Joyent®

Cloud Analytics in Joyent SmartDataCenter

A Joyent White Paper

Powerful Troubleshooting and System Tuning Tools that Save Money, Drive Revenue, Improve Efficiency, and Generate Happier, More Loyal Customers

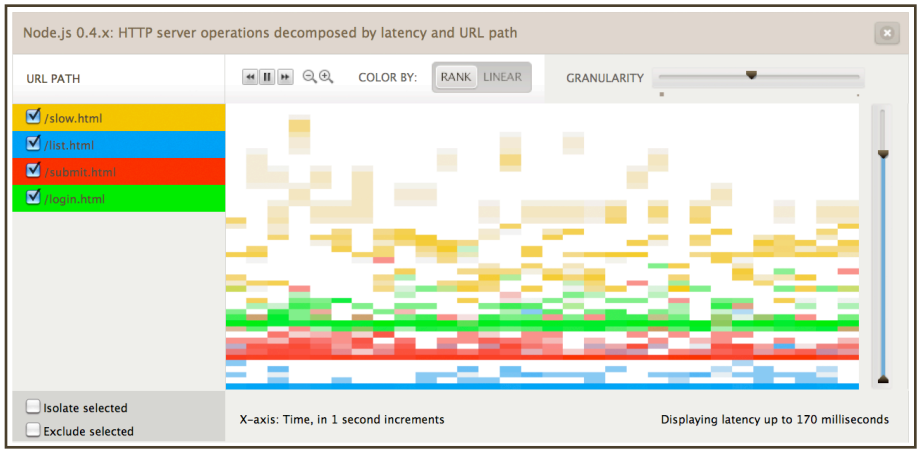

Unique visualization tools with point-and-click filtering capabilities makes it easy to spot patterns and outliers for key performance metrics like latency.

Executive Summary

For service providers running data centers, understanding performance problems can require expensive time and expertise. Complex IT architectures and application stacks make it difficult to trace root causes of problems to a line of code or a bad server node. Existing analytics packages that focus on averaging of processes rather than identifying specific process bottlenecks fail to spot outlying events that may be the cause or the foreshadow of application slowdowns that cause poor customer experience. In multi-tenant cloud environments where many more virtual machines are housed on a single server and application and network topographies change constantly, performing good analytics analysis is even more difficult.

SmartDataCenter's Joyent Cloud Analytics technology uses revolutionary visualization capabilities, customizable UIs, and "on-the-fly" instrumentation to simplify IT system and application analysis. Further, Joyent Cloud Analytics delivers a host of unique capabilities. These include the ability to visualize vertical cascading latencies at multiple levels of an application stack (e.g. CPU, disk I/O, file system, NoSQL, HTTP, and MySQL) simultaneously, the ability for operators to measure latencies in single digit milliseconds for earlier troubleshooting, and a comprehensive point-and-click UI that can both filter down to focus on specific areas of concern quickly or allow operators to click on a single point of latency and drop immediately into the responsible process anywhere from the CPU to the kernel to application. Operators may use Joyent Cloud Analytics to visualize sources of latency across their entire virtual data center or only focus on specific portions or even on a single customer's Joyent SmartMachine.

Joyent Cloud Analytics can help service providers reduce time spent troubleshooting complex issues from days or hours to minutes. In root cause analysis and latency analytics, the vast majority of time is spent locating the source of the problem. Joyent Cloud Analytics allows service providers and their customers to dramatically reduce the time spent locating the subsystem issue. This ability leads to happier customers, faster resolutions of support tickets, elimination of finger pointing between service providers and customers as to the source of problems, and, ultimately, empowers customers to do their own troubleshooting, reducing the operating expenses both in terms of headcount and man hours spent by the customers and service providers. Joyent Cloud Analytics has further uses in SLA verification, application and performance tuning, and security. The following White Paper discusses Joyent Cloud Analytics and its use cases.

Real-Time Analytics Are Hard

Today's Data Center and Application Stack Are More Complicated Than Ever

Tracing problems back to the source in today's typical Internet-driven application IT infrastructure is like finding a needle in a digital haystack. There are simply too many moving parts. Devices, device drivers, file systems, storage systems, compute nodes, OS, middleware, firewalls, multiple Web servers, load balancers and databases all contribute to an environment of dizzying complexity. This is partly because analytical tools have not kept pace with rising architectural complexity. Finding the corrupt database table that is spewing out server queries is extremely time consuming when the traffic from that table is buried in the white noise of data coursing through the data center architecture or web-based application stack.

For most service providers, analyzing and identifying what has gone wrong within a server rack, a database, or a customer's application stack can take hours, if not days or weeks. In the cloud, real-time analytics and troubleshooting is even more complex. **In cloud environments, virtual servers mean dynamic IP addresses, data center locations, and zones within data centers as customers power up and then power down virtual architectures as part of on-demand computing strategies.** Hardware failures could occur in a wide variety of geographical locations. Identifying when and where a customer is on a server is a tricky task unto itself, let alone troubleshooting that server.

For service providers and their end customers, Joyent Cloud Analytics technology offers a new visual analysis paradigm for troubleshooting, operational analytics, application tuning and security.

The Joyent Analytics Difference

Until now, data center operators have relied on external software packages from companies like Splunk, Exceptional or SysLog for data

center analytics and “dashboards” to help spot sources of latency. Latency is the most critical concern for cloud and virtual data center operators as well as for their customers. Web application performance and end customer user experience are entirely dependent on reducing latency. What’s more, operators of public (and private) clouds face new challenges in providing analytics both for internal service provider purposes and to customers running their own mission-critical applications. Both operators and customers need to understand performance metrics in order to manage their businesses and properly architect their current and future IT footprints. Moreover, extremely detailed analytics are essential for proper application tuning to ensure that cloud delivery of content and application information to end-users is rapid and responsive.

Joyent’s Node Control offers a comprehensible visual overview for users of the No.de service.

Joyent built Cloud Analytics specifically to address these needs and to bring a new visual paradigm that represents the complexities of existing cloud analytics tools and to empower data center operators and customers alike to perform far more detailed analytics using simple point-and-click heat maps and GUI-driven filters. As compared to

Syslog, Splunk and other log-based analytics tools, Joyent Cloud Analytics allows for far easier and faster root cause analysis of sources of latency. Further, by creating an easy way to extend these same visual analysis capabilities to end customers, Joyent has created a true differentiator for cloud service operators and taken what formerly was a costly business problem and turned it into a valuable business solution.

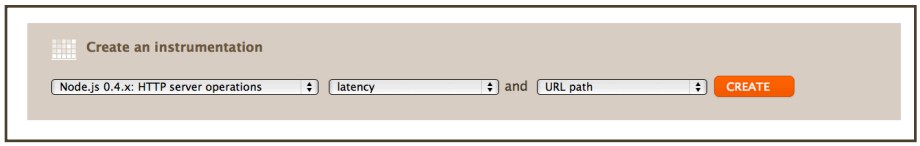

The browser-based interface allows for easy creation of instrumentations to analyze key performance metrics.

The visual presentation of latency heat maps in Joyent Cloud Analytics, combined with drop-down and check-box filters, can quickly instrument a view showing any combination of over 70,000 measured processes. Operators can identify a slow process represented as a color-coded pixel block on the screen, click on that pixel block, and instantly drop down into the specific process that is the source of the latency visualized in the block. This granularity is unmatched by any other analytics package. What's more, Joyent Cloud Analytics is the only analytics software that can measure latencies down into the single-digit milliseconds. This granularity allows customers to further cut latency in already fast applications or to identify trending latencies before they impact end users or slow application performance.

What's Different About the Cloud?

Lots. First, real-time analytics for the cloud must be far more flexible in a heavily virtualized environment because customers are lighting up and then shutting down instances all the time. Building analytics that are primarily focused on hardware or even static virtual server spaces does not work in a dynamic cloud environment because customers' IT footprints change on a near daily business. For true Cloud Analytics to work well, the end customer and their application stack must be the focus of the analytics package, rather than the data center or the virtual server. In a similar vein, a service operator will need far more fluid

metrics for their own internal uses because the size and nature of accounts may fluctuate dramatically, even over the course of a single day as customers add then diminish capacity in brief or sustained bursts to accommodate normal use cycles of their cloud infrastructure.

For example, an eCommerce provider focused on the U.S. may cycle down considerable capacity when North America is asleep, or the launch of a game company's new online multi-player game hosted in the virtual data center may go "viral" and require more processing power.

Cut Through the Fog of Complexity With Ubiquitous Transparency

To attain the level of granular visibility and tracking capabilities required for a true customer-centric view, and to easily analyze and visualize the rapid changes in virtual infrastructure, Joyent built an entirely innovative Cloud Analytics interface and flexible layer of abstraction. This interface leverages DTrace, a dynamic framework for troubleshooting kernel and application problems in production environments. DTrace is an open source technology (and recipient of a Wall Street Journal Technology Innovation Award) that works at any level in the application stack: from the kernel to the web server to the database to the scripting framework to the storage software. DTrace can also provide unparalleled horizontal analytics coverage, visualizing operations across a service provider's entire virtual data center while visualizing compute processes across a customer's entire cloud infrastructure. DTrace exerts negligible system drag and uses minimal system resources (less than 1%).

Joyent SmartDatacenter also offers DTrace API hooks for customers or service providers who want to design their own analytics. Operators can blend Joyent Cloud Analytics with other analytics output to construct overlapping heat maps that generate color-coded or chart-based representations of core processes and functions inside a data center, giving a richly detailed yet united view of many different systems. The key is rendering all the disparate flows of data into visualizations, which

tap into the human brain's powerful ability to spot patterns in visual information.

Visualize, then Analyze

The visualization capabilities of Joyent Cloud Analytics offer many relevant use cases. Analytics packages have traditionally reported measurements averaged over a time period. This means that outlying data points - which may be signs of trouble - are often buried in the noise, making troubleshooting far more difficult and time consuming. This problem has become particularly pronounced in the cloud era. A customer may be running any number of application packages, as well as their own code, on top of a virtual OS or other datasets and machine images offered by PaaS and IaaS providers.

A moving line or bar chart only displays the moving average. Averaging can hide problems and sources of latency. A single slow process may not greatly impact an average chart but it could tremendously impact users if that process is relevant to eCommerce transactions on a site, for example. In contrast, the pixel-based heat maps that visualize every process and displayed by Joyent Cloud Analytics makes it simple to create rolling system maps of astounding granularity, clearly calling out outlying points of latency. With this capability operators and customers cannot only perform trouble shooting but also:

- Finely tune applications by tweaking parameters and measuring before and after performance down to the single-digit milliseconds
- Identify mounting sources of latency before they become obvious to end users, allowing the operator or customer to either spin up additional compute resources and move users or to quickly troubleshoot the problem
- Verify whether the source of latency is the operator (hardware, virtual OS layer) or the customer (application, runtime environment)
- Definitively verify SLA compliance or breaches for virtually any parameter
- Quickly identify cascading problems resulting from a security breach
- Identify what breaks in an application stack under load test conditions during pre-production shakedowns.

Analytics as Revenue Augmentation and Cost Cutting Tool

Traditionally, analytics are used to make long-term business decisions and plan for resource allocations. Data center operators have used analytics for troubleshooting, but the sheer output of data coming from many analytics packages has actually increased the need for headcount rather than reduce it. So, CIOs or data center operators have rarely viewed analytics packages as a mechanism to cut operating expenses, reduce headcount or drive revenue. With Joyent Cloud Analytics, operators can actually stand the traditional paradigm on its head by converting analytics from a cost center to a service differentiator, account management and cost reduction tool.

Service providers may offer Joyent Cloud Analytics, its granular UI control, and incredible real time insight, as a feature that differentiates their cloud environment. In a world where cloud environments are rapidly multiplying, true differentiation that offers hard-to-replicate yet highly valuable features represents a powerful sales and marketing tool. On the cost reduction side, Joyent Cloud Analytics should significantly reduce the number of support cases as many clients problem solve their latency issues without engaging the service provider. In addition, service providers will spend fewer hours per support case, allowing for a smaller support team to service a greater number of customers.

Joyent Cloud Analytics is also an excellent account management tool, enabling efficient and meaningful collaboration between service provider and client as they drive towards rapid problem resolution.

Advanced instrumentation and telemetry, combined with monitoring and alerts, set the conditions for significant improvements in Quality of Service and resource utilization. Finely tuned applications run on less hardware. Customers with screaming fast performance are happier and more loyal. Plus customers utilizing Joyent Cloud Analytics are less likely to take up either support or excess hardware resources. The ultimate promise of Joyent Cloud Analytics is: faster resolution of customer problems, happier customers, lower support costs and better overall application performance.

SmartDataCenter's Joyent Cloud Analytics

Transparency in Complex Systems that Reduces Costs, Improves Service and Customer Experience

That virtually every organization will deploy cloud computing is now a forgone conclusion. According to a recent survey of 1,200 IT professionals by computer hardware and software distributor CDW Corporation, 84% of organizations surveyed had deployed at least one application in the cloud and 28% indicated they were actively using cloud computing to power parts of their IT infrastructure. More and more service providers are moving quickly to create clouds to take advantage of this fast growing market. An opaque cloud that is hard to analyze and troubleshoot, however, can be a serious time drain and cost center, not to mention a flash point for customer anger (remember: failures are always the service provider's fault). By creating a flexible, scalable, ubiquitous analytics system, Joyent SmartDataCenter and Joyent Cloud Analytics provides the most transparent cloud on the market that allows both service providers and their customers to know everything. This, in turn, gives services providers a powerful tool for marketing, maintaining and operating their cloud.

Please visit www.Joyent.com for more information.