastructure performance monitoring needs. Questions you need to consider today include:

- Can my performance monitoring platform handle the massive increase in network traffic and still perform with speed at scale?
- Am I able to monitor new devices as they come online, regardless of the communication standard or source of performance metrics?
- Can I achieve granular visibility of network traffic down to the second?
- Am I able to monitor hybrid cloud environments with dashboards that encompass physical/virtual and cloud/on premise KPIs, all from the same screen?
- Can my performance monitoring platform render IPv4 and IPv6 metrics in the same graph?

## About SevOne.

SevOne provides the world's most scalable infrastructure performance monitoring platform to the world's most connected companies. The patented SevOne Cluster™ architecture leverages distributed computing to scale infinitely and collect millions of objects. It provides real-time reporting down to the second and provides the insight needed to prevent outages. SevOne customers include seven of today's 13 largest banks, enterprises, CSPs, MSPs and MSOs. SevOne is backed by Bain Capital Ventures. More information can be found at www.sevone.com. Follow SevOne on Twitter at @SevOneInc.