RC25377 (WAT1304-077) April 26, 2013
Computer Science

IBM Research Report

Application Management Services Analytics

Ying Li, Ta-Hsin Li, Rong Liu, Jeaha Yang, Juhnyoung Lee
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 208
Yorktown Heights, NY 10598
USA

Research Division
Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

Application Management Services Analytics

Ying Li, Ta-Hsin Li, Rong Liu, Jeaha Yang and Juhnyoung Lee
IBM T. J. Watson Research Center
Yorktown Heights, NY, U.S.A.
{yingli, thl, rliu, jeaha, jyl}@us.ibm.com

***Abstract*— Enterprises often maintain many IT applications to support their business. Application Management Services (AMS) aim to maintain high levels of service quality and availability by restoring normal application service operations and minimizing negative business impact. In this paper, we present the AMS Analytics System for improving the productivity and quality of delivery for AMS practices. Issues regarding IT applications are formally referred as IT incidents or tickets, which are an important vehicle for measuring quality of AMS. IT incident ticket analytics, an important component of of the analytics system, measures workload variability, resource productivity and delivery performance using algorithms from statistics, queuing theory, data clustering and signal processing. The AMS Analytics System provides a standardized, integrated analytics platform supporting AMS delivery. It is built on a Web platform using a set of standard open stack software, enhanced with advanced analytics. Since its initial release, we have applied the AMS Analytics System to several dozens of real-world enterprise users, receiving very positive feedback.**

Keywords— IT incidents; IT incident tickets; IT incident analysis; Application management services; incident management

I. INTRODUCTION

Application Management Services (AMS) is a type of IT services offerings that brings together an industrialized, globally integrated approach to help companies strategically manage their application portfolios in support of business goals. It includes application implementation, application help desk, application maintenance and support, and application monitoring. The objective of the application management process is to restore a normal operation of applications as quickly as possible and to minimize the negative impact on business operations, thus ensuring that the best possible levels of service quality and availability are main-tained [5]. 'Normal service operation' is often defined as service operation within a service-level agreement (SLA). An incident (a.k.a ticket) is any event which is not part of the standard operation of a service and which causes, or may cause, an interruption to or a reduction in, the quality of that service [3]. Incidents are the result of failures or errors in the IT infrastructure.

Application maintenance cost is a significant portion of the IT portfolio budget, and organizations struggle to reduce the IT support burden. Application support teams are faced with year after year budget reductions and ever increasing complexity of the portfolio. Historically, there has been very little analytics based information by which to drive solutions to these challenges. New analytics-based application maintenance system must enable the account resources to make well informed decisions. It should help the support team to improve the productivity, quality of delivery and enables better solutioning for maintenance and support services. It should provide analysis of incidents, problem records, service requests, supply and demand data to evaluate where and how effort is spent, and provide implementable recommendations for improved service performance and more effective resource utilization.

In this paper, we present the AMS Analytics System which was designed to improve the productivity and quality of delivery of the Application Management Service practitioners. The teams can seamlessly leverage the capabilities of the analytics platform covering four major areas:

• *Incident Ticket Analytics* which helps AMS delivery by measuring workload variability, resource productivity and delivery performance using algorithms from statistics, queuing theory, data clustering and signal processing.

• *Investigative Analytics* that uses stochastic analysis and machine learning algorithms to uncover root causes behind spikes in ticket volumes and changes in delivery performance.

• *Knowledge Management* that mines historical resolution information from ticket data and captures knowledge from experts, and then enables the reuse of resolution knowledge in the form of searchable database or interactive dialogs in practice [1].

• *Resource Management* that helps plan resource needs within and across accounts for enhancement and support work by analyzing demand pipelines, ticket data and resource availability.

Among the areas, this paper focuses on the IT Incident Ticket Analytics aspect of the system in particular. In Section II, we start by explaining the AMS ticket analysis process. In Section III, we describe several examples of Incident Ticket Analysis from the AMS practice. We will explain how the analyses can help imrpoving the productivity and quality of delivery as well as enable effective solutioning for the support teams. Section IV describes the system architecture and implementation of the AMS Analytics System in its entirety including the overview of the other analysis areas as well as the incident ticket analytics. Section V summarizes previous work in the areas of IT incident analysis and management. It also explains how the presented work is different and how it addresses gap left by the previous work. Finally, in Section VI, conclusions are drawn and future work is outlined.

II. APPLICATION MANAGEMENT SERVICES

The general process of AMS Analytics system starts with collecting and cleasing incident and related data. A number of reports are generated to give different performance metrics. Then statistics analyses are executed to understand the strength and weakness of performance and also provide actionable recommendations to guide service delivery.

A. General Incident Ticket Analysis Process

Figure 1 shows the general process of conducting IT incident ticket data analysis. Specifically, the first step is to collect data, which could include ticket data, resource data, and account operations data. Usually, the account team will help collect data following certain templates. The data will then be reviewed by analysts to ensure that fields are properly mapped and understood, and that there are sufficient amount of data for conducting a meaningful analysis. If necessary, this data collection step can be repeated.

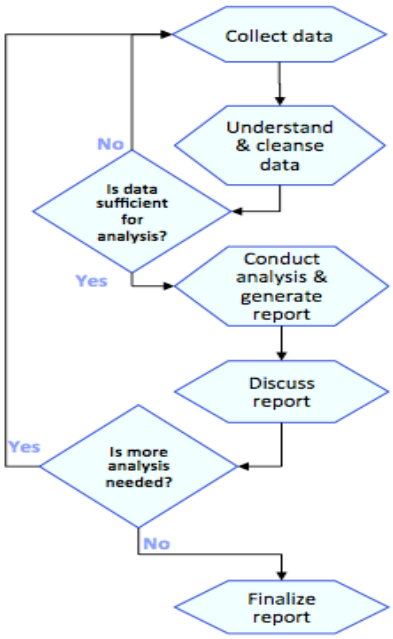

Figure 1: General incident ticket analysis process

The analyst will then conduct the data analysis using appropriate tools and generate a set of reports, which contain information and messages important for the accounts. These reports are then reviewed with the account team in terms of findings, recommendations and potential action items. This is also the time that the account team asks questions, offers insights and provides feedback. Based on the discussion, a need to generate further analysis reports could be raised, which may or may not require extra input data. Otherwise, the reports will be finalized and delivered to the account team.

B. IT Application Incident Analysis

Below, we briefly discuss the set of operational analytics that are developed to mine incident ticket data and offer insights on the quality of application management process and the efficiency and effectiveness of actions taken in the corrective maintenance. Specifically, the analytics are grouped into the following two categories.

1. *Basic operational analytics*: it calculates KPIs (key performance indicators) and their trends directly from the ticket data. The analytics performed includes ticket volume and incident resolution time (i.e., the time spent solving a ticket) trends over time by any ticket group, volume share and work load of any ticket group or individual resource, SLA (service level agreement) performance, ticket arrival distribution by time of the day or day of the week, and ticket resolution performance in terms of the trends in arrivals, completions and backlogs.

2. *Advanced operational analytics*: it dives deeper into ticket data and applies techniques such as clustering, modeling and simulation to further derive business intelligence and identify potential saving opportunities. The analytics performed includes ticket resolution effort estimation, resource load analysis, ticket volume forecasting, resource utilization measurement, resource sharing and pooling, staffing analysis, group right-sizing, and cross-skilling recommendation.

In the traditional practice, reporting KPIs is typically done by stand-alone legacy systems which ofter do not communicate with each other. By integrating operations together, the AMS Analytics System makes tracking and monitoring all KPIs easier & transparent. The system also enables automated and routine examination and warning of anomalous behaviors of incident tickets such as an unanticipated spike or increasing trend in volume or resolution time.

The system is enhanced with the advanced analytics to help manage the AMS operations. For example, ticket volume forecasting based on statistical models provides an important piece of information to the equation of future resource needs and performance expectations. Resource right-sizing based on queuing models enables the capacity planner to optimize resource utilization without sacrificing SLA commitments. An additional analysis capability for resource sharing and pooling identifies opportunities to improve business processes for higher productivity.

III. APPLICATION MANAGEMENT ANALYSIS

Since its initial release, we have applied the AMS Analytics System to several dozens of real customer accounts receiving very positive feedback. In this section, we present some of the reports that we generated for some customers, along with the insights we derived and the feedback we obtained. Note that all data that are reported here have been modified from their original values, as a way to protect their confidentiality.

A. Ticket Volume Distribution

Figure 2 shows the ticket volume distribution by the opening time of ticket arrival dates, breaking down by different severity levels. As we can see, tickets especially of low severity, have been arriving around the clock. In fact, we are observing more tickets arriving between 6pm and 3am, as opposed to the normal daytime schedule. This could be explained if the account has service centers in different parts of

the world, thus receiving tickets at wee hours and late nights due to the time zones. If that is not the case, then it is likely that there are certain applications constantly running in after hours. Consequently, to achieve a better service performance, the account should arrange people to work with night shifts so as to handle after-hour tickets (especially those critical ones) as soon as possible.

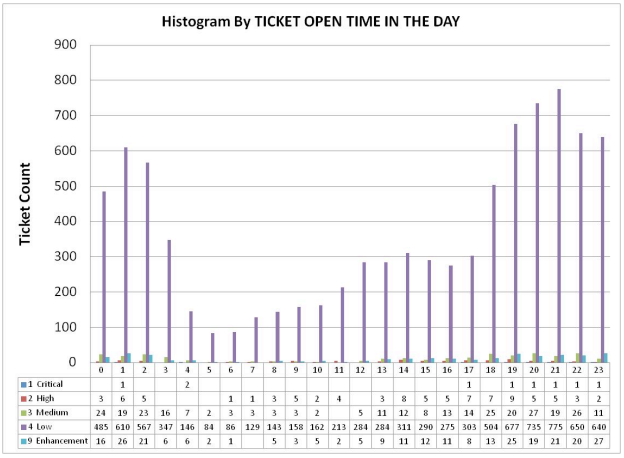

Figure 2: Ticket volume distribution by time-of-day

B. *Ticket Backlog Analysis*

Figure 3 shows the volume distribution of ticket arrival, completion and backlog by month, where the backlog for a particular month is measured as the difference between the number of arriving tickets and resolved tickets in this month, plus the backlog from the previous month. From the figure we see that the arriving ticket volume has been on a general rising trend since January 2012. While the ticket completion curve has been closely following the arriving curve since March 2012, yet due to their gap at beginning of the year (i.e. from January to March 2012), the backlog has also been increasing over the time. Based on this observation, we recommend the account to staff more resource to clear up the backlog, if such backlog has caused severe SLA (Service Level Agreement) breaches. On the other hand, if it is impractical to get extra help, then training the current resources for additional skills and/or extending their work hours properly are good ways to provide possible solutions.

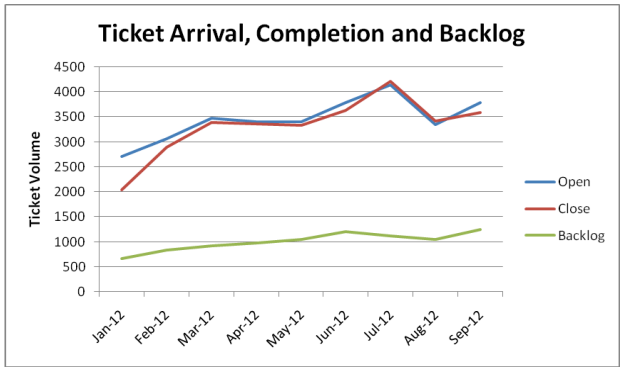

Figure 3: Distribution of ticket arrival, completion and backlog by month

C. *Resolution Time Analysis*

Figure 4 shows the resolution time (in calendar days) over a one-year period. As we can see, the resolution time was on a decreasing trend from 2011-10 to 2012-04, yet it started climbing up all way to 2012-09 with an increase of 28%. This may indicate that the team had been overloaded and could not keep up with the ticket arriving rate. It could also be possible that the problem nature of tickets have changed so that the existing skills of resources were not good enough, resulting in longer resolution time. This chart should be viewed along with ticket volume trend, backlog trend, as well as ticket application/category charts, to identify the right cause.

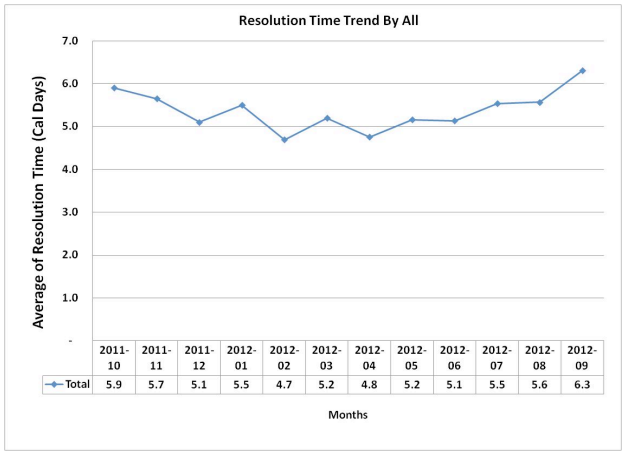

Figure 4: Resolution time trend

D. *Analysis of Resources on Tickets*

For account managers, it is important to keep track of the utilization of their human resources, not only to monitor their productivity, but also to assist workload balance, cross-skill training and resource planning. Figure 5 shows an example of such resource utilization distribution based on their effort spent on ticket handling. From the figure, we see that around 25% of people are fully utilized and 88% of people have utilization above 0.6, which is very encouraging. Nevertheless, there are around 12% of people who spent very small amount of time on ticket (with utilization rate below 0.4).

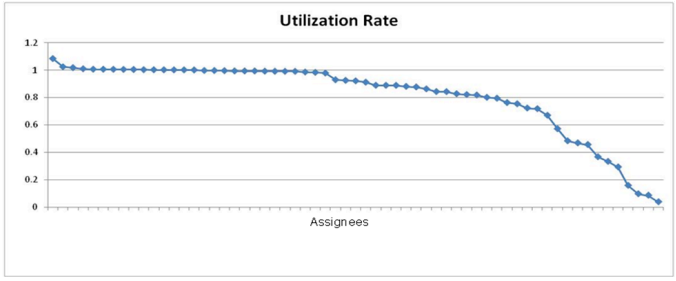

Figure 5: Utilization distribution of assignees

When we reviewed this finding with the account, the team explained that some assignees are actually managers thus they are not fully dedicated to the ticket handling. Moreover, the account is using a different tool for recording tickets of specific applications, which were thus not included into our analysis, consequently, our utilization measurement of those related

assignees would not be accurate enough. Such discussions were very good, as it not only helped us to better understand the account operation, but also motivated the account team to think deeper, dig out potential reasons and find good solutions to improve their operation performance.

E. Anomaly and Trend Detection

The AMS Analytics System makes it easy for automated routine examination of ticket data and detection of anomalous behaviors such as outliers and trends. The anomaly detection is another example of advanced operational analysis capabilities provided by the AMS Analytics System. Figure 6 shows an example where a stochatic process control (SPC) method is used to detect an outlier in a monthly series of ticket volumes. In this example, the mean and standard deviation are calculated from historical volumes and a warning is issued whenever the current volume exceeds a limit which equals 3 standard deviations above the mean (e.g., 3 sigma rule).

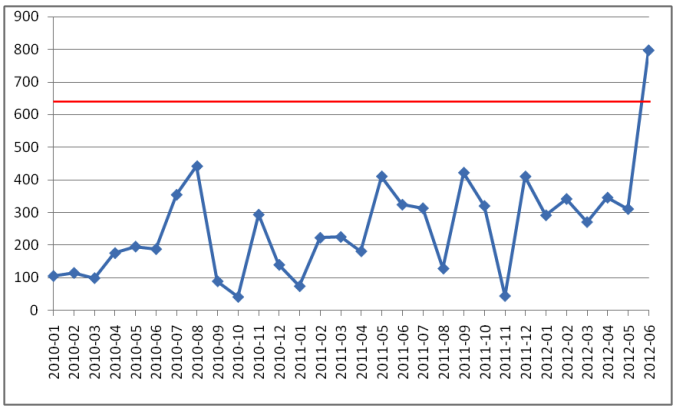

Figure 6: Detection of outliers in monthly ticket volume

Another example is shown in Figure 7, where an increasing trend is detected using a statistical linear regression technique. This method is applied to the volumes in the most recent two-year period. A warning is issued whenever an increasing trend is found to be (a) statistically significant with the consideration of the inherent volatility and (b) practically significant in the sense that the increasing trend is rapid enough to cause concerns about deterioration of resolution time.

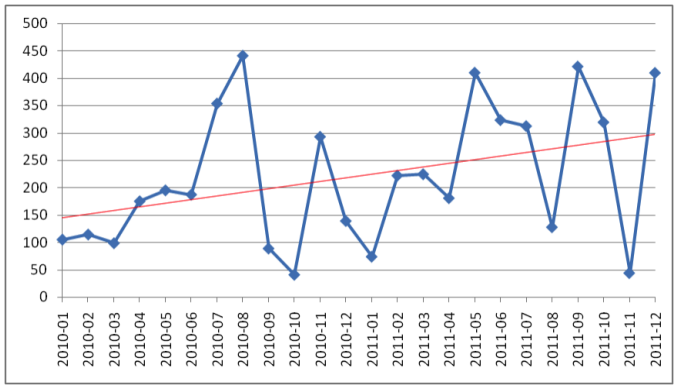

Figure 7: Detection of linear trend in monthly ticket volume

IV. SYSTEM ARCHITECTURE

While we focused on the Incident Ticket Data Analysis in the previous sections, in this section, we will discuss the system architecure in its entirety providing additional analysis capabilities including the Investigative Analytics, the Knowledge Management and the Resource Management. Figure 8 shows the architecture of the AMS Analytics System at a high level, showing different layers including the data, tools, operations and presentation, and different types of users of the system who bring in different analysis use scenarios for the system. The design objective of the system architecture is to provide a standardized, integrated analytics platform supporting AMS delivery that is built by using the standard open stack software on a Web platform, enhanced with advanced analytics, to improve productivity and quality of delivery of the AMS practice.

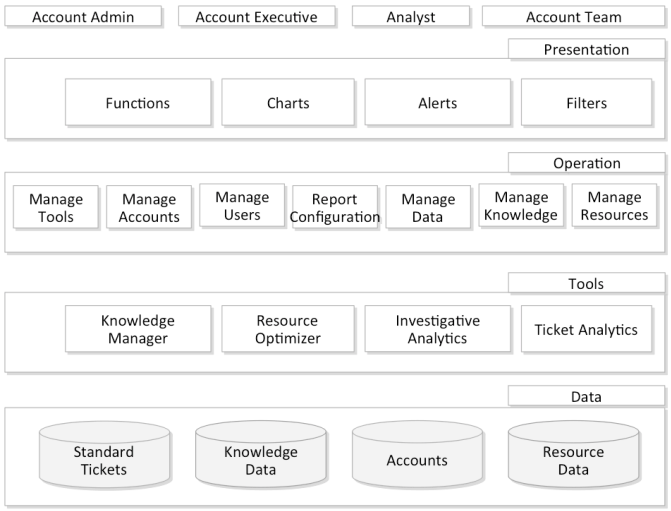

Figure 8: Architecture of the AMS Analytics System

Incident ticket data is usually proprietary to the system that records, and is often stored in different formats from one ticket system to another. Collecting these fragmented data under the standard data model, and enabling tools to operate only on the unified and single data model, simplifies tool development and maintenance. However, this approach puts more burden on the data loading process because it requires mapping the original incident data into the standard data model.

In the Data layer of the architecture in Figure 8, ticket data in the standard format consist of a set of predefined commonly used columns along with several account specific, general purpose columns. A few account specific columns can be designed to store unstructured text data for text analytics. Ticket data should be read-only. Associated with the raw ticket data is derived analytics data created by analytics tools. The Account Data stores information describing AMS customer accounts, the users and profiles, subscription, data upload history, and even metadata to further describe ticket data for complex analysis. The Knowledge Data contains historical ticket resolution information either from the ticket data and authored by SMEs (Subject Matter Experts) in form of dialogs

to find relevant resolutions. This resolution data is used by the Knowledge Manager tool to incrementally and systematically to improve the delivery productivity by reducing ticket resolution time. The Resource Data contains information related to the service delivery resources inclduing skills, roles, and working hours to help plan to improve productivity.

In the Tools layer, the Ticket Analytics tool is for reporting the analysis described in Sections 3 and 4. The current implementation is on the IBM Cognos Business Intelligence Server (reference). It lets experienced analyst users to configure the analytics reports with valid and meaningful options for the report parameters and filters. Less sophisticated users and account executives can take advantage of predefined reports and dashboards by merely selecting predefined parameters and filters. The Investigative Analytics tool can help an account manage and track IT problems for root causes. The process often involves several phases of investigation. The tool provides various diagnosis algorithms to narrow down to a problematic subset of incident data. Also, it ensures a continuous and seamless experience in the incident diagnosis process. The Knowledge Management tool can be used by the incident resolution practitioners. By capturing and authoring resolution knowledge for frequently reoccurring incidents and facilitating the reuse of the knowldege by the practitioners in a time-constrained environment, the tool helps improve the AMS delivery practice. The Resource Management tool helps the delivery executives understand the utilization of the sdelivery resources and plan the resource needs within and across customer accounts.

All the tools on the AMS Analytics System uses a common APIs to access to the Data layer in a consistent way. Security is an important aspect of this approach. The system enables multi-tenancy for the tools by requiring an account ID on the API, which ensures that data for an account is restricted to authorized users for the account. User privilege is determined at login and used throughout the session. Users are tagged to an account, and so their access is bound to the designated account with requisite privileges.

In the Operations layer, the Manage Data component handles the data uploading to the system. Figure 9 shows the data upload process. The raw ticket data is usually acquired in a spreadsheet file from the account's ticket system. As the first step of the data upload, the account administrator maps the raw data to the common data model for the present and subsequent upload. Certain columns such as severity or priority need to be mapped to the standard set for cross account analytics. Data upload to the standard data model requires careful data validation such as conversion of fragmented date formats to the standard format, text length, data types, etc. Once the mapping is defined, data is securely copied over to the Data Warehouse server where the data is cleansed, verified and loaded into the database. Data warehouse tools such as IBM Infosphere Information Data Stage Server import data directly from its original source like database or from various type of files, and help transforming data into the standard data model via ETL process. Customer accounts often require sanitizing data to mask sensitive information, for example, e-mail address, URLs, employee names and phone numbers. Data retention is another important aspect of data management. The AMS Analytics System maintain two consecutive years amount of data to detect and show significant trends. Monthly, with the load of each new dataset, data prior to two years is erased from the system.

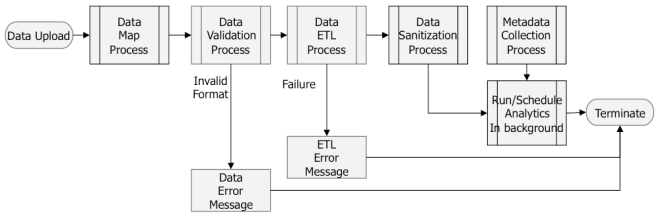

Figure 9: Data upload process

The Manage Tools component in Figure 8 helps get new tools on board to the system and monitor their status and maintenance. Various analysis capabilities that are developed separately can be deployed and integrated onto the AMS Analytics System through the five integration points:

- *Data*: Because the AMS Analytics System has the standard data model and standard APIs to access data, a new tool must be developed to work with the standard data model.
- *Security*: Authentication data should be added to data accessing APIs to control access to account data and for a Single-Sign-On capability across tools on board.
- *View*: User views must be managed based on the supported user roles. Account administrators can determine which reports or functions each user should have access to.
- *Runtime*: Execution of a tool's functionality should be determined and scheduled in a coordinated way through the Application Server runtime.
- *Logs*: The system provides a way to collect messages from each tool such as tool usage information, tool running status, and other system messages.

The AMS Analytics System manages tool subscription of accounts and activate account based on the contract status. The Manage Users component helps an account manage its users and helps users quickly find and access information to tailor their views according to their business needs. The Report Configuration component allows each account to select pre-generated reports that fits their account needs, and customize reports if necessary. The Manage Knowledge component allows the Subject Matter Experts to author incident resolution knowledge so that the incident resolver practitioners can reuse the knowledge to improve the productivity. The Manage Resources component helps an account manage and prepare data about resources handling incidents in association with the use of the Resource Optimizer tool.

Finally, the Presentation layer of the AMS Analytics System takes care of the delivery and formatting of analytics the tools generates for further processing or display. It provides the users with a analytics catalog view with a configirable set of parameters and filters, and a set of notification capabilities.

V. RELATED WORK

IT Service Management is a new area that concerns contribution of IT to the customer's business. ITSM is different from the traditional technology-centered approaches to IT management and business interaction [4, 9, 10, 11]. Incident Management is an IT service management (ITSM) process area [5]. The objective of Incident Management is to restore a normal service operation as quickly as possible and to minimize the impact on business operations, thus ensuring that the best possible levels of service quality and availability are maintained. Information Technology Infrastructure Library (ITIL) is a set of practices for ITSM that focuses on aligning IT services with the needs of business [3]. The main processes of incident management include incident detection and recording, classification and initial support, investigation and diagnosis, resolution and recovery, incident closure, and incident ownership, monitoring, tracking and communication. The AMS Analytics System presented in this paper concerns mostly about the operational analysis, investigation and diagnosis of incidents in ITSM.

In addition, the AMS Analytics System provides advanced analysis capabilities in knowledge management, resource management and operational forecast analysis model. Forecast analysis or predictive analysis is an area of statistical analysis that deals with extracting information from data and using it to predict future trends and behavior patterns [2, 6]. The core of predictive analytics relies on capturing relationships between explanatory variables and the predicted variables from past occurrences, and exploiting it to predict future outcomes. Due to its significant importance in business, incident diagnosis and the application of predictive analytics has been discussed among practitioners through forums and other media [4]. However, the problem is tackled only recently with a rigorous academic method.

While we concentrated on the Incident Ticket Analysis of the AMS Analytics System, we also prepare for publication on the other areas: the Investigative Analytics, the Knowledge Management and the Resource Optimization. The Dialog Manager which is a Knowledge Management tool in the System was presented in [1]. The Dialog Manager is a conversational Web-based tool that helps organizations manage procedural knowledge which is of particular importance for service delivery organizations. It helps standardization of operating procedures in order to deliver a consistently high level of service. In turn, it provides a systematic approach to capturing, managing and disseminating procedural knowledge. It captures procedural knowledge in the form of dialogs that serve as interactive guides by offering a visual knowledge representation and an integrated content management system. It also provides a sophisticated conversation management mechanism that records and manages the activation status of nodes in a conversation.

VI. CONCLUDING REMARKS

Application Management Services brings together an industrialized, globally integrated approach to help companies strategically manage their application portfolios in support of business goals. Application Management Services, which ensure high levels of service quality and availability, has a significant challenge primarily because enterprises often maintain many applications in shared dynamic IT environments. In this paper, we presented the AMS Analytics System for improving the productivity and quality of delivery for AMS practices. The Incident Ticket Analytics of the system measures workload variability, resource productivity and delivery performance by using algorithms from statistics, queuing theory, data clustering and signal processing. The AMS Analytics System architecture provides a standardized, integrated analytics platform supporting AMS delivery that is built on a Web platform using a set of standard open stack software, enhanced with advanced analytics. Since its initial release, we have applied the AMS analytics system to several dozens of real customer accounts, receiving very positive feedback. Our future work is to harden the AMS Analytics System for a large-scale deployment to provide Application Management Services to hundreds of Fortune 500 enterprises covering for the entire life-cycle of the Application Management Serviecs including incident detection, recording, diagnosis, investigation, resolution, resource management, and incident closure.

ACKNOWLEDGEMENT

This research paper would not have been possible without the support of many people. The authors wish to express gratitude to Krishna Ratakonda, Kaan Katircioglu, Rangachari Anand, Rakesh Mohan, Srikanth Tamilselvam, Bikram Sengupta, Vinay Reddy, Shantanu Godbole, Vibha Sinha, David M Pozniak, Michael Sandin, Lisa Smith and others at IBM for their help and support.

REFERENCES

[1] Anand, R.; Lee, J.; Miyamoto, K.; Mei L.; Li Q. (2013). The Dialog Manager. Submitted for publication.

[2] Bishop, C. (2006). Pattern Recognition and Machine Learning, Springer ISBN 0-387-31073-8.

[3] Clifford, David; van Bon, J. (2008). Implementing ISO/IEC 20000 Certification: The Roadmap. ITSM Library. Van Haren Publishing. ISBN 908753082X.

[4] IT Service Management Forum – USA. http://www.itsmfusa.org/

[5] ITIL Incident Management - The ITIL Open Guide, http://www.itlibrary.org/index.php?page=Incident_Management.

[6] Mitchell, T. (1997). Machine Learning, McGraw Hill. ISBN 0-07-042807-7.

[7] Ozgur A, Cetin B, Bingol H: "Co-occurrence Network of Reuters News" (15 Dec 2007) http://arxiv.org/abs/0712.2491.

[8] Salton, G., Wong, A., and Yang, C. S. (1975). "A Vector Space Model for Automatic In-dexing," Communications of the ACM, vol. 18, nr. 11, pages 613–620.

[9] Van Bon, J. ed. (2002). IT Service Management: An Introduction. Van Haren Publishing. ISBN 9080671347.

[10] Van Bon, J. (Editor) (2002). The guide to IT service management. Addison Wesley. ISBN 0201737922.

[11] Van Bon, Jan; Verheijen, Tieneke (2006), Frameworks for IT Management, Van Haren Publishing, ISBN 9789077212905.