Leveraging Domain-Specific Software Architectures for Classifying Cloud Service Abstractions

T.S. Mohan[1] Nenad Medvidovic[2] Chris A. Mattmann[2, 3]

[1]*Ecom Research Labs*
Infosys Technologies
Bangalore, India
subramanian_mohan@infosys.com

[2] *Computer Science Department*
University of Southern California
Los Angeles, CA 90089, USA
{neno,mattmann}@usc.edu

[3] *Jet Propulsion Laboratory*
California Institute of Technology
Pasadena, CA 91109, USA
mattmann@jpl.nasa.gov

Cloud computing service abstractions have been broadly classified at three levels—as relevant to IT managers (as Infrastructure as a Service, IAAS), developers (Platform as a Service, PAAS), or end users (Software as a Service, SAAS). However, scientists and business enterprises have found it difficult to choose the right cloud services for leveraging the cloud's disruptive economies of scale: elastic computational power or storage. Reference architectures for the cloud have been particularly lacking, as have specifications of the key interactions between various cloud services and deployment configurations. This is likely due to the organic emergence of the cloud. Ten years ago, the notions behind cloud conflicted with best practices of exotic and expensive hardware and software, as opposed to low-cost, commodity alternatives. But, as Google, Yahoo, and other web innovators have shown, those days are over; clouds can scale economically using commodity hardware and specialized cloud software. The time is ripe for an explicit software engineering focus when considering the cloud.

While clouds are new, several best practices and specific lessons from the world of Grid Computing [1] have a great potential for adoption. The notion of Domain Specific Software Architectures (DSSA) [2] when applied to grids enables the articulation of service component interactions both at configuration time and at runtime [3]. Using a cloud DSSA, we uncover several interesting parallels between grids and clouds. Specifically we look into key examples from Amazon web services (representative of IAAS), Microsoft Azure and Google App Engine (representative of PAAS), and Gmail or Hotmail email service (representative of SAAS) to illustrate the power of applying DSSA. Given the plethora of cloud computing service offerings, we argue that DSSA could represent a plausible model for comparative analysis, helping to illustrate tradeoffs in various cloud service offerings.

References

[1] C. Kesselman, et al. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *Intl' Journal of Supercomputing Applications*, pp. 1-25, 2001.

[2] W. Tracz, et al. Software Development using Domain- Specific Software Architectures. *ACM SEN*. pp. 27-38, 1995.

[3] C. Mattmann, J. Garcia, I. Krka, D. Popescu and N. Medvidovic. The Anatomy and Physiology of the Grid Revisited. In *Proceedings of the Joint Working IEEE/IFIP Conference on Software Architecture & European Conference on Software Architecture*, pp. 285-288, Cambridge, UK, September 14-17, 2009.

Brief Bios

T.S. Mohan is a Principal Researcher in Ecom Research Labs, E&R Unit of Infosys Technologies in Bangalore. His areas of research interests include High Performance Computing – Grid and Cloud Computing as well as large-scale software architectures and systems.

Nenad Medvidovic is an Associate Professor of Computer Science in USC and Director of USC's Center for Systems and Software Engineering. He is a co-author of the textbook *Software Architecture – Foundations, Theory and Practice*, Wiley 2010. He is a recipient of the NSF CAREER and IBM real-time innovation awards.

Chris Mattmann is a senior computer scientist at NASA's Jet Propulsion Laboratory, working on instrument and science data systems on earth science missions and informatics tasks. Mattmann is also an Adjunct Assistant Professor in the Computer Science Department within the USC Viterbi School of Engineering. His research interests are primarily software architecture and large-scale data-intensive systems.