

Technical Disclosure Commons

Defensive Publications Series

May 10, 2018

DIGITAL CAMPAIGN TRACKING WITH RESPECT TO APPLICATION INFRASTRUCTURE

Piyush Mathur
Cisco Systems, Inc.

Bhupinder Singh
Cisco Systems, Inc.

Siddhartha Dutta
Cisco Systems, Inc.

Arpita Agarwal
Cisco Systems, Inc.

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

Mathur, Piyush; Singh, Bhupinder; Dutta, Siddhartha; and Agarwal, Arpita, "DIGITAL CAMPAIGN TRACKING WITH RESPECT TO APPLICATION INFRASTRUCTURE", Technical Disclosure Commons, (May 10, 2018)
https://www.tdcommons.org/dpubs_series/1190



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

DIGITAL CAMPAIGN TRACKING WITH RESPECT TO APPLICATION INFRASTRUCTURE

AUTHORS:

Piyush Mathur
Bhupinder Singh
Siddhartha Dutta
Arpita Agarwal

CISCO SYSTEMS, INC.

ABSTRACT

Digital campaigns are often derailed from their goals due to amiss or often miscalculated aspects of the impact of the campaign on application infrastructure. To mitigate this problem, an approach is described herein which unifies these otherwise disjoint domains of digital marketing and application monitoring analytics. A solution is provided whereby a user may perform capacity planning for upcoming campaigns and perform real time correlation on campaign and infrastructure Key Performance Indicators (KPIs), thus maximizing revenue and optimizing infrastructure cost. Through machine learning algorithms which are running on top of correlated data from both the domains, a plethora of actionable and valuable insights may be unraveled.

DETAILED DESCRIPTION

In today's age of technology, companies drive engagement, conversions, traffic, and revenue by running many digital marketing campaigns. The success of these campaigns is of utmost importance to the firm as it ties in with the overarching goals of the organization.

Although multiple factors can contribute to the derailment of a campaign from its goals, one oft-ignored factor is how these campaigns affect the infrastructure health of the application. These campaigns usually introduce a spike in application usage and reveal application infrastructure to be under-provisioned.

To mitigate the adverse effect of the campaign to an application, firms currently deploy an ad-hoc process. This approach involves capturing key data points of user traffic from campaigns and key performance metrics from an Application Performance Manager

(APM), and attempting to correlate the two. This is limited because the whole process is usually manual, accounts for a limited historical data set, is non-adaptive to continuous changes in campaign strategy or application infrastructure, and is mostly pivoted on prior human experience.

Figure 1 below provides examples illustrating these problems.

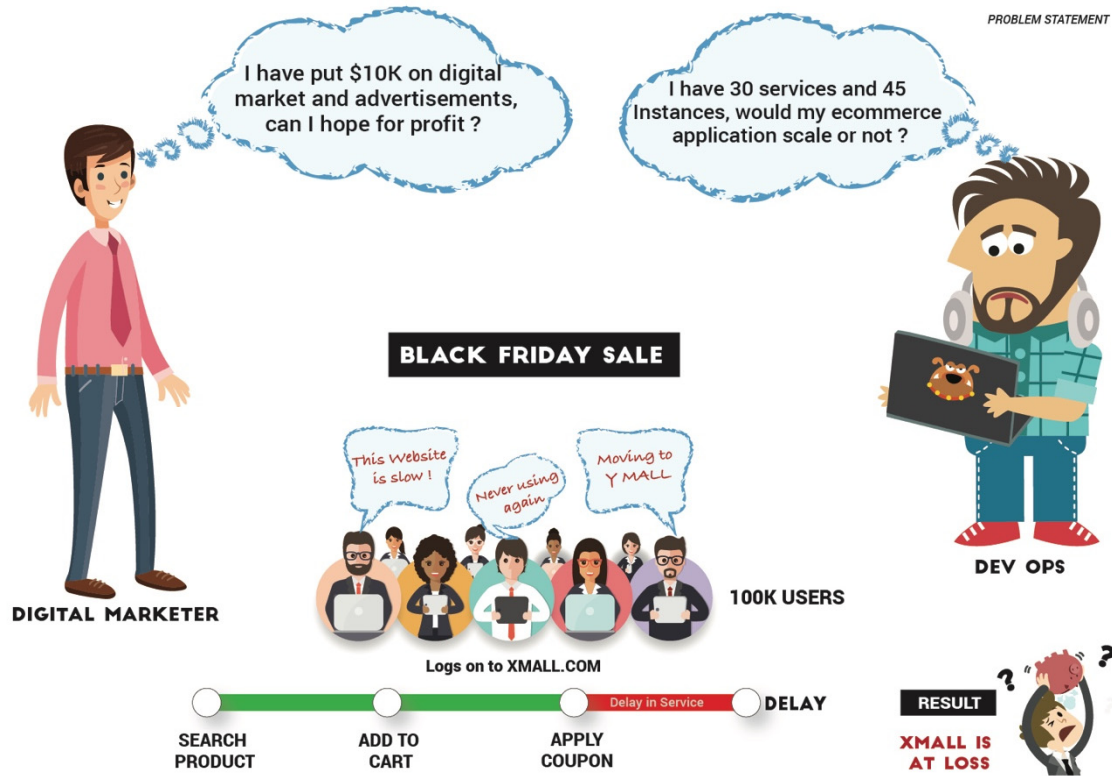


Figure 1

Once all these data points are collected, these two domains are mapped together and machine learning algorithms are applied to unravel a plethora of actionable insights, as illustrated in Figure 3 below.

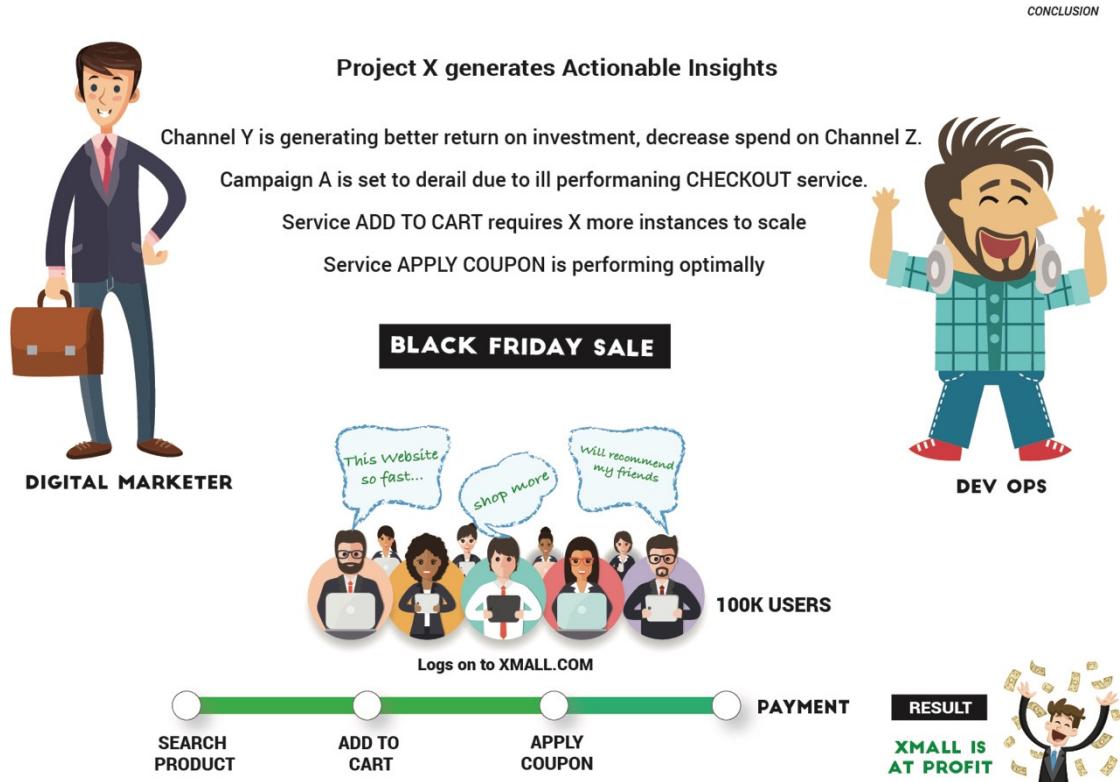


Figure 3

A robust, automated, end-to-end data collection process is inevitable before any correlation can be performed at scale.

The campaign’s data points may be collected from multiple direct sources, thereby allowing a user to analyze whether the campaign is on track and, if not, take corrective measures. Metadata regarding campaigns includes budget, reach, and spend data points for a campaign, which can be collected directly from the source where these campaign are running by tapping into their respective Application Programming Interfaces (APIs). Real-time data from campaigns includes how many users were brought in by a campaign. The behavior of these users is further tracked using beacons on the client side of the application. Key data points collected from here include revenue generated by these users, average time spent by these users per application page and drop off pages, etc.

Metadata regarding application infrastructure may also be collected, and services in the application may be discovered. This can be performed using a custom/third party APM tool. Benchmarks for all the services against key performance metrics may be provided by the dev-ops team initially and adjusted over time by machine learning algorithms. Key performance metrics include call per minute and average response time, and should be reported for each service.

To correlate data points in an end-to-end fashion, they must be glued together. This is achieved by embedding a unique Global Unique Identifier (GUID) in the link marketed by a campaign. This marks the starting point of a user journey. Once the user clicks on the link and lands in the application, this GUID is captured and maintained throughout the user session. Whenever a user performs an action, this GUID is also sent to the backend, which reports it to the APM as a metric. Thus, the load on a campaign may be grouped per application service per campaign, allowing the impact of campaign on the application infrastructure to be mapped.

The solution is designed to correlate and apply machine learning algorithms real-time to generate actionable insights. As user traffic data on a campaign is mapped, the load from a campaign or group of campaigns may be directly calculated. One key metric that may be captured throughout a user's journey in an application is revenue generated by that user. It is known from which campaign a user is brought in, and by correlating these two items and the spending of the campaign, this may be obtained as part of the metadata of the campaign. This can drive profit per campaign per channel.

Machine learning correlation and insights may enable identification of the best performing channel. This powerful insight may be provided by correlating historical data points of the number of users a channel is usually able to bring in by spending a certain amount of money. This insight can further be used to identify the channel that generated the most revenue with the allocated budget. Machine learning may also enable a cost versus benefit analysis. Often, the cost of the additional infrastructure required to support the additional load of the user generated by the campaign is ignored. This insight helps a user identify this cost and map it back to the profit calculation of a campaign.

By correlating the load per service and service benchmarks, infrastructure insight regarding the number of instances that would be required to keep up with the additional

load by campaigns may be provided. In addition, combining the data points specific to a campaign (e.g., return on investment), user traffic and other such metrics which are of importance to campaign goals, and key performance metrics of infrastructure (i.e., how the campaign is affecting the infrastructure), may provide a health number for the campaign and generate an alert when the health falls. When the health of a campaign falls below a certain threshold due to the inability of the infrastructure to scale up to an additional load, an alert may be generated and the user may immediately stop spending on the derailed campaign until the situation improves.

For predictive correlation and analysis, a set of mapped data points may be collected and run through machine learning algorithms. This allows for capacity planning to account for the spike in load from upcoming campaigns and the infrastructure needed to support it. The historical load may be analyzed by generating similar campaigns and giving the benchmark for services in the application which can predict the number of instances of a particular service that would be required. Similar campaigns may be identified automatically based on data points for channel, budget, name, time of year, etc.

It may be also be predicted that if the spending is increased to X, then Y additional users will be brought to the application. If the application infrastructure is ill provisioned, it is bound to have poor performance.

As part of user journey tracking, data points may be collected based on the user demographic. This may be correlated to products a user looks at or buys. From this, it may be inferred which user demographic tends to buy certain products. Factoring this information point into campaign planning could be of great use. The demographic data may also be used to predict cloud spots required for infrastructure scaling.

With automated data collection, mapping, and correlation at scale, the process of managing a campaign may be simplified. Through this data collection and mapping process, actionable insights are provided in real time and key data points are identified. These data points can be further run through various machine learning algorithms to introduce a predictive nature to this unified picture.

Vertical business benefits from improved machine learning models for better insights, improved advertisement targeting based on the insights, and automated budget allotment and campaign halting. Vertical APM also benefits from proven machine learning

models for better insights, as well as service instance scaling on the basis of geographical load and intuitive failure analysis.

Digital marketers can often become frustrated because their systems cannot scale due to the additional load generated by big budget campaigns. Enterprises can take significant losses due to application performance issues (e.g., users unable to apply a coupon code during checkout because of coupon service performance issues). In one real-world example, a large e-commerce company launched a mega sale after investing a large amount of money on advertisements and social media campaigns. This campaign produced large amounts of network traffic but due to poor infrastructure management, the site crashed, resulting in loss of revenue and brand image.

As described herein, this real-life problem is addressed using APM. Capacity planning helps the dev-ops team plan the infrastructure based on estimated traffic for upcoming campaigns. The techniques described herein provide real-time tracking of every dollar spent on online campaigns on various digital marketing platforms and the traffic each marketing platform is able to bring to the application. This helps estimate revenue generated by a particular campaign and also provides key metrics to estimate how many extra instances a service (e.g., apply coupon or checkout) requires in order to handle the extra load. While tracking the flow, it also provides real-time insights to the campaign manager to manage the budget across various marketing platforms to obtain maximum revenue, and to the dev-ops team to optimize infrastructure cost.

Figure 4 below illustrates an example dashboard tailored to provide an overall picture of all campaigns and how the infrastructure and services are responding to campaign traffic. This focuses on showcasing metrics such as dollars invested, the number of running campaigns, the health of campaigns, and their respective growths along with overall revenue generated, both monetary and footfall.

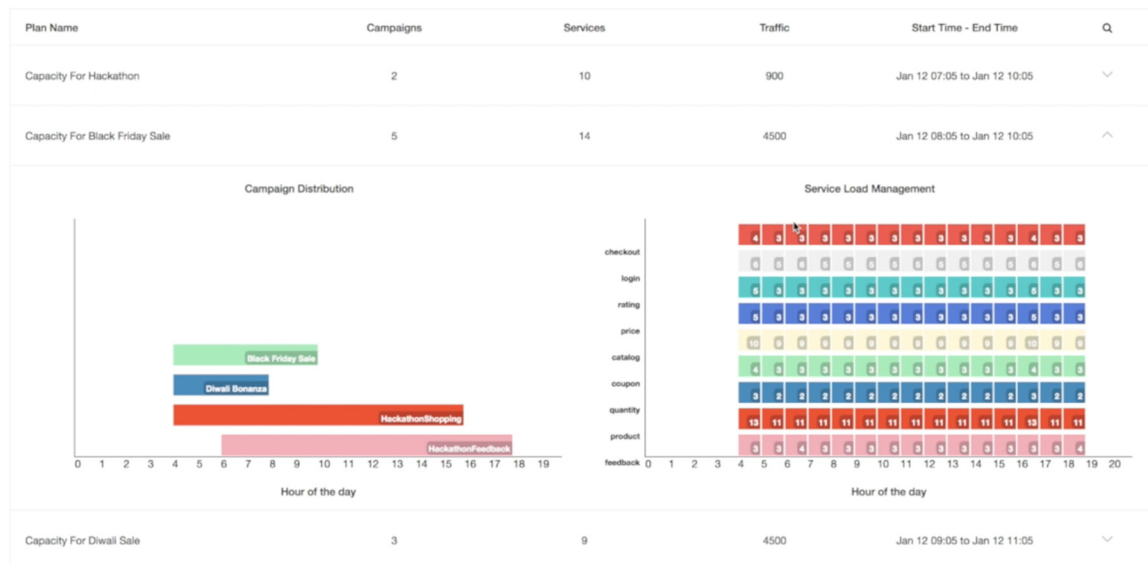


Figure 4

The GUID for tracking is passed into each service as a header, creating a path which is used to understand user behavior. Machine learning algorithms are executed to generate real-time insights for optimizing the budget spent strategy across various digital marketing platforms. To prepare for extra traffic, a nascent version of capacity planning can be used to predict the number of instances per service required over a particular interval of time. The load management widget is useful for providing the dev-ops team with insights for performing infrastructure planning for the future. This may be further enhanced by integrating with various cloud APIs to scale the system automatically in the future.

In 2017, over \$35 billion were spent on global social media advertising, and a 20% increase in anticipated per year by 2019. The techniques provided herein enable marketers as well as dev-ops teams to scale their campaigns accordingly.

In summary, an approach is described herein which unifies the otherwise disjoint domains of digital marketing and application monitoring analytics. A solution is provided whereby a user may perform capacity planning for upcoming campaigns and perform real time correlation on campaign and infrastructure Key Performance Indicators (KPIs), thus

maximizing revenue and optimizing infrastructure cost. Through machine learning algorithms which are running on top of correlated data from both the domains, a plethora of actionable and valuable insights may be unraveled.