Technical Disclosure Commons

January 02, 2018

Identifying Outages In User-Facing Products Based On Search Statistics

David Coles

Yehuda Karlinsky

Gil Ran

Shlomo Urbach

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

IDENTIFYING OUTAGES IN USER-FACING PRODUCTS BASED ON SEARCH STATISTICS

ABSTRACT

A method is disclosed for identifying outages in user-facing products based on search statistics. The method includes collecting search queries related to a user-facing product from a search engine. The method further includes generating statistics based on the collected search queries. The statistics may indicate spikes based on the product and failure-related terms associated therewith. Based on the statistics, the spikes for searches related to outage of the product are identified. Further, on identifying the spikes, alerts may be issued to an appropriate technical team related to the outage.

BACKGROUND

Companies running web services invest a lot in monitoring their products to quickly identify problems. In particular, the companies spend a lot of effort trying to determine if the services are available for their users, for example, they may try to monitor as close to the user as possible. While companies would want to measure availability directly from the users, client-side monitoring poses many challenges. In some instances, serving problems might escape the monitoring tools for many reasons, such as: monitoring the wrong thing, monitoring is broken, not having monitoring capability, service is fine but there is a problem between the user and the product, returning a result that looks fine but is actually broken, e.g., serving an empty page. In other instances, the problem manifests in a subset of use cases, which aren't covered by the monitoring tools.

DESCRIPTION

A method is disclosed for identifying outages in user-facing products based on search statistics, as illustrated in FIG. 1. The method includes collecting search queries related to the user-facing products from a search engine, at block 102. For instance, most major problems in a web service or product may encourage users to search regarding the outage of the product in search engines. A company may collect search queries related to a product X from the search engine. The search queries may include product X and failure-related terms, such as "product X down", "product X outage", "product X error", "product X failure", "X not working", or "X downtime".

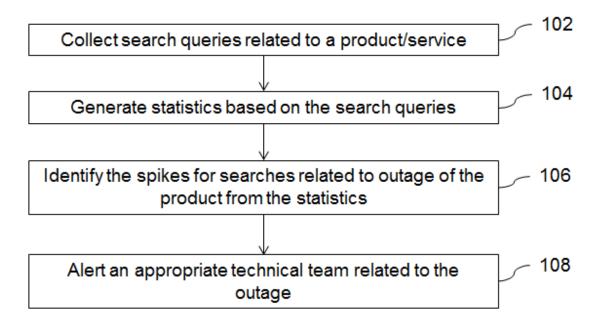|	
Collect search queries related to a product/service	— 102
↓	
Generate statistics based on the search queries	— 104
↓	
Identify the spikes for searches related to outage of the product from the statistics	— 106
↓	
Alert an appropriate technical team related to the outage	— 108

FIG. 1: Method of identifying outages in a user-facing product based on search statistics

The method further includes generating statistics based on the collected search queries, at block 104. The statistics generated may indicate spikes based on the product and the failure-related terms. Based on the statistics, the spikes for searches related to outage of the product are identified, at block 106. If search statistics are available within minutes, these spikes may also be identified within minutes. Any well-known algorithm for identifying a spike in a timeline series

may be used. Further, on identifying the spikes, alerts may be issued to an appropriate technical team related to the outage for necessary action, at block 108.

In one aspect, the spikes at block 106 may be filtered based on specific feature of the product such as regional basis, or language. In another aspect, language differences may be overcome by having the search-statistics collection system map search terms into "entities" (e.g. KG entities), thus allowing the spike to be for an <entity, failure term>, instead of <search term, failure term>. In some aspects, social media networks may also provide insight into a user's negative experience of a product, by searching for "product X down" in public conversations.

The above method eliminates the requirement to know what type of errors to expect as one needs to think of the different keywords users might use to search for it. This method could work for any product with a critical remote component, for example, mobile apps, cloud-based services, cable/set-top boxes, ISPs, mobile phone providers. The techniques used herein may be applied using any search engine stats or data scraped from other public web facility, which show how often a particular search-term is queried.