

## EDITORIAL

## Cloud scalability: building the Millennium Falcon

## 1. CLOUD SCALABILITY PROBLEM

Warship construction is a time-consuming, complicated business. The original inception, funding, design, creation of prototypes and training of personnel alone can take years. The actual construction is not typically much faster. The expenses are excessive in both funds and highly specialized labor.

As could be expected, the pressure on starship architects is enormous; once a vessel has been built, the Empire is committing itself to that vessel for the next several decades. At some point, any changes – even trivial ones – in the vessel's design can cost literally billions of credits and thousands of extra man-hours.

This is the point we are at today in cloud computing; the pre-construction and initial phases are completed, much experience has been accumulated, but some inherent cloud features are still causing some trouble. And providers stress their engineers to fulfill their ever-mounting expectations, especially those related with scalability.

This can be understood: the illusion of a virtually infinite computing infrastructure/platform capable of providing an automated on-demand self-service is one of the paramount features of the cloud [1, 2] along with security (after all, no one aims for another massive and expensive Death Star vulnerable to a single X-Wing). Scalability is responsible for making any particular service something more than 'just an outsourced service with a prettier marketing face' [3]. This particular feature pushes cloud constructors to introduce changes to optimize resource consumption while preserving the performance of the deployed application.

Cloud scalability is also an issue that is still poorly understood. Many open questions that call for new research that will eventually incorporate new insight into already running or newly built systems remain. State-of-the-art technologies in cloud scalability typically focus on handling several replicas (service clones) of the image and load balance requests among them ([4] or Amazon's EC2<sup>1</sup>) or federating clouds (infrastructure clones) to increase the pool of available resources [5, 6]. In some sense, these approaches can be compared with Corellian corvettes: they prove the concept in a quick and agile manner, but they are relatively vulnerable during huge business level loads (keeping our analogy with Star Wars, you would not face them with a Star Destroyer). Few are the examples of academic approaches that have reported reaching the scale of Amazon's infrastructure in number of virtual machines. Sharing the lessons learned in that endeavor is still pending.

This special issue covers some of the most relevant trends on scaling cloud infrastructures and platforms. Readers will gain insight on what are the required steps toward optimizing their own clouds to support more concurrent users or operations while minimizing the usage of resources. These articles are also in the interest of those wondering what elements would be good to have as users trying to obtain the maximum scalability for their applications. Section 2 of this document lists some of the strategies that researchers are working on to improve the scalability of clouds, whereas Section 3 briefly describes the contributions presented in this special issue.

## 2. WAY BEFORE THE CLONE WARS

Current state-of-the-art technology in cloud system scaling places us well before 27,000 bBY (before the Battle of Yavin<sup>2</sup>). The Star Forge, a giant automated shipyard created by the Rakata

<sup>1</sup>Available: <http://aws.amazon.com/autoscaling/>. Last visited: March 2012.

<sup>2</sup>Available: [http://starwars.wikia.com/wiki/Battle\\_of\\_Yavin](http://starwars.wikia.com/wiki/Battle_of_Yavin). Last visited: March 2012.

(also known as ‘the builders’), drew energy and matter from a nearby star,<sup>3</sup> which, when combined with the power of the Force,<sup>4</sup> was capable of creating an endless supply of ships, droids,<sup>5</sup> and other war supplies. Current cloud technologies have just started to aim for similar levels of automation.

In the same fashion, as described by Vaquero *et al.* [7], virtual machine replication techniques based on automated rules [8] with the support of load balancers is already mature, as indicated by products such as Amazon’s CloudScale. However, relevant features are still missing like having customized load-balancing strategies in most public cloud vendors. Also, the possibilities to create a personalized virtual network on top of the existing physical network are still very limited. This may be due to the reluctance of network administrators to introduce innovations that may disrupt an essential system, which is required to work 24/7.

Strategies for scaling the infrastructure are typically narrowed down to (i) expanding the size of the underlying hardware or (ii) replicating the available computational/storage substrate, rather than optimizing processes. The same way that creating bigger space ships does not necessarily make them faster or more agile, cloud scaling also needs to rely on smarter approaches that optimize underlying resource consumption. In this special issue, some of these fresher approaches are presented.

In the next section, we present related works that keep a very similar purpose: bringing cloud scalability to the next generation.

### 3. WHAT IS HAPPENING IN THE STAR DRY DOCKS?

Researchers and engineers are working on several approaches to improve the scalability of infrastructure clouds. Many of these lines try to evolve cloud systems so they are smart enough to make a better usage of the available underlying resources upon a variety of conditions, for instance when workload type is a determinant on the performance of the service [9].

Hardware expansion and optimization are among the most active research and engineering themes. Mechanisms to improve the performance of the employed systems in the light of the huge amounts of data to be processed are really common these days [10]. New trends aim to go beyond horizontal scalability or system optimization by including new elements that can be beneficial when it comes to perform compute-intensive tasks.

The bottleneck in cloud systems may not be in the computing infrastructure itself (see [7] and references therein) but in the storage and the network. Intra-data center networking is one of the most active research areas in the cloud arena. New protocols, systems, cabling mechanisms or [11–13] are under study to solve static network assignment, poor server-to-server connectivity (data center switching layer overload), resource fragmentation (caused by popular load-balancing techniques such as destination NATting) and proprietary hardware that scales out. For instance, techniques to optimize data center cabling to maximize bisection bandwidth while minimizing latency for data center applications are a trendy topic [14, 15].

Unlike smaller and more controlled cluster environments, a cloud data center may include a diverse variety of workloads and noisy applications (with regards to their resources usage). Although some seminal works are already part of the state of the art [16], a tighter control on resource allocation and a synergistic integration across the cloud stack [17] are still needed.

Getting to know the vessel is essential to find flaws in its design or construction and possible improvement points. In spacecraft building, improving the vehicle or the on-board monitoring and control systems is an important determinant of the success at war. Thus, modern cloud systems (akin of a Millennium Falcon with top-of-the-line sensor arrays to detect distant Imperial ships before they ever notice) need to predict conflictive situations before they actually occur. Advanced monitoring systems that are capable of delivering the right information to the right decision-making module without creating a huge overload are required [18–20]. Much effort has (and will be) been

<sup>3</sup>Available: [http://starwars.wikia.com/wiki/Abo\\_\(star\)](http://starwars.wikia.com/wiki/Abo_(star)). Last visited: March 2012.

<sup>4</sup>Available: [http://starwars.wikia.com/wiki/The\\_Force](http://starwars.wikia.com/wiki/The_Force). Last visited: March 2012.

<sup>5</sup>Available: <http://starwars.wikia.com/wiki/Droid>. Last visited: March 2012.

devoted to offer appropriate channels for event filtering/aggregation in a way that is meaningful for the application.

Exactly as it happens with mother ships such as Naga Sadow's, there is a need to connect different clouds in an efficient and secure manner. Inter-cloud federation is a topic of raising importance that affects networking, monitoring, scheduling and resource management in general.

And finally, scalability can be reached by developing cloud services/subsystems that provide specific functionalities that ease the construction of scalable applications, in the same way that building a fleet requires ships that are suitable to a variety of different tasks.

Cloud development and adoption were led by a hop in technology and a drive of industrial partners toward a more scalable environment. It seems that scalability begets scalability, and the cloud has now to face a huge amount of data to be measured, collected and analyzed. The huge scale of the underlying infrastructure and platforms, and the vast amount of data generated by users running on top of the cloud grants new and most interesting scaling challenges that will need to be addressed. These are exciting times indeed for 'the builders' of the new Millennium Falcon.

#### 4. BUILDING THE NEXT GENERATION MILLENNIUM FALCON

The articles included in this special issue seem to lead the way where nearly coming research efforts are headed. Here, we try to gather some of the most prominent works related with cloud scalability. Given the variety of research lines addressing this same problem (some of them listed in the previous section), the heterogeneity of the topics covered by the papers in this special issue is not surprising. Table I shows an overview of the challenges addressed by these articles.

Running general-purpose computation on graphic processing units as a mechanism to expand current vessel capabilities is the aim of the work by Expósito *et al.* [21]. This is a revolutionary approach within the classic vertical scalability approach in conventional data centers.

Beyond vertical scalability, there is another trend of techniques that tries to optimize the performance of the on-board system. Distributed Virtual Machine Scheduler describes a resource-scheduling framework, where reconfigurations are enabled by partitioning with a minimum of resources necessary to find a solution to the reconfiguration problem. Quesnel *et al.* [22] propose an algorithm to handle deadlocks that may appear because of the partitioning policy.

Interferences in the stellar positioning and navigation systems terribly affect resource usage and, therefore, vessels' performance. Thus, Barrett *et al.* [23] try to optimize resource usage in the light of interferences between virtual machines on the same hardware, going one step beyond existing approaches that are typically based on setting a threshold, performing badly on unexpected circumstances. These two works and similar systems [8] represent a stage forward toward the creation of droids capable of operating the spacecraft in a nearly optimal manner.

Pattern recognition for similar events is a main concern for automated 'cruise control' in spacecrafts. Storing and handling relevant events in a scalable manner are essential to recognize these patterns. HDKV [24] is one such mechanism that helps with a mechanism to reduce searching time over large data sets that could move spacecraft cruise control one step forward.

A big part of spacecraft engineering has to do with probing the system during its full lifecycle to gather a wealth of information patterns for later analysis. Using the statistical information of these

Table I. Challenges addressed by the contributions in this special issue

Challenge	Paper
Hardware expansion	Usage of GPU power [21]
System optimization/resource allocation	Control deadlock in resource allocation [22]
	Reduce inter virtual machine interferences [23]
Advanced monitoring	Reduce search time in cloud data sets [24]
	Enhance pattern analysis in large data sets [25]

GPU, graphic processing unit.

patterns for fine tuning the system is a huge data-intensive task, which requires especially scalable platforms and algorithms. The work by Rizvandi *et al.* [25] presents a relevant example of such research line.

## 5. CONCLUSION

The Rakata created the Star Forge, but history demonstrated that without proper control, this technological marvel came at a terrible cost. The Star Forge became a fusion of technology and dark side energies that began corrupting the Rakata to gain the immense power it required to operate and ultimately caused the collapse of the Rakatan Empire. In the same way, dealing with the already impressive scale of the data and operations of the cloud requires well-thought mechanisms that are also capable of coping with the exponential increase in data generation expected in the coming years. This special issue paves the way toward understanding new trends and open challenges in the short to mid terms.

### Disclaimer

The opinions herein expressed do not represent the view of HP Labs. The information in this document is provided as is, no guarantee is given that the information is fit for any particular purpose. The above companies shall have no liability for damages of any kind that may result from the use of these materials.

## REFERENCES

1. Buyya R, Yeo CS, Venugopal S, Broberg J, Brandic I. Cloud computing and emerging it platforms: vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Computer Systems* 2009, vol. **25**, no. 6, pp. 599–616. DOI: 10.1016/j.future.2008.12.001
2. Vaquero LM, Rodero-Merino L, Cáceres J, Lindner M. A break in the clouds: towards a cloud definition. *SIGCOMM Computer Communications Review* 2009, vol. **39**, no. 1, pp. 50–55. DOI: 10.1145/1496091.1496100
3. Owens D. Securing elasticity in the cloud. *Communications of the ACM* 2010, vol. **53**, no. 6, pp. 46–51. DOI: 10.1145/1743546.1743565
4. Rodero-Merino L, Vaquero LM, Gil V, Galán F, Fontán J, Montero R, Llorente I. From infrastructure delivery to service management in clouds. *Future Generation Computer Systems* 2010, vol. **26**, pp. 1226–1240. DOI: 10.1016/j.future.2010.02.013
5. Buyya R, Ranjan R, Calheiros R. Intercloud: utility-oriented federation of cloud computing environments for scaling of application services. In *Proceedings of ICA3PP 2010: The 10th International Conference on Algorithms and Architectures for Parallel Processing*, 2010, pp. 19–24. DOI: 10.1007/978-3-642-13119-6\_2
6. Vecchiola C, Chu X, Buyya R, Aneka: a software platform for NET-based cloud computing. *Advances in Parallel Computing* 2009, vol. **18**, pp. 267–295. DOI: 10.3233/978-1-60750-073-5-267
7. Vaquero LM, Rodero-Merino L, Buyya R. Dynamically scaling applications in the cloud. *ACM Computer Communication Review* 2011, vol. **41**, no. 1, pp. 45–52. DOI: 10.1145/1925861.1925869
8. Moran D, Vaquero LM, Galan F. Elastically ruling the cloud: specifying application's behavior in federated clouds. In *Proceedings of the 4th IEEE International Conference on Cloud Computing (CLOUD 2011)* 2011, pp. 89–96. DOI: 10.1109/CLOUD.2011.53
9. Marshall P, Keahey K, Freeman T. Elastic site: using clouds to elastically extend site resources. In *Proceedings of IEEE International Symposium on Cluster Computing and the Grid* 2010, pp. 43–52. DOI: 10.1109/CCGRID.2010.80
10. Borthakur D, Gray J, Sarma JS, Muthukkaruppan K, Spiegelberg N, Kuang H, Ranganathan K, Molkov D, Menon A, Rash S, Schmidt R, Aiyer A. Apache Hadoop goes realtime at Facebook. In *Proceedings of the 2011 international conference on Management of data (SIGMOD '11)*, 2011, pp. 1071–1080. DOI: 10.1145/1989323.1989438
11. Greenberg A, Hamilton J, Maltz DA, Patel P. The cost of a cloud: research problems in data center networks. *SIGCOMM Computer Communications Review* 2008, vol. **39**, no. 1, pp. 68–73. DOI: 10.1145/1496091.1496103
12. Greenberg A, Hamilton J, Jain N, Kandula S, Kim C, Lahiri P, Maltz DA, Patel P, Sengupta S. VL2: a scalable and flexible data center network. In *Proceedings of the ACM SIGCOMM 2009 conference on Data communication (SIGCOMM '09)*, 2009, pp. 51–62. DOI: 10.1145/1594977.1592576
13. Benson T, Akella A, Maltz DA. Network traffic characteristics of data centers in the wild. In *Proceedings of the 10th annual conference on Internet measurement (IMC '10)* 2010, pp. 267–280. DOI: 10.1145/1879141.1879175
14. Mudigonda J, Yalagandula P, Mogul JC. Taming the flying cable monster: a topology design and optimization framework for data-center networks. In *Proceedings of the 2011 USENIX conference on USENIX annual technical conference (USENIXATC'11)*. 2011.

15. Curtis AR, Carpenter T, Elsheikh M, López-Ortiz A, Keshav S. REWIRE: an optimization-based framework for unstructured data center network design. *INFOCOM* 2012. pp. 1116–1124. DOI: 10.1109/INFOCOM.2012.6195470
16. Sotomayor B, Santiago-Montero R, Martín-Llorente I, Foster I. Virtual infrastructure management in private and hybrid clouds. *IEEE Internet Computing* 2009, vol. 13, no. 5, pp. 14–22. DOI: 10.1109/MIC.2009.119
17. Papazoglou MP, Vaquero LM. Knowledge-intensive cloud services: transforming the cloud delivery stack. *Knowledge Service Engineering Handbook*, Taylor & Francis Group: 6000 Broken Sound Parkway NW, Suite 300 Boca Raton, Florida, USA; 2012, chapter 19, pp. 449–494.
18. Clayman S, Galis A, Chapman C, Toffetti G, Rodero-Merino L, Vaquero LM, Nagin K, Rochwerger B. Monitoring service clouds in the future Internet. *Towards the Future Internet – Emerging Trends from European Research*, IOSPress: Nieuwe Hemweg 6B 1013 BG Amsterdam, Netherlands; 2012, pp. 115–126. DOI: 10.3233/978-1-60750-539-6-115
19. Ciuffoletti A. Monitoring a virtual network infrastructure: an IaaS perspective. *SIGCOMM Computer Communications Review* 2010, vol 40, no. 5, pp. 47–52. DOI: 10.1145/1880153.1880161
20. Shao J, Wei H, Wang Q, Mei H. A runtime model based monitoring approach for cloud. In *Proceedings of the 2010 IEEE 3rd International Conference on Cloud Computing (CLOUD '10)* 2010. *IEEE Computer Society*, pp. 313–320. DOI: 10.1109/CLOUD.2010.31
21. Expósito E, Taboada G, Ramos S, Touriño J, Doallo R. General-purpose computation on GPUs for high performance cloud computing. *Concurrency and Computation: Practice and Experience* 2013; 25(12): 1628–1642.
22. Quesnel F, Lèbre A, Südholt M. Cooperative and reactive scheduling in large-scale virtualized platforms with DVMS. *Concurrency and Computation: Practice and Experience* 2013; 25(12): 1643–1655.
23. Barrett E, Howley E, Duggan J. Applying reinforcement learning towards automating resource allocation and application scalability in the cloud. *Concurrency and Computation: Practice and Experience* 2013; 25(12): 1656–1674.
24. Zhou W, Han J, Zhang Z, Xu Z, Dai J. HDKV: Supporting efficient high-dimensional similarity search in key-value stores. *Concurrency and Computation: Practice and Experience* 2013; 25(12): 1675–1698.
25. Babaii N, Taheri J, Zomaya A. A study on using uncertain time series matching algorithms in map-reduce applications. *Concurrency and Computation: Practice and Experience* 2013; 25(12): 1699–1718.

LUIS MIGUEL VAQUERO

*Hewlett-Packard Labs, Cloud and Security Lab, Bristol, Avon, UK*

E-mail: alzindiq@gmail.com

LUIS RODERO-MERINO

*Gradiant, Networks and Applications, Vigo, Pontevedra, Spain*

RAJKUMAR BUYYA

*University of Melbourne, Department of Computing and Information Systems, Melbourne,*

*Victoria, Australia*