Can Machine Learning Prevent Application Downtime?

Nimble Storage uncovers the true cause of application disruptions and slowdowns through installed-based learning

Table of contents

Business users expect immediate access to data, all the time and without interruption. But reality does not always meet expectations. IT leaders must constantly perform intricate forensic work to unravel the maze of issues that impact data delivery to applications. This performance gap between the data and the application creates a bottleneck that impacts productivity and ultimately damages a business' ability to operate effectively. We term this the "app-data gap."

Not only does the app-data gap impact application performance and availability, it also forces businesses into a heavily reactive mode. In the worst case, downtime occurs resulting in a fire drill complete with all hands on deck and sleepless nights. In the best case, user complaints start a troubleshooting process that may go in circles between the storage, virtual machine (VM), networking, and application teams blaming each other. This cycle is dangerous for the business, as IT leaders have little opportunity to spend time on value-add initiatives, and IT as a whole is perceived as a barrier to business productivity rather than a key partner in increasing competitiveness.

Harness the power of Big Data

- Identify poor performance early, before users perceive an impact
- Reduce the effects of a problem
- If any one customer encounters a problem, prevent others from hitting the same issue
- Continually improve performance and availability for an entire population

End finger pointing

In order to address the app-data gap, it is important to understand its root cause. Data centers are made up of multiple hardware and software layers, including networks, servers, storage, hypervisors, operating systems, and applications. Within each layer, there can be multiple individual components from multiple vendors (for example, a single application might rely on several different databases each running on a different platform). Although all the components should be designed to work together, there is an enormous amount of complexity caused by the sheer volume of components and interactions that take place between them. This complexity, and the fact that an application's performance is limited by the slowest component or interaction between components, is the root cause of the app-data gap.

Data science unravels complexity

The traditional approach to closing the app-data gap is for IT to develop operational expertise in each piece of software and hardware, purchase monitoring tools, and employ teams of people to streamline overall operations. However, optimizing end-to-end performance has remained challenging even for the best run IT departments. An alternative approach to closing the app-data gap is to take advantage of data science and machine learning to harness Big Data gathered from thousands of sensors across every piece of the data center.

Harness the power of Big Data

• **Identify poor performance early, before users perceive an impact**—Machine learning is
  used to identify high performing or healthy environments by analyzing performance metrics
  gathered across a large population. This creates a baseline that can be customized to a
  specific environment and used to identify poor performance, automatically providing remedies
  before it becomes problematic.

• **Reduce the effects of a problem**—Correlate vast amounts of information across the
  infrastructure to detect and rapidly identify the root cause and resolve the problem.

• **If any one customer encounters a problem, prevent others from hitting the same
  issue**—Once a problem and root cause have been determined, its signature can be used to
  identify other customers who might be affected. A rule can be created to prevent the issue
  from occurring or automate its resolution if it does occur.

• **Continually improve performance and availability for an entire population**—Evolve
  software releases to optimize performance and availability based on machine learning across
  all customers. Using data science and machine learning is an extremely effective method of
  flagging potential issues and abnormal behavior, recommending steps to return an environment
  to peak health and continually improving the performance and availability of an environment.

The same data science can also deliver some in-depth insights into the cause of the app-data
gap. This report analyzes data collected by **Nimble Storage, a Hewlett Packard Enterprise
company**, and points to the most pertinent hurdles affecting the speed at which companies
access the data that powers applications.

Key findings
Storage is normally the first suspect when identifying the culprit for the app-data gap.
However, less than half of the time, the root cause of issues is directly related to the storage.

- Fifty-four percent of problems resulted from issues with configuration, interoperability,
  and not using best practices that were unrelated to storage.

- Forty-six percent of issues detected were related to storage issues including hardware
  and software issues, software update assistance, and occasionally performance issues.

Methodology

Nimble analyzed more than 12,000 anonymized cases documenting examples of app-data
gap related issues. The data was collected from a vast range of IT infrastructure across more
than 7500 customers. Nimble aggregated and analyzed this data using its predictive analytics
platform, InfoSight, which collects 30 to 70 million sensor data points a day from across the
infrastructure in which Nimble array is deployed. This data provides a comprehensive and
granular view into each infrastructure. It should be noted that in 90% of cases, issues detected
in these environments were identified and remedied by InfoSight before the customer even
recognized an issue was present.

**Top problems contributing
to the app-data gap**

1   Storage related

2   Configuration issues

3   Interoperability issues

4   Non-storage best practices
   impacting performance

5   Host, compute, VM

Figure 1. Most prevalent problems that lead to data and application delays

Causes of the app-data gap are found across the infrastructure stack

The causes of the app-data gap are not isolated to certain parts of the IT stack. Rather, the culprits are spread across the entire stack. Figure 1 breaks down the most prevalent problems that lead to data and application delays.

- **Storage-related issues (46%)**—These comprise of hardware and software issues, software update assistance, and occasionally performance issues. Examples are failed drives (predictive and proactive replacements) and automated software fault analysis with update recommendations.

- **Configuration issues (28%)**—Without predictive analytics, all configuration issues would be extremely difficult to identify and resolve.

- **Interoperability issues (11%)**—These issues tend to be related to setup configuration with Windows®, Exchange, and application-level networking. Examples include personnel not following Microsoft® SQL best practices, such as log and database volumes not being separated or MPIO setup on Windows.

- **Non-storage best practices impacting performance (8%)**—These issues can be related areas such as unaligned I/O and networking configuration, including multipathing not being set up correctly or incorrect MTU.

- **Host, compute, or VM related issues (7%)**—These are issues relating to hosts (Linux®, VMs, and so on) as well as setup configuration issues. The challenges encountered include incorrect virtual network configuration, host-side iSCSI setup, UCS setup, and under provisioned hosts.

Flash alone will not prevent downtime

These findings show that 54% of problems that can lead to the app-data gap are not due to storage issues. Many IT administrators' first instinct is to presume the storage environment is at fault and are prompted to buy faster storage. However, fast flash alone will not fix non-storage related problems. Not being able to quickly pinpoint the problem leads to a wide range of consequences, including time wasted chasing the wrong issue, extended downtime, user frustration, and missed business goals.

While 46% of the problems are storage related, the vast majority of these can be automatically identified and resolved by using predictive analytics (for example, drives that have a high probability of failure are proactively replaced by an automated procedure). Non-storage configuration issues are the next highest cause of the app-data gap, making up to 28% of the problems. This is because of individual product complexity as well as the number of different components (hardware and software) that need to work together in a typical IT environment. Each product comes with its own set of recommended best practices. These configurations or best practices often assume a specific environment, and if there is a variance from the exact assumed environment (even something as simple as a version release) the configuration may need adjustment. Further, the many permutations and combinations for an entire IT stack makes it cost prohibitive for any one vendor to track, let alone test all combinations exhaustively.

Compounding the problem is the fact that many vendors build point monitoring and troubleshooting solutions. However, these will typically only show a slice of the problem or how a problem relates to only one part of the infrastructure stack. This is where predictive analytics techniques that have visibility across the entire stack can effectively identify issues no matter where they originate.

Figure 2. InfoSight includes cross-stack analytics to rapidly root cause and resolve performance problems

Infrastructure solutions should utilize data science and machine learning

To boost performance and significantly reduce chances of downtime in the environment, companies should change how they evaluate key infrastructure products. Evaluating solutions based solely on speeds and feeds or price is no longer adequate. Nor is it sufficient to rely on traditional models of infrastructure reliability and high availability, which primarily rely on redundancy of each component but do little to ensure that all components interoperate correctly.

Companies should instead validate solutions that utilize machine learning and predictive analytics to perform the following capabilities:

• **Downtime prediction**—Infrastructure must be able to predict potential causes of slowness and downtime well before they occur.

• **Downtime automatic prevention**—Once predicted, tools should be able to prevent the adverse situation automatically through machine learning. Traditional infrastructure comes with reactive monitoring, which provides little relief other than flagging the problem.

• **Prescriptive resolution**—For the rare occasion where the infrastructure cannot automatically prevent an issue, it should lead you to a clear and prescriptive resolution—the days of looking up online forums, documentation, and calling support to help resolve issues, together with the long delays involved, are over. It leads to loss of productivity and significantly slows downtime to resolution.

• **Rapid root-cause analysis**—For rare occasions where no automatic prescription is available, it should rapidly identify the root cause so that the problem can be quickly resolved. Traditional root cause analysis involves numerous cycles of troubleshooting, problem recreation, capturing of logs, engineering analysis, and weeks of time and frustration between the customer and the vendor. In many cases, without predictive analytic capabilities, difficult problems can take weeks or months to resolve.

• **Cross-stack application of analytics**—The predictive analytics capability should include the intimate knowledge of and the ability to collect information across the infrastructure stack. If a product is not analyzing interactions across the ecosystem, it is missing out on a big part of the picture and a major cause of the app-data gap. This in turn leads to either frequent false alerts, missing early indicators of issues, or the inability to rapidly resolve issues when they do occur.

- **Analytics-driven tech support**—Ask the vendor about its support practices. Advanced analytics are able to eliminate the need for frontline, level-1 and level-2 support engineers. Frontline engineers spend most of their time documenting the issue, collecting data, and performing initial triage—all of which can be automated through predictive analytics. With the right analytics-driven support approach, and for the small percentage of problems that require the need to talk to an engineer, a customer can immediately reach a level-3 engineer who has precollected telemetry and deep knowledge on how to rapidly resolve even the most complex issue.

- **Measured availability metrics**—When evaluating a vendor, ask to see a report of their measured availability. This should not be a theoretical availability number based on a system's design, but rather it should be measured in real-world environments across an entire customer base.

Freeing IT organizations for proactive initiatives

Data science and machine learning, when used together in a **predictive analytics solution**, improve performance and availability of applications by closing the app-data gap. The benefits of predictive analytics should be integrated into the infrastructure components and should be provided at no additional cost. Employing leading edge machine learning technologies to manage the infrastructure not only makes the business more productive but also frees up the IT organization to partner with the business on high value-added initiatives.

Learn more at
hpe.com/storage/nimblestorage

Make the right purchase decision. Click here to chat with our presales specialists.

Hewlett Packard Enterprise