Integrated Lifecycle of IT Services in a Cloud Environment

Karuna P Joshi, Tim Finin, Yelena Yesha,
University of Maryland, Baltimore County, Baltimore, USA

Virtualized Service models are now emerging and redefining the way information technology is delivered to end users. Current research is focused on specific pieces like service discovery, composition etc. There is no holistic view of what would constitute a lifecycle of virtualized services delivered on a cloud environment. In this paper, we propose an integrated methodology that covers the entire service lifecycle. We have divided the IT service lifecycle on the cloud into five phases of requirements, discovery, negotiation, composition, and consumption. We describe each phase and it's sub-phases in detail along with the information that needs to flow between them. We also propose some key metrics that should be tracked for each phase.

Categories and Subject Descriptors: H.3.5 [**Online Information Services**]: Web-based services, XML/XSL/RDF

General Terms: Management, Design, Security, Standardization.

Additional Keywords and Phrases: Services; Methodology; lifecycle

1. INTRODUCTION AND RELATED WORK

The development and maintenance of Information Technology (IT) which was previously either in-house or outsourced is now being replaced by a new delivery model where businesses purchase IT components like software, hardware or network bandwidth as services from providers who can be based anywhere in the world. The service is acquired "on demand". In such scenarios, multiple providers often collaborate to create a single service for an organization. Delivery of IT service is moving away from a single provider mode, and is increasingly based on the composition of multiple services and assets (technological, human, or process) that may be supplied by one or more service providers distributed across the network. Often, a single provider in the cloud seemingly provides the service, and each component service might be virtualized and participate in many composite service orchestrations. The service, in effect, is *virtualized on the cloud*. This virtualized model of service delivery [XU et al 2004] allows easier customization, better utilization and greater responsiveness.

A key barrier preventing organizations from successfully managing services on the cloud is the lack of an integrated methodology for service creation and deployment that would provide a holistic view of the service lifecycle on a cloud. In this paper we present a methodology to address the lifecycle issue for virtualized services delivered from the cloud. We use semantically rich descriptions of the requirements, constraints, and capabilities that are needed by each phase of the lifecycle. These can be reasoned over to automate the phases guided by high level policy constraints provided by consumers, service customers, or service providers. To judge if the lifecycle is progressing successfully, we have proposed metrics for each phase of the lifecycle. This methodology will enable practitioners to create and deploy virtualized services, and measure their success using the associated metrics.

There has been some work done on related issues. [Papazoglou and Van Den Heuvel 2006] have proposed a methodology for developing and deploying web services using service oriented architectures. Their approach, however, is limited to the creation and deployment of web services and does not account for virtualized environment where services are composed on demand. Providers may need to combine their services with other resources or providers' services to meet consumer needs. Other methodologies, like that proposed by [Bianchini et al., 2006], do not provide this flexibility and are limited to cases where a single service provider provides one service. [Zeng et al., 2003] address the quality based selection of composite services via a global planning approach but do not cover the human factors in quality metrics used for selecting the components. [Maximilien and Singh, 2004] propose an ontology to capture quality of a web service so that quality attributes can be used while selecting a service. While their ontology can serve as a key building block in our system, it is limited by the fact that it considers single web services, rather than service compositions. [Black et al., 2007] have proposed an integrated model for IT service management. Their model is limited to managing the service from the service provider's perspective.

Authors' address: Computer Science and Electrical Engineering Department, University of Maryland, Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250, USA. E-mail: {kjoshi1, finin, yeyesha}@cs.umbc.edu

2. PROPOSED METHODOLOGY

For our methodology, we divide the IT service lifecycle on the cloud into five phases. In sequential order of execution they are requirements, discovery, negotiation, composition, and consumption. Figure 1 illustrates our proposed service life-cycle in detail and includes the sub-phases.

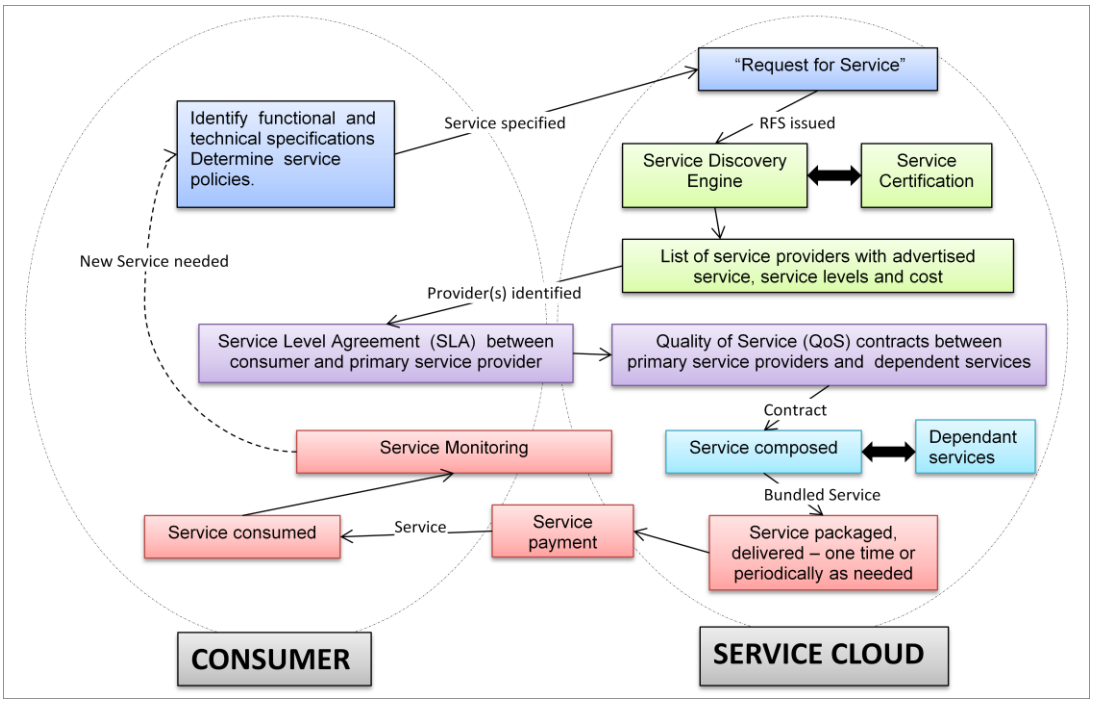

Figure 1: The IT service lifecycle on a virtualized cloud comprises of five phases: requirements, discovery, negotiation, composition and consumption. The figure also illustrates the sub phases and the flow of information between them.

A. Service Requirements

In the service requirements phase the consumer details the technical and functional specifications that a service needs to fulfill and also specifies non-functional attributes such as characteristics of the providing agent, constraints and preferences on data quality, service compliance and required security policies for the service. Depending on the service cost and availability, a consumer may be amenable to compromise on the service quality.

In the Specification sub-phase, the consumer identifies the domain, functional and technical specifications for the service. Functional specification describe in detail what functions/tasks should a service help automate. The technical specifications lay down the hardware, operating system, application standards and language support policies that a service should adhere to. Specifications also list acceptable security levels, data quality and performance levels of the service agent and the service software. Service compliance details like required certifications, standards to be adhered to etc. are also identified

Once the consumers have identified and classified their service needs, they will issue a "Request for Service" (RFS). This request could be made by directly contacting the service providers. The consumer can send the RFS to a few service providers that s/he is comfortable with and get quotes from them for the services. Alternatively, consumers can utilize a service search engine on the cloud to procure the desired service.

B. Service Discovery

In this phase, service providers that offer services matching the specifications detailed in the RFS are searched (or discovered) in the cloud. The discovery is constrained by functional and technical attributes defined, and also by the budgetary, security, data quality and agent policies of the consumer. If the consumer elects the option to search the cloud then the discovery of services is done by using a services search/discovery engine. This engine runs a query against the services registered with a central registry or governing body and matches the domain, data type, functional and technical specifications and returns the result with the service providers matching the maximum number of requirements listed at the top.

C. Service Negotiation

Service negotiation phase covers the discussion and agreement that the service provider and consumer have regarding the service delivered and its acceptance criteria. The service delivered is determined by the specifications laid down in the RFS. Service acceptance is usually guided by the Service Level Agreements (SLA) that the service provider and consumer agree upon. SLAs define the service data, delivery mode, agent details, compliance policy, quality and cost of the service. A consumer may compromise on data quality if it ensures cost saving or may be agreeable to have provider advertisements displayed on the screen if the service is delivered at reduced/no cost. SLAs will help in determining all such constraints and preferences and will be part of the service contract between the service provider and consumer.

At times, the service provider will need to combine a set of services or compose a service from various components delivered by distinct service providers in order to meet the consumer's requirements. The negotiation phase also includes the discussions that the main service provider has with the other component providers. The primary provider will also have to negotiate the Quality of Service (QoS) with the secondary providers to ensure that SLA metrics are met. The final product of the negotiation phase is the service contract between the consumer and primary provider and between the primary provider and the component (or secondary) providers.

D. Service Composition, Orchestration

In this phase one or more services provided by one or more providers are combined and delivered as a single service. Service orchestration determines the sequence of the service components. Many times what is advertised as a single service by a provider could in turn be a virtualized composed service consisting of various components delivered by different providers. The consumer neither knows that the service is composite, nor needs to care. The provider will have to monitor all the other services that it is dependent on (like database services, network services etc.) to ensure that the SLAs defined in the previous phase are adhered to.

E. Service Consumption and Monitoring

The service is delivered to the consumer based on the delivery mode (synchronous/asynchronous, real-time, batch mode etc.) agreed upon in the negotiation phase. After the service is delivered to the consumer, payment is made for the same. The consumer then begins consuming the service. In this phase, consumer will require tools that enable quality monitoring and service termination if needed. This will involve alerts to humans or automatic termination based on policies defined using the quality related ontologies that need to be developed. The **service monitor** sub-phase measures the service quality and compares it with the quality levels defined in the SLA. This phase spans both the consumer and cloud areas as performance monitoring is a joint responsibility. If the consumer is not satisfied with the service quality, s/he should have the option to terminate the service and stop service payment.

3. SERVICE METRICS

To ensure successful discovery, composition and consumption of the services, it is essential to track performance of each phase of the lifecycle. While determining the phases of our methodology, we have also identified the performance metrics associated with each phase. Table 1 lists some of the key metrics that should be tracked to ensure high service quality. There will also be industry or domain specific metrics that the consumers/providers might need to track. A framework [JOSHI et al. 2009] that enables service administrators to automatically track service quality while the service is being consumed will be very beneficial.

Table 1: Some essential metrics that will be tracked for the service lifecycle

Quality Metrics	**Phase**	**Definition**
Data quality	Requirements, Discovery	The quality of the data delivered by the service.
Cost	Requirements, Discovery, Consumption	Cost of the service to the consumer - can be measured as a fixed price cost or on a time and materials basis. While discovering a service, there may be a cost constraints imposed by the consumer.
Security	Requirements, Discovery	Required security/permission levels for the service
Service Gap	Discovery	The gaps that exist between the consumer's requirements and functionality of the services available off the shelf.

Certificate	Discovery	Certification of the service provider to be able to meet service requirements and constraints. Issued by an independent body.
SLA	Negotiation, Consumption	Service Level Agreement between consumer and primary provider. Includes security policy and data quality policy.
QoS	Negotiation, Consumption	Quality of service agreement between primary provider and component providers.
Delivery mode	Consumption	Service delivered in real-time, batch mode or as a one-time service.
Payment options	Negotiation, Consumption	Service payment will be up-front or on a periodic basis (monthly, quarterly, annual etc.). Depending on the option selected, the service will be delivered before or after payment.
Coupling	Composition	Coupling determines how dependent the service is on other services or resources for its delivery [JOSHI et al. 2009]
Cohesion	Composition	Cohesion measures the extent to which related aspects of a requirement are kept together in the same service, and unrelated aspects are kept out.
Reliability	Consumption	Reliability tracks the service quality to ensure the service functionality and data accuracy is maintained
Performance	Consumption	Track the service performance that includes throughput, latency and Response Time.
Consumer Satisfaction	Consumption	Periodically the provider tracks (via survey, opinion poll etc.) if consumers are satisfied with the service.

4. ONGOING WORK

In this paper we have defined the integrated lifecycle for IT services on the cloud. To the best of our knowledge, this is the first such effort, and it is critical as it provides a "big" picture of what steps are involved in deploying IT services. This methodology can be referenced by organizations to determine what key deliverables they can expect at any stage of the process. We also hope that it will enable the academia and the industry to be in the "same page" when they speak about IT services on the cloud.

In our ongoing work, we are developing the ontology to capture the steps and metrics we have identified in the lifecycle using semantic web languages.

REFERENCES

BIANCHINI, D., DE ANTONELLIS, V., PERNICI, B., AND PLEBANI, P., 2006 Ontology-based methodology for e-service discovery, International Journal of Information Systems, The Semantic Web and Web Services, Volume 31, Issues 4-5, June-July 2006, pp 361-380

BLACK, J. et al, 2007 An integration model for organizing IT serviceManagement, IBM Systems Journal, VOL 46, NO 3, 2007

JOSHI K., JOSHI A., YESHA Y., KOTHARI R., A Framework for Relating Frontstage and Backstage Quality in Virtualized Services, UMBC Tech Report TR-CS-09-01, May 2009

MAXIMILIEN, E.M. AND SINGH, M., 2004, A Framework and Ontology for Dynamic Web Services Se-lection, IEEE Internet Computing, vol. 8, no. 5, pp. 84-93, Sep./Oct. 2004

PAPAZOGLOU, M., AND and VAN DEN HEUVEL,W.,2006, Service-oriented design and development methodology, International Journal of Web Engineering and Technology, Volume 2, Number 4, 2006, pp. 412 – 442

XU, M., HU, Z., LONG, W., AND LIU W , 2004, Service virtualization: Infrastructure and applications - The Grid: Blueprint for a New Computing Infrastructure By Ian Foster, Carl Kesselman, Morgan Kaufman, 2004

ZENG, L., BENATALLAH, B., DUMAS, M., KALAGNANAM, J., SHENG, Q., 2003 Quality driven web ser-vices composition, Proceedings of the 12th international conference on World Wide Web, 2003, pp 411 – 421