Chapter 9
AppScale:
Open-Source Platform-as-a-Service

Chandra Krintz
UC Santa Barbara, USA

Chris Bunch
UC Santa Barbara, USA

Navraj Chohan
UC Santa Barbara, USA

ABSTRACT

In this chapter, the authors overview AppScale, an open source cloud platform. AppScale is a distributed and scalable cloud runtime system that executes cloud applications in public, private, and hybrid cloud settings. AppScale implements a set of APIs in support of common and popular cloud services and functionality. This set includes those defined by the Google App Engine public cloud as well as those for large-scale data analytics and high-performance computing.

The goal with AppScale is to enable research and experimentation into cloud computing and to facilitate a "write once, run anywhere" programming model for the cloud, i.e., to expedite portable application development and deployment across disparate cloud fabrics. In this chapter, the authors describe the current AppScale APIs and the ways in which users can deploy AppScale clouds and applications in public and private settings. In addition, they describe the internals of the system to give insight into how developers can investigate and extend AppScale as part of research and development on next-generation cloud software and services.

INTRODUCTION

AppScale is a scalable, distributed, and fault-tolerant cloud runtime system that we have developed at the University of California, Santa Barbara as part of our research into the next generation of programming systems (Chohan, et al., 2009; Bunch, et al., 2010a; Bunch, et.al, 2010b). In particular, AppScale is a cloud platform, i.e. a platform-as-a-service (PaaS) cloud fabric that executes over cluster resources. The cluster resources underlying AppScale can be managed with or without virtualization, e.g. Xen, KVM, or via popular cloud infrastructures including Amazon EC2 and Eucalyptus.

DOI: 10.4018/978-1-4666-0098-0.ch009

The AppScale platform virtualizes, abstracts, and multiplexes cloud and system services across multiple applications, enabling *Write-Once, Run-Anywhere (WORA)* program development for the cloud. In addition to simplifying application development and deployment using cloud systems and modern distributed computing technologies, AppScale brings popular public cloud fabrics *on-premise*, i.e., to private clusters. To enable this, we emulate key cloud layers from the commercial sector, and do so (1) to engender a user community, (2) to gain access to and to study real applications, and (3) to investigate the potential implementations of and extensions to, public cloud systems using open-source technologies.

The initial APIs that AppScale emulates are those of Google App Engine. Google App Engine is a public cloud platform (a complete software stack) that exports scalable and elastic web service technologies via well-defined APIs to network-accessible applications. These APIs implement messaging, key-value data storage, map-reduce, mail, and user authentication, among other services. The platform facilitates easy asynchronous multi-tasking, web server support, elasticity, and resource management. Using Google App Engine, developers debug and test their programs using an open-source Software Development Kit (SDK) provided by Google that implements non-scalable versions of the APIs. Developers then upload their code and data to Google clusters and use Google cluster resources and services on a free (up to some fixed set of per-resource quotas) and *pay-per-use* (resource rental) basis.

AppScale is API-compatible with Google App Engine. As such, applications that execute on Google App Engine can also execute on AppScale without modification, using private cluster resources or public cloud infrastructures. Our API and service implementations are scalable, distributed, fault tolerant, and facilitate high-performance and highly available service access. To enable this, we leverage mature open-source technologies as extensively as is possible. AppScale implements multiple language runtimes (Java, Python, Ruby) as application frontends and a wide range of open source database technologies (key-value and relational) as options for its internal, system-wide datastore. AppScale is not a replacement for Google App Engine or any other public cloud technology. Instead, AppScale is a robust and extensible research infrastructure and private cloud platform that provisions the resources it is allocated scalably across multiple applications.

The AppScale platform also exports services and APIs other than those provided by Google App Engine. These technologies are important for application domains beyond those of web services, including data analytics and data and computationally intensive applications. AppScale exports these technologies as services, i.e. AppScale "service-izes" libraries, tools, and software packages, including MapReduce, X10, R, and MPI. In addition, as it does for the Google App Engine APIs, AppScale provides automatic configuration, deployment, and distribution for these technologies and facilitates their elasticity, load balancing, and fault-tolerance.

Since AppScale provides a software layer between applications and the implementations of cloud APIs and services, we are also able to employ existing cloud services for the implementations. As such, AppScale is a *hybrid* cloud platform—a programming system through which an application can access services from different clouds (a mix of public and private) at different times or concurrently. Developers can use this cloud hybrid support to move data between clouds, e.g., for data analytics, backup of application state and data, to reduce public cloud costs (to use lower-cost alternatives), and to "fall out" from a private cloud with limited resources to a public cloud, on-demand. Developers can also employ AppScale's hybrid support to provide a greater degree of fault-tolerance and availability: in the unlikely but possible scenario in which a public cloud fails or becomes inaccessible, users can still