



Modeling Distributed Ensembles of Virtual Appliances in the Cloud

Student: Xabriel J. Collazo-Mojica, PhD Student, FIU

Research Advisor: S. Masoud Sadjadi, FIU

CI-PIRE Partner Advisors: Dilma Da Silva, IBM, and Fabio Kon, USP

I. Research Overview and Outcome

Problem Statement:

• Realizing distributed ensembles of virtual appliances (DEVAs) in the cloud is a common task for solution providers in today's Internet. Designing these compositions require a good understanding of the underlying details such as the software installation or the network configuration.

• These ensembles are typically deployed in a cloud layer called Infrastructure as a Service (IaaS). Each IaaS provider has different APIs to configure the virtual appliances and their interdependencies, requiring the user to learn the details of a new API if they would like to migrate.

• Research has pointed the need for better tools for composition in the cloud [1]. Clearly, an easier to understand model can help non-experts in cloud computing to develop solutions in this domain.

Motivation:

• We envision that different IaaS APIs and providers will continue to proliferate.

• Non-expert users such as Web developers and CS students should not be bothered by the configuration details of DEVAs, and they should be able to deploy working solutions by themselves and without the help from virtualization and IaaS experts.

• We have identified specific concepts which could use better abstractions than the ones in current IaaS implementations by vendors such as Amazon [2] or GoGrid [3].

```
<?xml ver="1.0" encoding="UTF-8"?>
<DEVA rev="1">
  <NodeList>
    <Node>
      <Appliance rev="1">
        <name>RoR Node</name>
        <guestOS ver="10.04">Ubuntu</guestOS>
        <dependencies>
          <dep ver="2.2.16">Apache HTTPD</dep>
          <dep ver="2.3.8">Ruby on Rails</dep>
          <dep ver="1.2.3">...etc...</dep>
        </dependencies>
      </Appliance>
    </Node>
    <Node>
      <Appliance rev="2">
        <name>MySQL DB</name>
        <guestOS ver="10.04">Ubuntu</guestOS>
        ...etc...
      </Appliance>
    </Node>
  </NodeList>
  <NodeRelationList>
    <NodeRelation>
      <n1Name>RoR Node</n1Name>
      <n1Int type="database-consumer">db</n1Int>
      <relationProperties>
        <prop key="MaxAllowedConn">5</prop>
        <prop key="etc">...</prop>
      </relationProperties>
      <n2Name>MySQL DB</n2Name>
      <n2Int type="database-provider">db</n2Int>
    </NodeRelation>
  </NodeRelationList>
</DEVA>
```

Simplified internal XML model of a DEVA.

Proposed Solution:

• We propose a modeling approach that is abstract enough to allow these virtual appliance ensembles to be seen as "appliances providing and/or consuming services from other appliances. This abstraction allows solutions to be:

- easily designed – we use a visual language to compose appliance-to-appliance service consumption.
- fast to deploy – we leverage existing IaaS APIs and machine image repositories.
- unaffected by IaaS vendor lock-in – since we build solutions from models, we could use different transformation for different IaaS realizations.

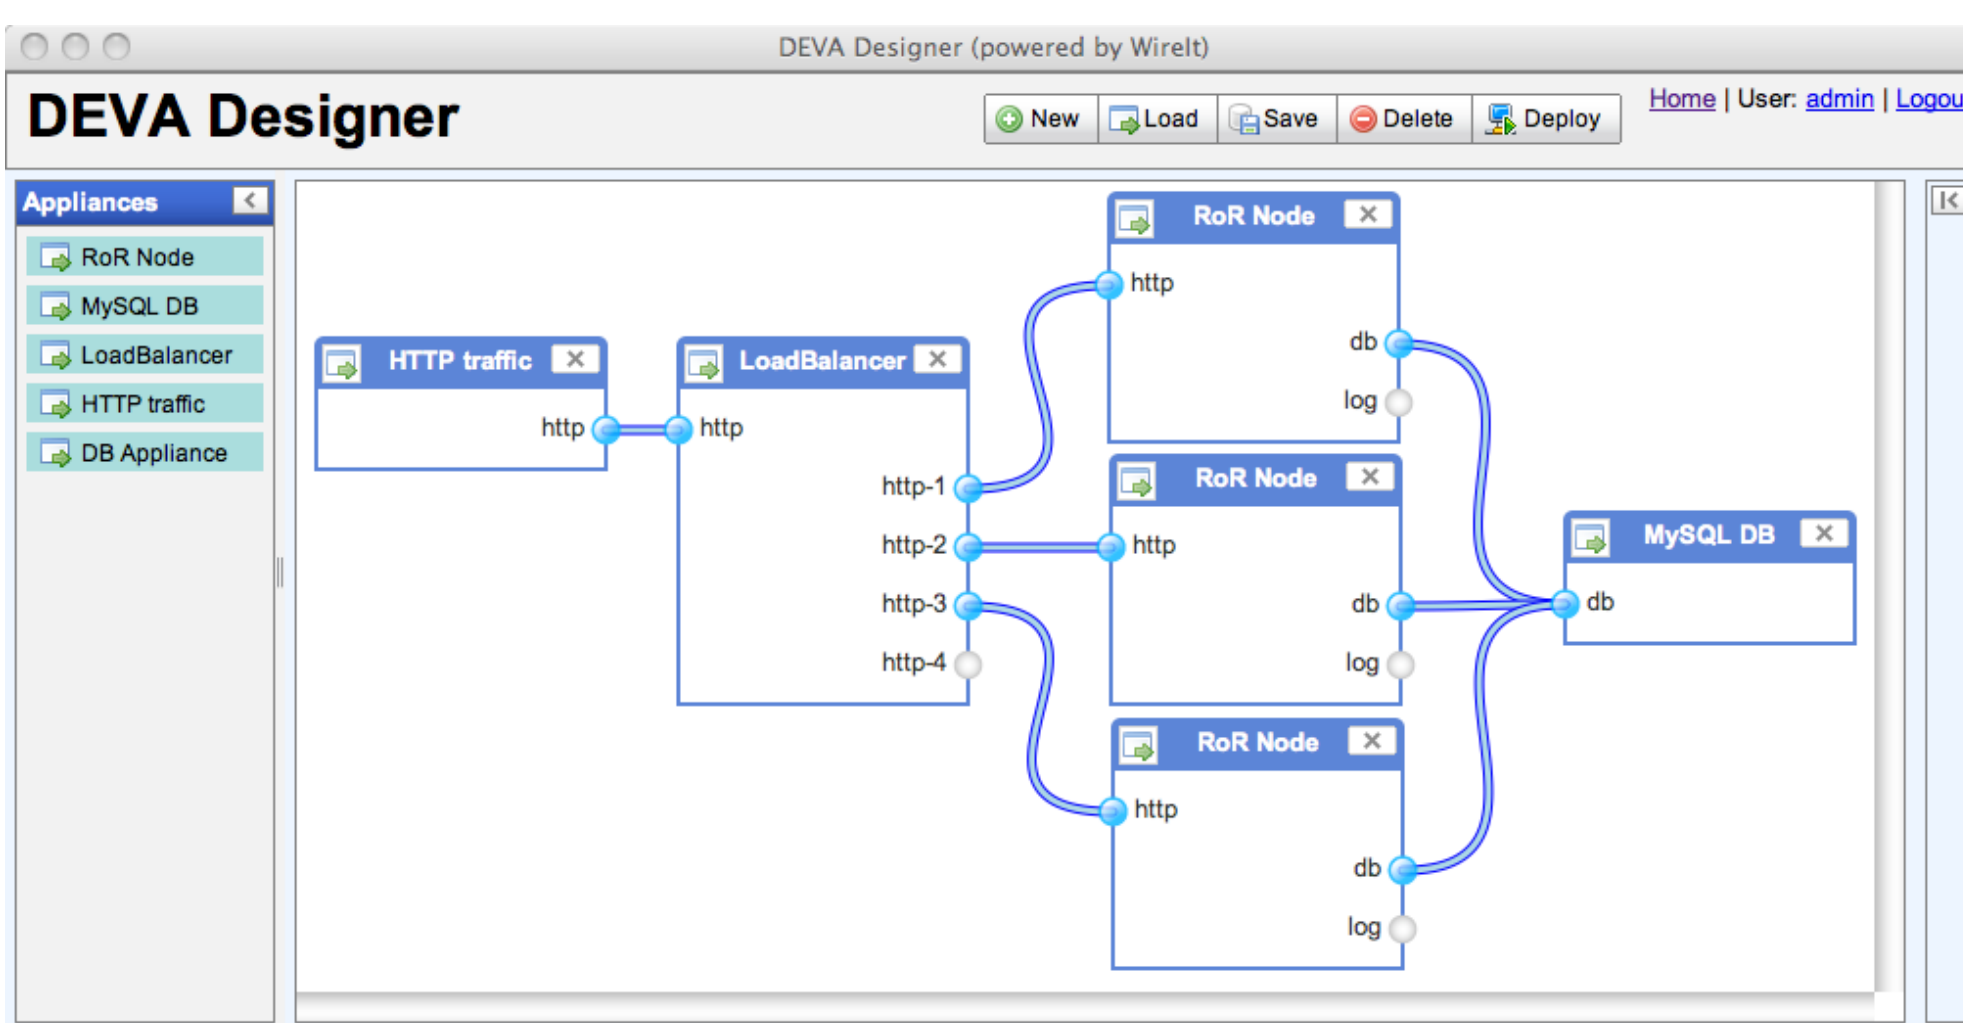
• We have developed a prototype visual designer for easy composition of these environments.

Remarks and Future Work:

• In our position paper [4], we presented the fundamentals of DEVAs modeling.

• We envision that our idea of abstracting away unnecessary configuration details from non-expert users will enable fast deployment of working systems through an automatic configuration process.

• For future work, in short term, we hope to have a prototype that can deploy fully working virtual environments with dynamic change support. In long term, we will put our efforts on providing various IaaS deployment choices.



A complete DEVA model being edited in our designer.



Possible model transformations for deployment on IaaS clouds.

References:

- 1 K. Sripanidkulchai, S. Sahu, Y. Ruan, A. Shaikh, and C. Dorai. Are clouds ready for large distributed applications? SIGOPS Operating Systems Review, 44(2), Apr 2010.
- 2 Amazon Elastic Compute Cloud, Aug 2010. URL <http://aws.amazon.com/ec2/>.
- 3 GoGrid Cloud Hosting, Aug 2010. URL <http://www.gogrid.com/cloud-hosting/cloud-servers.php>.
- 4 Collazo-Mojica et al. Virtual Environments: Easy Modeling of Interdependent Virtual Appliances in the Cloud. SPLASH 2010 Workshop on Flexible Modeling Tools (2010).

II. Industry and International Experience

IBM TJ Watson Research Center – New York, USA

- Obtained valuable industry experience, and attended excellent presentations.
 - This experience will definitely serve for when I'm in the job hunting process.
- Ideas constantly challenged by top-notch researchers.
 - This was the best part, as you learn how to critique your own work.
- Had the opportunity to live in NY and to visit NYC.



Universidade de São Paulo – São Paulo, Brazil

- Obtained a good understanding of the research topics of interest in USP.
 - This puts me in a good position for future collaborations.
- Produced a position paper of great quality jointly with our collaborators.
 - This paper and the feedback I got from it could shape my research in the next few years.
- Had the opportunity to immerse myself in a new culture, to live in São Paulo, and to visit Rio.



III. Acknowledgement

We appreciate the discussions held with David Villegas. The material presented in this poster is based upon the work supported by the National Science Foundation under Grant No. OISE-0730065. This work was also supported by the US Department of Education under P200A090061, and by IBM. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation, the US Department of Education, or IBM.