



ICT-287510 RELEASE A High-Level Paradigm for Reliable Large-Scale Server Software A Specific Targeted Research Project (STReP)

D7.4 (WP 7): RELEASE project Workshop report

Due date of deliverable:	07.01.2014
Actual submission date:	02.01.2014

Start date of project: 1st October 2011 Lead Contractor: University of Glasgow Duration: 36 Months Revision: 0.1

Purpose: This deliverable describes the workshop we had on the 6th of December, 2013.

Results: Work was summarized during the presentations and tutorials on the workshop which helps to exploit the result of the project.

Conclusion: As a result of the project language extension and tools were developed, which are suitable for the purpose of RELEASE project.

<Mark appropriate dissemination level with an asterisk>

	Project funded under the European Community Framework 7 Programme (2011-14)						
Dissemination level							
PU	Public						
PP	Restricted to other programme participants	(including the Commission Services)					
RE	Restricted to a group specified by the consortium	(including the Commission Services)					
CO	Confidential only for members of the consortium	(including the Commission Services)					

RELEASE project Workshop report

Author: eva.bihari@erlang-solutions.com

1 Executive Summary

RELEASE workshop was organized on the 6th of December in the ESL office after Code Mesh conference with the goal of presenting actual results of the project and discussing their uptake. The timing and location made possible to involve people who attended on Code Mesh to join RELAESE workshop as well.

Code Mesh, the Alternative Programming Conference, focuses on promoting useful nonmainstream technologies to the software industry. The underlying theme is "the right tool for the job", as opposed to automatically choosing the tool at hand. Several talks were provided around Erlang, which was chosen for the implementation of RELAESE activities as well. To schedule the RELEASE workshop after Code Mesh made possible to extend the participants knowledge with ongoing resource activities.

Talks on Code Mesh which were around Erlang, and connected tools (links to the presentations will be updated when these will be available – at the moment of writing this report only a very limited number of them are available):

Tutorials:

Francesco Cesarini: Erlang - A Language for programming Reliable Systems

Omer Kilic: Tinkering with the Raspberry Pi using Erlang

Talks:

Keynote: Jonas Bonér and Francesco Cesarini: Going Reactive: New and Old Ideas for your 21st Century Architectures

Peer Stritzinger: Grisp: Building Spaceships with Erlang (slides)

Bruce Tate: Snake Bitten - Danger and Misfortune in the Evolution of Programming Languages

Simon Thompson: Refactoring Functional Programs: Past and Future (slides)

Joe Armstrong: Fault Tolerance 101

Kresten Krab Thorup: Erjang - With the JVM under the Hood (video, slides)

Torben Hoffmann: Game of Threads - You Spawn or You Die, (video, slides)

Andrew Crump: Deploying the Languages of the Future on Cloud Foundry (slides)

Andreas Olofsson: Are You Ready for 1000-Way Parallelism on a Single Chip? (slides)

The workshop gathered 18 participants; they followed all the presentations and the tutorial as well.



The agenda of the workshop was the following:

Time	Title	Presenter	WP
9:00-9:30	Overview of RELEASE	Phil Trinder	
9:30-10:00	VM & Scalability - part 1	Kostis Sagonas	WP2
10:00-10:30	Coffee break		
10:30-12:00	VM & Scalability - part 2	Kostis Sagonas	WP2
12:00-13:00	Lunch break		
13:00-14:15	Scalable Distributed Erlang	Natalia Chechina	WP3(6)
14:15-14:30	Break		
	Deploying Erlang apps in the Cloud using	Enrique	WP4(6)
14:30-15:45	Wombat	Fernandez	
15:45-16:00	Break		
16:00-17:00	Profiling	Huiqing Li	WP5(6)

RELEASE workshop kicked off by Phil with an overview of the project scope, objective and strategies, followed by presentations of the work packages. (See: 4.2.1)



The order of the presentations followed the structure of the work packages, Kostis Sagonas, work package leader of WP2, was given the first detailed presentation (See:4.2.2). After giving introduction to Erlang and its concurrency model he explained how the Erlang VM (Virtual machine) is working. After describing the architecture of the 'bencherl' tool (1), he also showed live examples how scalability measurements can be done.

The next topic was to explain the improvement was done with the scalability of ETS (Erlang Term Storage) between OTP R11B and R16B and also gave a walkthrough of ETS benchmarks.

As the next layer within the project is to deal with the Erlang language extensions and introduce SD (Scalable Distributed) Erlang, Natalia gave a presentation covering the above-mentioned topics. (See 4.2.3) She explained what are the requirements and limitations of distributed Erlang, and gave an introduction to Orbit – a tool used as an example to prove the benefits of using SD Erlang. Then she provided an overview how SD Erlang is working, what is the definition of S_groups and what are the functions were introduced to support S_groups. Then she compared Distributed Erlang Orbit and SD Erlang Orbit even providing code examples.

Following the theoretical presentation there were a tutorial, where participants were able to use SD Erlang language extension on top of Erlang/OTP R15B03 (3) and create their own S_groups on several nodes, check their registered names, check s_groups related information and run SD-Orbit (2). This tutorial was useful to "make the participants hand dirty", i.e. better understand how the language extension is supposed to work.



The next session was provided by Enrique Fernandez, who explained how Erlang applications to be deployed on the cloud with Wombat (an orchestration tool). (See: 4.2.4) He drew us through the basics of heterogeneous and semi-explicit deployment and described the different steps and inputs that need to be provided by the user. A live demo closed the presentation when a Riak node was successfully deployed to Amazon W3 cloud.

The last topic of the workshop was to understand what multicore codes does, how it makes use of the resources and how bottlenecks in the system can be found. Huiqing Li explained how the above questions could be investigated by using tracing and profiling tools. (See 4.2.5) First she summarized what are the already existing tools and gave a live demo about **P**ercept the offline **Er**lang **c**oncurr**e**ncy-**p**rofiling **t**ool (4).

Then she explained what are the enhancements are introduced with Percept2, and how to use the tool (a live demo was showed as well). A case study demonstrated how bottleneck was found within the system using Percept2, finally a summary about ongoing development work was provided.

After the presentations the participants had the possibility to raise questions, for later reference the presentations are available in the project github repository:

https://github.com/release-project/midproject-workshop

2 Conclusion

A questionnaire was sent out to the participants asking feedback 6. Generally participants were satisfied with the event, so it fulfilled it's purpose. For detailed feedback is available see 4.3

An other session is planned to be organized in June, 2014 at Stockholm in connection to the Erlang User Conference, which is the biggest yearly event for Erlang developers. This will help the project to exploit the results towards a wider group of potential users.

Change L

Version	Date	Comments
0.1	09/12/2013	First version submitted to internal reviewers
0.2	20/12/2013	Updated

3 References

- 1. bencherl git@github.com:softlab-ntua/bencherl.git
- 2. SD Orbit: <u>https://github.com/release-project/benchmarks</u>
- 3. SD Erlang install: https://github.com/release-project/otp/tree/sd-erlang-tutorial
- 4. Percept2: https://github.com/RefactoringTools/percept2
- 5. Frank Lubeck and Max Neunhoffer. Enumerating Large Orbits and Di- rect Condensation. Experimental Mathematics, pages 197–205, 2001.
- 6. RELEASE workshop feedback questionnaire: https://www.surveymonkey.com/s/ReleaseWorkshop

4 Appendix

4.1 Orbit

An Orbit is a symbolic computing kernel and is a generalization of a transitive closure computation [5]. To compute Orbit for a given space [0::X] a list of generators g1; g2; :::; gn are applied on an initial vertex x0 2 [0::X]. This will create new numbers (x1:::xn) 2 [0::X]. Repeatedly, generator functions are applied on each of the new numbers and this continues until no new number is generated.

For detailed information see: <u>https://github.com/release-</u> project/benchmarks/blob/master/AboutOrbitBenchmark.pdf

4.2 Abstracts

4.2.1 RELEASE overview

This presentation introduced the RELEASE project, outlining it's scope, aims, objectives, work package structure, with a brief sketch of achievements to date.

RELEASE is a European Union 7th Framework Programme research project (an EU FP7 STREP). The project runs between October 2011 and September 2014, the total funding is Euro 3.6M, and it employs approximately 25 Researchers. The project addresses the software aspect of the multicore challenge. It aims to scale the radical distributed actor-programming paradigm to build reliable general-purpose software, such as server-based systems, on massively parallel machines (100 000 cores).

The objectives and associated Work Packages (WPs) of the project are

- 1. To Scale Erlang by co-designing
 - an extension, Scalable Distributed (SD) Erlang (WP 3)
 - the BEAM virtual machine (WP 2)

for reliable scalability.

- 2. To develop a scalable virtualisation infra-structure (WP4) that creates, manages and dynamically scales super-clusters of heterogeneous clusters/cloud instances
- 3. To develop improved software development technologies for large scale Erlang software (WP5)
- 4. To evaluate SD Erlang and associated tools/methodologies using case studies on an IBM Blue Gene (WP6)

The Work Packages are related as follows:



4.2.2 Scalability and the Erlang VM by Kostis Sagonas

In this talk we will take a look at how Erlang's concurrency model is supported by the virtual machine of the Erlang/OTP system and how its virtual machine has evolved during the last few years to support the scalable execution of Erlang programs on modern

multicore machines. In addition, we will present and demonstrate BenchErl, a userextensible, open source scalability benchmark suite for the Erlang/OTP system and its applications. The last part of the talk will present in detail the scalability characteristics of the Erlang Term Storage: its underlying architecture, current status and scalability of its implementation, the concurrency tuning options it comes with, and some aspects of the implementation of ETS that can be still improved.

The talk will include live demos and will give the chance to participants to download BenchErl and play with it.

4.2.3 Scalable Distributed Erlang by Natalia Chechina

In this talk we present and demonstrate Scalable Distributed (SD) Erlang -- an extension of Distributed Erlang functional programming language for reliable scalability. We start by discussing distributed Erlang limitations and introducing scalable groups (s_groups). The s_groups enable scaling of Erlang nodes by eliminating transitive connections, i.e. a node may belong to multiple s_groups, and each s_group node has transitive connections with the nodes from the same s_groups and non-transitive connections with other nodes. Then we refactor a distributed Erlang program into an SD Erlang program using Orbit benchmark -- a generalization of a transitive closure computation. That is we group nodes in s_groups, add gateway processes and nodes that enable message passing between worker nodes from different s_groups without a direct connection between the nodes, and introduce an additional hashing to identify an s_group where a corresponding part of the distributed hash table is stored. We conclude the talk by brief introduction of the ongoing work on the semi-explicit process placement in scaled networks.

4.2.4 Deploying Erlang apps in the Cloud using Wombat by Enrique Fernandez

Wombat is the cloud deployment tool for distributed Erlang applications developed in the fourth work package - a.k.a. the Scalable Virtualisation Infrastructure work package - of the RELEASE project. This session aims at demonstrating that the developed tool is capable of automating the deployment of real-word distributed Erlang applications on virtual environments (i.e. clouds). More concretely, this session will showcase the automatic deployment (and bootstrapping) of a Riak cluster - Riak is a well known NoSQL DDBMS - on the virtual infrastructure offered by one of the most popular public Infrastructure-as-a-Service providers, Amazon EC2.

4.2.5 Profiling by Huiqing Li

The number of cores in modern computer system is increasing rapidly, and Erlang's support for multi-core system is also being continuously improved. To take full advantage of the computing power provided by multi-core Erlang, new Erlang applications need to be written with more parallelism to assure that there are enough processes ready to run, and existing sequential Erlang applications can be refactored to introduce more parallelism and therefore scalability. However, very often one might find that an application does not scale as expected.

In this talk, we present the Erlang concurrency-profiling tool Percept2, which is an enhanced version of the original Erlang concurrency-profiling tool Percept. Built on top of Percept, Percept2 provides not only process profiling information, but also scheduler and function profiling information. Visualization and navigation through data is made easy in

Percept2, and the performance has also been improved by the parallelization of tool itself.

With a number of live demos, We show Percept2 in action and demonstrate Percept2 can be used to find Erlang application level bottlenecks, and also how the tool can be used to provide guidance and feedback on the parallelization of existing sequential Erlang applications.

4.3 Feedback after the workshop

4.3.1 Questions

The following marks were used:

- 1 "Strongly disagree," or the lowest, most negative impression
- 2 "Disagree"
- 3 "Neither agree nor disagree," or an adequate impression

4 - "Agree"

5 - "Strongly agree," or the highest, most positive impression

1. What's your overall satisfaction with the workshop?

1	2	3	4	5
			\bigcirc	\bigcirc

2. Workshop venue

How do you rate the workshop regarding:

	1	2	3	4	5
Comfortable					
Well located		\Box			
Food and refreshments were adequate					
3. Workshop Programme	1	2	3	4	5
Workshop objectives					
Relevance of the subjects		\Box			\Box
Availability of documentation	\Box				
4. Speakers/ Facilitators	4	0	0	4	F
	I	2	3	4	Э

ICT-287510 (RELEASE) 07 January 2014

	1	2	3		4	5
Knowledge of the subject				C		
Clarity of the talk				0		
Availability to answer questions				C		
5. Participants						_
1	2	2	3	4		5
Motivation	C)				
Level of participation)				
6. Workshop outcomes			-			_
		1	2	3	4	5
The objectives were met						
The results are useful for my a	ctivity					
The workshop was a good op to learn about how distributed work and what supports will b provided by the newly develop Erlang tools and language eler	s O					

4.3.2 Answers

6 participants filled in the questionnaire. Timing might have the reason of the low number of responses as we are very close to the end of the year and vacation season.

Q1 What's your overall satisfaction with the workshop?



	1	2	3	4	5	Total	Average Rating
(no label)	0%	0%	16.67%	16.67%	66.67%		
	0	0	1	1	4	6	4.50



Q2 Workshop venue How do you rate the workshop regarding:



	1	2	3	4	5	Total Respondents
Comfortable	0%	0%	16.67%	66.67%	16.67%	
	0	0	1	4	1	6

2/10

RELEASE workshop feedback questionnaire

Well located	0% 0	0% 0	16.67% 1	16.67% 1	66.67% 4	6
Food and refreshments were adequate	0% 0	0% 0	0% 0	33.33% 2	66.67% 4	6



1 2 3 4 5

	1	2	3	4	5	Total Respondents	
Workshop objectives	0%	0%	0%	33.33%	66.67%		
	0	0	0	2	4	6	
Relevance of the subjects	0%	0%	0%	16.67%	83.33%		
	0	0	0	1	5	6	
4/10							

RELEASE workshop feedback questionnaire								
Availability of documentation	0%	0%	16.67%	50%	33.33%			
	0	0	1	3	2	6		

13



1 2 3 4 5

	1	2	3	4	5	Total Respondents
Knowledge of the subject	0%	0%	0%	16.67%	83.33%	
	0	0	0	1	5	6
Clarity of the talk	0%	0%	0%	50%	66.67%	
	0	0	0	3	4	7
			6/10			·

RELEASE workshop	feedback	k questionnaire
------------------	----------	-----------------

Availability to answer questions	0%	0%	0%	0%	100%			
	0	0	0	0	6	6		



1 2 3 4 5

	1	2	3	4	5	Total Respondents
Motivation	0%	16.67%	16.67%	16.67%	50%	
	0	1	1	1	3	6
Level of participation	0%	16.67%	16.67%	33.33%	33.33%	
	0	1	1	2	2	6





1 2 3 4 5

	1	2	3	4	5	Total Respondents
The objectives were met	0%	0%	0%	50% 3	50%	6
	-	-	-		-	-

9/10

DELEASE	workshop	foodbook	questionnoire
RELEASE	WORKSHOD	leeuback	questionnaire

		-				
The results are useful for my activity	0% 0	16.67% 1	0% 0	16.67% 1	66.67% 4	6
The workshop was a good opportunity to learn about how distributed systems work and what supports will be provided by the newly developed Erlang tools and language elements	0% 0	0% 0	16.67% 1	0% 0	83.33% 5	6