Proactive Performance Optimization of IT Services Supply-Chain Utilizing a Business Service Innovation Value Roadmap

Short Paper for a work-in-progress

Ethan Hadar
Corporate Technical Strategy,
Distinguished Engineer
Senior Vice President
CA Technologies, Inc.
Herzelia, Israel
ethan.hadar@ca.com

Jason Davis
Sr Principal Services Architect
CA Technologies, Inc.
Berkshire, GB
jason.davis@ca.com

Donald F. Ferguson
Chief Technology Officer,
Distinguished Engineer
Executive Vice President
CA Technologies, Inc
New York, NY
donald.ferguson@ca.com

Abstract— **Business investment in IT is increasingly linked to IT delivering new or enhanced services that leverage the capabilities of the cloud. The challenge that IT faces especially in cloud environment, is to continuously assess and proactively optimize the performance and quality of supporting IT services being delivered. This evaluation must be performed in the context of the overall Service Level Agreements (SLA) of the composite application and subsequently, the underlying dynamic compound IT services. Our paper presents a proactive optimizer solution that provides on-going service improvement driven by the regular evaluation alternatives of the performance if individual services. Through a Business Service Innovation value roadmap, the Enterprise Architect can model and assemble candidate services for composite applications and automatically use tools to deploy and assure the performance of the composite application. Using our solution, over time the architect can manage the composite application by comparing the quality delivered against simulated alternatives and make recommendations for change. Consequently, using the solution presented in this paper, IT changes are aligned to the challenge to leverage the constantly improving quality of supply-chain IT service while maintaining or reducing costs.**

Keywords: Business Service Innovation; cloud optimization; cloud quality; IT services supply-chain

I. INTRODUCTION

In cloud computing, competition among service providers is affecting the flexibility and dynamics of possible combinations of underlying services within the IT Service supply-chain. This flexibility enables proactive, predictive optimization for both cost reduction and revenue increase [4] of the underlying supporting services.

Business transactions flow through the composite application underlying services and supporting hardware and software resources. Leveraging cloud computing, these services can conceptually be replaced with smart self-service and automation tools [5].

The composite application owner's challenge is to constantly innovate and improve their business service quality while reducing cost. The owner can achieve this goal by proactively adapting and optimizing the composite application's underlying supply-chain services, and composite IT systems. The optimization recommendations are considered based on qualified and quantified metrics. Through frequent recalculation, this dynamic adaptation can drive down Operational and Capital Expenditure (OPEX and CAPEX), raise Service Level Agreements (SLA) quality, and adhere to increased security and privacy compliance needs [1][7].

This paper suggests a system that assists the Enterprise Architect in replacing existing IT supply-chain cloud services according to business needs. The suggested refactoring changes (replacements) to existing IT services and process are based on configured goals for improving internal SLA. The system utilizes services from internal and external, private and public clouds (SaaS, PaaS and IaaS).

Section 2 of this paper presents the conceptual lifecycle framework, the Business Service Innovation upon which our system is structured. Section 3 presents the value for consumers and users of our solution. Section 4 details the necessary conceptual elements, followed by Section 5 that describes the technical modules that implement this solution. Section 6 highlights the value of our solution with associated needs for extensions and future work.

II. STATE OF THE ART

In previous work, we proposed a data store that contains normalized metrics of the services quality based on a Complex Event Processing (CEP) engine [3].

The CEP system is extended in this paper into the "proactive performance optimizer" solution that compares and proposes alternatives to the underlying services of the IT supply-chain. The proposed system's main principles follow a predetermined Business Service Innovation value roadmap (Figure 1), which structures our systems' management steps.

 ISBN: 978-1-61208-216-5

The Business Service Innovation steps are *Model-Assemble–Automate-Assure* and overall *Manage* the evolution, interwoven with IT *Security*.

The detailed steps are:

1. *Model* and *Assemble* a list of applied service elements (composite IT systems and composite applications), and present alternative options to the supporting supply-chain services (candidate services).
2. *Automate* the deployment and *Assure* the quality all underlying services (the ones that participate in the composite application, as well as the candidate services that are not currently part of the composite applications but could replace an existing service.)
3. *Manage*, regularly evaluate the quality and indicate/highlight if an alternative candidate service is superior to the one currently in use, as compared to predefined filter and search criteria.
4. Return to the first step, and accept or reject (manually or automatically) the recommended changes by remodeling the change.

Less common proactive optimization for reducing costs will be to replace a single element with high performance attributes that have a high cost, due to the nature of the high quality SLA, with a lower quality one. This might be the case when this instance (service) is coupled with other transactions, which have a much lower aggregated SLA due to other services and components. The overall SLA is the lowest common denominator, therefore paying a high price for a quality service that does not get used could be considered wasteful.

This example illustrates just one dimension of optimization, in which the cost alternative of the underlying services that participate in a composite transaction can be replaced at any given point in time.

Naturally, there are many other dimensions for proactive optimization such as increasing availability, improving load, increasing speed for change, better robustness and more. Compliance and liability as well as and insurance coverage are additional examples for improvement and replacement [6].

III. THE SYSTEM VALUE

Unlike existing process and composite application design systems that structure a service or a process from a stable state environment point of view, the proposed system constantly improves and evolves a flexible business process. The system operates in a cloud-computing environment that is categorized by a changing environment and service composition possibilities, due to an open market and ease of change. In addition, the proactive optimization system is applicable for non-cloud services as well, and offers the same solution in a system that does not change as often as cloud alternative.

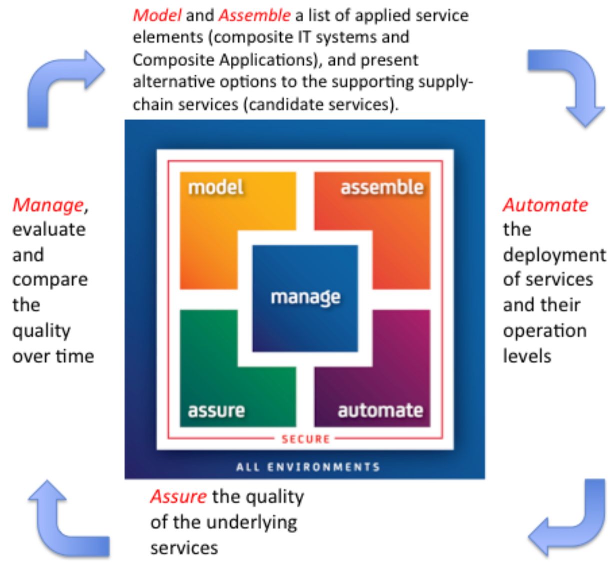

Figure 1: Business Service Innovation value roadmap of the proactive optimizer.

Thus, the proactive optimization of cloud service performance aspects provides several unique value propositions.

- Ongoing suggestions for changes to existing well-designed solutions that could be improved through modifications to underlying services (proactive optimization).
- Notifying the designer of plausible alternatives for existing consumed services.
- Improvement of a monitored service which may be applicable for a certain customer, but considered inappropriate for a different customer due to consumer-specific SLA and quality goals.
- The solution couples the abilities to monitor and compare thresholds of public cloud services (or any service for that matter) with predetermined service levels of consumers, as well as scanning alternative similar services for optional replacements.
- Apply an agile approach for regular incremental and iterative improvement of processes composite applications in production.
- Rationalize the change of the entire portfolio of composite application based on an overall aggregated quality rather than on the underlying single service.

IV. CONCEPTUAL STRUCTURE

This paper presents a refactoring service that continuously monitors possible underlying IT services within the context of a supply-chain of IT services, supporting a composite application. Based on complex event processing (CEP) and predefined threshold metrics, refactoring service triggers assessment of the suggested changes to the optimized services. The system's main modules are SLA thresholds and triggers, and the detection of needed/recommended change.

 ISBN: 978-1-61208-216-5

A. SLA Thresholds and Triggers

The SLA's that are associated with the supplied services that are part of the overall composite application determine the level of aggregated service or SLA that can be achieved. Improvement of the SLA for an overall service is driven by the SLA's or OLA's (operation level agreement) that relate to the underlying service. The system can either monitor a single most impactful supplied service as a candidate for improvement, or, monitor the overall SLA/OLA [6] with a mathematical weighting of the individual contribution of each supply-chain IT service. With a focus on improving the overall service, rather than a single service, the system provides a mechanism that allows the user to set goals for SLA improvement on the overall composite application, and sub-divide it in to the internal services derived OLA/SLA/goals. In the case of self-adaptation, these rules will trigger a suggestion for a change to a service, in the form of a Change Recommendations or external Service Design process.

B. Detection of change in Quality of Service using Complex Event Processing

The Complex Event Processing engine continuously scans internal and external clouds for detection of quality changes to composite elements such as SaaS, IaaS or PaaS that make up an overall composite application. These monitoring tools for service assurance scan for alternatives for improved metrics values.

Once data is collected, the CEP [3] system correlates the information gathered using specific formulas to determine the overall improvement or degradation of quality of service that is being delivered.

C. What-if modeler

For each of the proposed alternatives, an aggregated overall potential SLA is presented. Several of these alternatives can be presented and maintained in the modeler component that captures the structure of the composite application, presenting, over time, the trends and possible quality levels.

V. IMPLEMENTATION STRUCTURE

This section presents the implementation modules and a prototypical usage scenario to the solution according to the Business Service Innovation value roadmap (Figure 2).

A. The implemenation modules

The participating modules [2] are:

- CA AppLogic - constructs and test composite IT systems that support the composite application. The module defines the architecture structure and IT system physical dependency, load balancing, and network configurations.
- CA Service Operations Insight (CA SOI) – monitors internal cloud services and implements the CEP system. This module also defines the behavioral dependency of the supplied services and measured SLA/OLA.
- CA Application Performance Management Cloud Monitor (CA APM) - monitors external cloud services.
- CA Business Service Insight– presents alternative suggestions for change based on measured reported metrics of the vendors as well as aggregated statistics based on surveying users.
- CA Automation Suite for Clouds – automates the changes in infrastructure provisioning and capacity.
- CA IT Process Automation Manager (CA ITPAM) – supports automation of changes in more complex structures, triggering federated identity provisioning, or incident management.
- CA Performance Optimizer - providing preconfigured optimization capacity changes to private datacenter (for infrastructure services).

B. A prototypical usage scenario

In order to realize the connectivity between the modules, consider prototypical activation by an Enterprise Architect.

In this scenario, the enterprise architect selects services, denoted as "atomic services" for use in the composite applications. The services are selected from a list in CA APM Cloud Monitor and CA Business Service Insight (external services) and from a list of infrastructure components (CA AppLogic). The infrastructure services are combined using the CA AppLogic modeler. The integrated external IT services are mapped on a behavioral model presented, and later on monitored, by the CA SOI modeling tool. The structure of CA SOI Composite Application model and the CA AppLogic Composite IT System model defines what underlying services are candidates for changes.

These replaceable services are frequently compared with other options that provide the same conceptual service, yet, currently provide worse SLA, or cost more. CA APM Cloud Monitor and CA Business Service Insight provide the list of compared services.

A comparative "what-if" structure of a potential alternative to the composite application is calculated using a simulator instance of CA SOI behavioral model. Note that this is not the production system version of CA SOI, rather a testing system. In this case CA SOI a combination of production services as well as alternative underlying services. The Enterprise Architect defines the candidate services that can be considered for replacement, and limits the search space for design alternative. A fully opened optimization is not practical, due to contractual limitations, as discussed in the next section.

On a regular basis, the overall SLA for the composite application is calculated for the production composite application, and is constantly compared to the candidate

 ISBN: 978-1-61208-216-5

composite application in which some of the supply-chain services are allowed to change. For every permutation possible, a relative aggregated SLA is presented to the Enterprise Architect, over time.

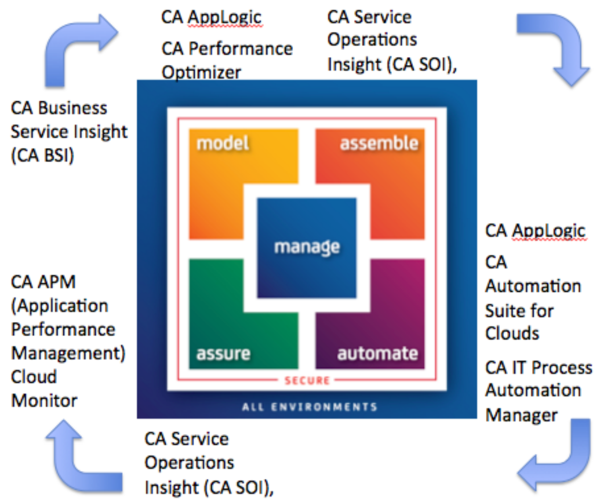

Figure 2 – Proactive optimizer Business Solution Innovation implementation modules.

If over a predefined interval the comparison shows improvement, change automation tools can implement an architect approved change using either CA Cloud Automation suite or ITPAM.

Accordingly, this procedure is repeatable, assuming services are added or removed, as captured on the contractual agreements defined in CA Business Solution Insight.

VI. Discussion and Conclusion

We presented our work-in-progress that provides a proactive optimization solution for enabling ongoing replacement of IT services in cloud environments. Activated according to a Business Service Innovation value roadmap, the solution leverages SLA performance measurements of existing production level applications and their underlying composite IT systems, compared against simulated and monitored alternatives. Consequently, our implemented solution reduced costs and/or improved quality, thereby addressing the challenges of the enterprise architect while providing the business rational for the change. If the change may be applicable to other situations, automation tools can activate the change over dynamic and elastic cloud environments.

However, change typically has associated cost and risk factors that can impact production systems. As a result, change activation should be performed only if the costs savings or ROI is higher and associated risk is lower than the existing state. With pure dynamic resource allocation management over virtual environments, these considerations can be eliminated.

The system does have its limitations; the computational change is not a complex optimization problem as it may be considered. The reason is that not all of the supply-chain IT services can change due to contractual agreements and limitations of liability. However, the complexity is noticeable when overall balancing of all the composite applications in the enterprise are considered, in particular in the domain of mashup and situational applications. Our future research work involves handling optimization for composite situational applications for the entire enterprise.

Even more, optimization is subjective to each customer, based on financial and quality based needs. Different internal consumers may require different service levels. As a result, the ability will be to best match a given SLA required level with supporting provider, changing the proactive optimization to a matchmaking algorithm or better yet, feasibility constraints target function.

Our future work is focusing on providing matchmaking optimization within feasibility box-constraints for assigning the best available supply-chain services

VII. References

[1] Blum D., Schacter P., Maiwald E., Krikken R., Henry T., Boer M., and Chuvakin A., “2012 Planning Guide: Security and Risk Management”, G00224667, Burton IT1 Research, 1 November 2011

[2] CA Web site for tools and products, http://www.ca.com/us/products.aspx, last accesed on May 10 2012

[3] Gal A. and Hadar E., book chapter: “Generic Architecture of Complex Event Processing Systems”, in “Handbook of Research on Advanced Distributed Event-Based Systems, Publish/Subscribe and Message Filtering Technologies”, edited by Annika Hinze and Alejandro Buchmann, IGI Global press, 2009

[4] Ferguson D.F. and Hadar E., “Optimizing the IT business supply chain utilizing cloud computing”, The 8th International Conference on Emerging Technologies for a Smarter World (CEWIT2011), page 1-6, Hyatt Regency Long Island, Hauppauge, New York, November 2-3, 2011.

[5] Hadar E., Connelly K., and Lagunova O., “Agile Evolution of Information Systems Using Model Driven Architecture of IT Services”, Proceeding of the “Architecture in an Agile world” workshop, October 25, OOPSLA 2009, Orlando, Florida, US, 2009

[6] Hadar E. and Danielson D.J., “Certified IT services in a box for cloud computing environments”, CLOSER 2012, 2nd international conference for cloud computing and service sciences, Porto, Portugal, 18-21 April 2012.

[7] Hadar E., Hadar I. and Ferguson D.F., “QDSL - Quality Domain Specific Language for cloud composite applications”, CLOSER 2012, 2nd international conference for cloud computing and service sciences, Porto, Portugal, 18-21 April 2012.