

INSIGHT

On-Demand Application Management in the Cloud — Turning a Traditional Model on Its Ear

Rona Shuchat

IDC OPINION

Small, medium-sized, and large enterprise customers are increasingly gravitating toward use of SaaS-based application management solutions to monitor and tune their Web applications running in private or public cloud environments. More specifically:

- ☒ New Relic, a two-year-old SaaS start-up, has been blazing its own trail as it meets demand for affordable 24 x 7 on-demand management service for production Web applications developed in Java, Ruby, or JRuby. Its recent announcements in 2Q10 — for its RPM Enterprise Edition, its integration with the Microsoft Azure platform, and its real-time thread profiling on live production Java applications — emphasize New Relic's shift to meeting enterprise demand for managing Web applications at scale. The integration of RPM with multiple cloud platforms, including Amazon Web Services (AWS), RightScale, Engine Yard, Heroku, GigaSpaces, GoGrid, and Speedyrails, exemplifies the fundamental bridge that is being built to allow customers to manage, monitor, and improve the performance of their Web applications in dedicated or public clouds.
- ☒ As application instances are being moved around or created dynamically in virtualized environments, traditional approaches to AM are giving way to SaaS solutions like RPM that can provide rapid Web access, assessment, and resolution of production-related Web application problems in real time.

IN THIS INSIGHT

This IDC Insight discusses New Relic, a SaaS start-up company that is focused on providing deeper levels of application management transparency to SMB, enterprise, and developer communities that are deploying Ruby- and Java-based production applications in virtual or dedicated environments. New Relic's RPM on-demand SaaS solution is targeted at a range of customers, including in-house development teams in large enterprises, SMBs, consultants and SIs, PaaS or cloud service providers (e.g., Heroku, GoGrid, Engine Yard), hosted infrastructure providers (e.g., Rackspace, Savvis), and full-scale outsourcers that can leverage this solution in conjunction with their other AM platform solutions.

SITUATION OVERVIEW

As the volume of Web applications designed and developed in RoR or Java grows across all sizes of business, the need to be able to test, manage, and monitor these applications with agile, accelerated operations management becomes critical. While adoption of agile/scrum development methodology is gaining traction among Java and Ruby developers, affordable application management tools have not necessarily kept pace with the rapid configuration and parallel operations management support required for these same applications. While there are several application performance management tools in the market, most are focused on assessing performance at the infrastructure layer. Although current APM dashboards typically provide visibility into server statistics, including availability and usage of database servers, application servers, and load balancers, they often fail to provide the ability for customers to quickly understand application performance behavior under load, rapidly identify the presence of inefficient code, and understand root causes for slow database response or failing transaction results.

New Relic has kept a sharp focus on developing a SaaS-based application management solution that is affordable and can be self-provisioned in minutes. It currently services more than 4,300 customers, with about 2,000 deployed in cloud configurations, private or public. Of this cloud group, approximately 50% of customers are deployed on Amazon EC2 and the other half are running in dedicated hosted environments with companies like Rackspace, Engine Yard, or Blue Box Group or in their own datacenters with use of the RPM service.

Customers extend from the very large enterprise (e.g., AT&T Interactive, CSC, Intuit, Sears, American Express Publishing, Washington Post, CBS Interactive) that are leveraging the service for their own internal applications to the very small business customer. According to New Relic, even though the majority of cloud customers were primarily smaller businesses just six months ago, the landscape has been rapidly shifting to the point where larger enterprises are now taking advantage of the cloud for new projects. With support for both Java and RoR, New Relic is seeing more enterprises looking to manage and monitor composite applications, where the front-end UI tier is done in RoR and the middle tier and back end are often Java based.

A number of New Relic's founding members and employees bring many man-years of AM software development expertise to this venture, having been involved in the evolution of Wily Technology, an ISV that was acquired by CA in 2006 for its advanced Web-based application management solution. New Relic has been backed by venture partners Trinity and Benchmark, with two rounds of funding raising \$11 million to date. The executive team includes the former founder of Wily Technology, Lewis Cirne, and the CTO, CFO, CMO, chief architect, and the business development lead from Wily.

Application Service Overview

New Relic's RPM SaaS solution has two component layers: the RPM agent and the RPM service. The agent is a small software module that is deployed within the Web application. Its purpose is to capture performance metric data and send it to the RPM service. The RPM service is hosted in New Relic's datacenter, and it provides the

aggregation, storage, correlation, and display of application performance metric data pulled from RPM agents. Using a browser, one can connect to the RPM service and monitor application performance in real time, spot and diagnose problems, or optimize an application for higher-level performance.

In addition to troubleshooting and root cause diagnosis, the RPM service offers proactive planning tools such as capacity planning, scalability analysis, database usage analysis, and Apdex scoring (an emerging standard for measuring customer experience from the application perspective). Apdex allows the IT and business owners of the application to set a threshold for the response time of the application to the end user's request.

Application behavior thresholds and use of alerts are developed using a number of criteria. Alerts may be triggered by one of the thresholds the customer has created such as an Apdex score violation or errors or response-time thresholds. The RPM system goes through a constant rescoring or reassessment of the application itself every minute as new data comes in and looks backward for a period of time and assesses what has been normal or has exceeded set thresholds. RPM performs a series of calculations that take into account response time, throughput, average Apdex score, user errors, and other common behaviors in the application and looks for anomalous behavior that will set off an alert. If the infrastructure fails downstream, for example, responses from payment gateways or from an external processing service would indicate excessively long or noncompleting transactions.

An RPM agent typically watches the transaction as it enters the application; measures its transit time from component to component; follows the transaction as it triggers calls to databases, external services, or other applications; assembles the performance data; and sends it to the RPM service. Details on background jobs that the application initiates such as credit card authorizations or email confirmations are offered.

When there are problems with a particular transaction, a stack trace can be requested, which provides an inventory of individual calls, methods, and classes for that transaction. Seeing the SQL statements behind the transactions and the identification of slow SQL queries can help a developer fine-tune the performance of an application. Ability to view performance by instance, by host, or by application cluster is available. An automatic application grouping capability allows a customer to logically combine data from multiple instances to see aggregate performance information from applications that span hosts.

RPM also monitors and tracks the number of times an application has been deployed, keeping track of who deployed it and when, helping those who manage application development to gain visibility into different elements including version control and performance of Apdex over time. Transparency into CPU and memory utilization, response time, assessment of new feature deployment, and performance in a rapid iterative manner is available.

New Relic continues to extend its application management features. In April 2010, New Relic introduced real-time thread profiling on live production Java Web applications. Historically, profiling has typically been used prior to deployment as part

of the application testing process because traditional profiling tools can introduce excessive performance overhead, which will slow the running application and impact the end-user experience. The availability of real-time thread profiling via the RPM service can help development teams find performance bottlenecks at the code level, identify flaws in the application stack (e.g., virtual machine, application server, integration middleware), reduce CPU utilization, and improve transaction response time, without the overhead drag of traditional solutions.

New Relic is also reaching out to extend its solution to support additional platforms. Its May 2010 announcement for support of Apache Solr, the open source search engine for the Lucene project (which is emerging as a new standard for Web application search), extends its value still further. These capabilities monitor Solr performance, cache hit rate and cache size, document cache, filter cache, and special tracking of index operations.

Pricing and Packaging

New Relic offers a number of package and subscription pricing options. Its RPM Lite offering is free and is targeted at start-ups, students, and nonprofit organizations. For small, medium-sized, and large enterprise customers, with larger environments and multiple application instances, higher subscription levels with more advanced feature sets are available. A paid RPM subscription improves on RPM Lite by providing longer, more detailed data retention and additional analysis capabilities such as transaction tracing, error detection, automated incident detection, scalability analysis, and custom dashboards, among other features. Bronze, silver, and gold subscription levels are available, with incremental services added. Pricing begins at \$50 per host per month for bronze, moves to \$100 per month for silver, and is currently \$150 per month for gold. Annual, monthly, or on-demand billing options are offered.

FUTURE OUTLOOK

As application instances are being moved around or created dynamically in virtualized environments, traditional approaches to AM have difficulty keeping pace. Complex APM tools have traditionally taken weeks to configure to support new applications being launched. A SaaS solution like RPM enables developers to easily access, monitor, and troubleshoot hosted applications in real time. The secret sauce of RPM is that it not only provides visibility into applications, but the agent sits lightly on the application and does not impose substantial drag or overhead on the transaction flow of the application, which can be a problem with other types of agent solutions.

With SaaS design also comes the ability to be highly responsive and in touch with customer needs. New Relic focuses on designing enhancements to its on-demand solution that are synchronized with end-user data requirements. Bug fixes, patches, and new features are pushed out on a weekly basis. To keep a close eye on end-user data needs, New Relic collects roughly 3 billion rows of application performance data daily from approximately 50,000 Java or Ruby virtual machines that supports identification of new application requirements.

New Relic, for example, shares its performance data with the core development team behind Ruby on Rails, enabling the development team to closely monitor how new features are performing as they move through the release cycle.

While RPM currently supports Java- and Ruby-based applications (support for Ruby was launched in its first year and support for Java was released in October 2009), it does not yet support .NET or PHP languages. While these languages are on its development road map, it is not yet clear which language will be targeted next. Small and medium-sized businesses and local and state government organizations are pushing for New Relic to support .NET. PHP is also very attractive due to the thousands of PHP Web applications that have been deployed to date.

Partnership Opportunities

While New Relic is partnered with a number of cloud platform vendors such as AWS, Heroku, GoGrid, and RightScale, among others, there are additional opportunities to partner or sell indirectly via other third-party providers. These include ASPs or hosted AM providers, hosted infrastructure providers, and traditional outsourcers that can leverage this SaaS solution as part of a broader set of tools in managing across customer application portfolios.

Third parties that focus on SaaS enablement of applications (e.g., Indian pure-plays and onshore global providers) can consider integrating RPM capabilities within their existing AM platforms to enhance customer-facing on-demand application support. Traditional on-premise ISVs with applications deployed globally can evaluate embedding this SaaS offering into their operational deployment solutions to accelerate support from tier 1 through tier 3 levels.

Cloud test providers like Soasta or Skytap might consider integrating with RPM for the functionality it offers in both application test and deployment scenarios, providing added intelligence into application performance metrics.

Deepest RPM integrations have occurred with cloud providers like Heroku or Engine Yard, where on-demand application performance transparency has been prioritized for its operational benefits and integrated value proposition.

Hosted infrastructure providers have been more resistant to adoption, partially because these providers have the greatest expertise at the infrastructure layer, in managing hardware, network, and storage, and are reluctant to get involved in deeper-level application management issues. Rackspace, for example, is operating more at the forefront, reflecting its own advancing cloud framework and its larger base of customers that have adopted the RPM solution. While the cloud application model implies that there are ready-built platforms to hold and rapidly deploy an application, it pushes responsibility for the operating system, the application server, the application container, and other choices on to the hosting provider. While IDC expects that some hosting providers will remain focused on the infrastructure layer, others are developing their own platform capabilities and will likely take on more responsibility for the application stack over time.

New Relic's affiliate program is also opening up opportunities for consultants, application development firms, VARs, Rails, and Java user groups to work the RPM solution into their client engagements. By providing application monitoring and troubleshooting as part of customer engagements, affiliates can target risk reduction during deployment, accelerate problem resolution, offer service-level agreements (SLAs), and focus on quality in service delivery as part of the overall client relationship.

Conclusion

While large enterprises have made significant investments in traditional APM tools, most have targeted the use of these costly tools at critical Web applications (e.g., online banking or critical Web ecommerce applications). The use of these products has been out of reach for many smaller customers due to their high cost of acquisition. While medium-sized to large enterprises continue to rationalize and improve upon their use of third-party AMS services, leveraging the RPM service in support of development/test requirements or for its core production-related AM value via third-party channels becomes an extremely attractive option. The need for visibility into application management and performance engineering to improve upon the end-user experience is only expected to grow.

The RPM service presents an alternative next-generation proposition for enterprises and third parties that are seeking greater application transparency into accelerated Web application development and deployment. Longer configuration cycles tied to more expensive on-premise AM tools have prevented large-scale adoption across all sizes of business. New Relic is betting on a highly affordable 24 x 7 SaaS AM solution that builds upon its deep AM experience.

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or Web rights.

Copyright 2010 IDC. Reproduction is forbidden unless authorized. All rights reserved.