

Voxxed Days CERN 2019



Report of Contributions

Contribution ID: 2

Type: **Talk**

Welcome & Keynote

Wednesday 1 May 2019 09:00 (1 hour)

Author: MARTIN, Robert "Uncle Bob"

Presenter: MARTIN, Robert "Uncle Bob"

Contribution ID: 3

Type: **Talk**

A Dozen Cool Things You can Do with JVM Languages

Wednesday 1 May 2019 10:30 (50 minutes)

Languages offer a lot more than syntax and compilers. They often have supporting libraries and special facilities that set them apart from other languages. Some languages offer special compiler support for a particular construct, like tail call optimization, for example. Others provide interesting library support or capabilities. In this presentation we will dive into 12 cool things we can do with different languages on the JVM, things that are either impossible or hard to do in Java, but are quite easy and useful to realize in other popular languages on the JVM. If mixing these languages is an option on your projects, you'll have dozen more reasons after this presentation.

Author: SUBRAMANIAM, Venkat**Presenter:** SUBRAMANIAM, Venkat

Contribution ID: 4

Type: **Talk**

The Importance of Fun in the Workplace

Wednesday 1 May 2019 10:30 (50 minutes)

An in-depth look at what makes software development a roller coaster where the highs of 0 compiler warnings are quickly cancelled out by the pain of long hours, bad requirements, endless configuration, clueless managers and a plethora of other issues which make death by a thousand cuts seem like a good idea... They will answer questions such as: "Why is programming often called an art despite having its underpinnings in formal logic?" "How can I rediscover the delight I felt when I first started coding?" "What's that rush I feel when my test passes? Am I addicted to TDD?" Combining Psychology, Philosophy and Computer Science, Dr Holly Cummins will present a series of practical tips to help you rediscover the euphoria that you felt the very first time a metal box in front of you came to life and cried out "Hello World".

Author: CUMMINS, Holly**Presenter:** CUMMINS, Holly

Contribution ID: 5

Type: **Talk**

The Error of our Ways

Wednesday 1 May 2019 11:30 (50 minutes)

What happens when things break? What happens when software fails? We regard it as a normal and personal inconvenience when apps crash or servers become unavailable, but what are the implications beyond the individual user? Is software reliability simply a business decision or does it have economic, social and cultural consequences? What are the moral and practical implications for software developers? And when we talk of 'systems', are we part of the 'system'? What about the bugs on our side of the keyboard? In this talk we will explore examples of failures in software and its application, and how they affect us at different scales, from user to society.

Author: HENNEY, Kevlin**Presenter:** HENNEY, Kevlin

Contribution ID: 6

Type: **Talk**

Modern Java: Change is the Only Constant

Wednesday 1 May 2019 13:30 (50 minutes)

In the past two years, we changed Java in three ways that we never have before: We modularized the platform, we removed some components, and we accelerated the pace of new releases. These changes aim to keep Java vibrant in an ever-changing world of competing platforms and new styles of application deployment, whether to the cloud or to app stores. We'll see why these changes are not as disruptive as you may think, and demonstrate some of the recent and future additions to the language and the platform.

Author: REINHOLD, Mark**Presenter:** REINHOLD, Mark

Contribution ID: 7

Type: **Talk**

Deconstructing Web frameworks for a more resilient code base

Wednesday 1 May 2019 15:50 (50 minutes)

In the wonderful world of frontend Web, the frameworks and their ecosystems come and go... For a lot of devs, this endless race is getting tiring. Nevertheless, we continue to make architecture decisions that will push us tomorrow to throw away and rewrite lots of code. Step by step, it gets very easy to be locked into a given framework/ecosystem and too easy to limit one's skills to what this ecosystem provides. We tend to forget the basics and most of the time it pushes us to write code where all layers and concerns are mixed up (UI, DOM access, data access, router, stage manager, bundling, build...). In this session, we'll deconstruct the different parts of modern Web frameworks : Why are we using those tools? Which problems do they solve? How can we avoid our code to be locked in and mixed up with those ecosystem? How can we produce future-proof code that resists against the hype and evolutions of the Web platform?

Author: SABLONNIÈRE, Hubert**Presenter:** SABLONNIÈRE, Hubert

Contribution ID: 8

Type: **Talk**

Deep Learning in Computer Vision

Wednesday 1 May 2019 14:30 (50 minutes)

When we look at the world our brain instantly turns the images we see into information, intuition and feeling. It does so with an enormous computational effort, using a network with overwhelming, unexplored architecture. We are not even close to create a machine of comparable capabilities, but step by step, improving both the hardware and algorithmic approach, we can make machine understand images they see. In this talk you will learn how to design and develop such a network. We will go from the basic machine learning concepts, through the advanced practical tips, up to using the most-recent state-of-the art architectures in practice. You will see examples on how insanely complicated problems can be solved using a portable computer. And in the meanwhile you will understand what is going on, why new layers are added and why new concept have to be introduced. While this presentation briefly introduces all the basic concepts, as we move on it quickly touches advanced and expert areas of the field.

Authors: KUBIAK, Błażej; KUDRYNSKI, Krzysztof**Presenters:** KUBIAK, Błażej; KUDRYNSKI, Krzysztof

Contribution ID: 9

Type: **Talk**

Bulletproof Java Enterprise Applications for The Hard Production Life

Wednesday 1 May 2019 11:30 (50 minutes)

It's one thing to develop Java Enterprise applications, to execute, and validate them in test environments. Running the application in production, however, is a whole other story. Production life is harsh, unforgivable and the ultimate verification whether your application is able to add value to its users. Besides knowing the Java Enterprise APIs, developers therefore also need to know how to meet production requirements. This session shows which non-functional requirements enterprise developers need to be aware of to build stable and resilient applications. We will see different resiliency approaches, such as circuit breakers, bulkheads, or backpressure, in action. We will implement these concepts using the Java Enterprise API, MicroProfile, and certain Java EE extensions. Furthermore, we will see, how especially service meshes enable us to add resiliency to multiple microservices in a transparent, redundancy-free fashion. Adding telemetry to our applications is crucial to gather insights about the application under production load. We will see how to enhance our microservices with concerns such as monitoring, tracing, or logging in a minimalistic way.

Author: DASCHNER, Sebastian**Presenter:** DASCHNER, Sebastian

Contribution ID: **10**

Type: **Talk**

A Deep Dive into Progressive Web Apps

Wednesday 1 May 2019 14:30 (50 minutes)

Progressive Web Applications or PWAs in short, utilize modern Web Technologies to create cross-platform applications that offer the best of web and native user experiences.

This talk begins with the introduction of the PWAs, advantages over the legacy approaches and covers the core elements that define a PWA accompanied by examples with popular libraries like Angular, React and Vue.

Author: CIVICI, Cagatay

Presenter: CIVICI, Cagatay

Contribution ID: 11

Type: **Talk**

Transform a Legacy Application with Kubernetes and Istio

Wednesday 1 May 2019 16:50 (50 minutes)

Kubernetes is often used to run micro-services based applications. Istio's service mesh adds even more power and flexibility to those applications. But what about your Good Old Monolithic Legacy Application? Should you feel left alone if you're not cloud-native? I'm going to demonstrate how to leverage Kubernetes and Istio to improve an existing application without touching a single line of its code. Because, you know, the original developers are long gone. Through lots of demos and zero slide, the attendees will learn to add Monitoring, Blue/Green Deployments, Canary builds, Fault tolerance...

Author: GAGEOT, David**Presenter:** GAGEOT, David

Contribution ID: 12

Type: **Talk**

Unleashing your musical creativity with Magenta.js

Wednesday 1 May 2019 13:30 (50 minutes)

People make things. You make things. Most of the time you make Very Serious Things™ that help your bosses sell more shoes or saxophones or those tiny coffee packets. And that's good, because you gotta eat and stuff. But something else has gotta eat, too: your brain. Music is creation, art is creation, code is creation. I think it's important to goof around with code sometimes, or make things that let other people goof around with code. And sometimes, that becomes the Very Serious Thing™. Magenta.js is one of these things; let's talk about it!

Author: DINCULESCU, Monica**Presenter:** DINCULESCU, Monica

Contribution ID: 13

Type: **Talk**

Closing Remarks

Wednesday 1 May 2019 17:45 (15 minutes)

Wrapping up an amazing day

Author: SILVA, João

Presenter: SILVA, João

Contribution ID: 14

Type: **Talk**

Interaction Protocols: It's all about good manners

Wednesday 1 May 2019 15:50 (50 minutes)

Distributed and concurrent systems can be considered a social group that collaborate to achieve collective goals. In order to collaborate a system of rules must be applied that affords good hygiene, fault tolerance, and effective communication to coordinate, share knowledge, and provide feedback in a trusted manner. These rules form a number of protocols which enable the group to act as a system which is greater than the sum of the individual components. In this talk we will explore the history of protocols and their application when building distributed systems. Protocols provide the foundation on which the quality attributes are delivered. Qualities such as performance, resilience, and security.

Author: THOMPSON, Martin**Presenter:** THOMPSON, Martin

Contribution ID: 15

Type: **Talk**

Recursion, Iteration, and JavaScript: A Love Story

Wednesday 1 May 2019 16:50 (50 minutes)

Recursion and iteration usually seem like epic rivals, on opposite sides of a fundamental tradeoff between machine- and human-efficiency. Iteration's stateful repetition can be more performant, but many coders prefer to work with the self-referential abstraction of recursion. But in this talk we'll see that with a little love from JS, these apparently star-crossed paradigms can actually unite to improve each other: custom Iterables allow recursion to power more human-efficient iterative code, and tail call optimization uses iteration to make recursive code more machine-efficient.

Author: VAKIL, Anjana**Presenter:** VAKIL, Anjana

Contribution ID: 16

Type: **Quickie**

The code is back to the database with Oracle MLE (GraalVM): Python and JavaScript in the DB

Wednesday 1 May 2019 12:25 (20 minutes)

Because of the lack of support for modern agile interpreted languages, the software architects have moved the logic execution out of the database. Oracle Labs is building a polyglot engine with GraalVM, and integrates it to the Oracle and MySQL RDBMS. Currently in beta, this scales the database applications by avoiding the CPU context switches between the procedural code and the persistence layer. Short presentation and demo on the beta.

Author: PACHOT, Franck**Presenter:** PACHOT, Franck**Session Classification:** Byte-sized talks from CERN engineers - Main Auditorium

Contribution ID: 17

Type: **Quickie**

The CERN Cloud Infrastructure: 5 years in 15 minutes

Wednesday 1 May 2019 12:45 (20 minutes)

In 2012, at CERN we started the deployment of our private Cloud Infrastructure using OpenStack. Since then we have moved from a few hundred cores to a multi-cell deployment spread between two data centres. After 6 years deploying and managing OpenStack at scale, we now look back and discuss the challenges of building a massive scale infrastructure from 0 to +300K cores. With this talk we will navigate through the history, architecture, tools and technical decisions behind the CERN Cloud Infrastructure.

Author: ABAD, Daniel (CERN)**Presenter:** ABAD, Daniel (CERN)**Session Classification:** Byte-sized talks from CERN engineers - Main Auditorium

Contribution ID: 18

Type: **Quickie**

Puppet + containers = multi-tenancy

Wednesday 1 May 2019 13:05 (20 minutes)

The CERN IT infrastructure consists of more than 40000 Puppet-managed virtual and physical machines located in two data centres. All the Puppet catalogs are served by a shared pool of almost 200 Puppet Servers regardless of the organisational unit from which the requests come. This deployment might lead to security risks as the number of nodes and personnel writing Puppet code increases.

The Configuration Team, as part of the IT department at CERN, is exploring ways to configure a multi-tenancy enabled Puppet Server deployment minimising the performance impact for users and the resource usage.

This talk will go over the different alternatives the team has considered and the results obtained so far.

Author: MORENO GARCIA, David (CERN)

Presenter: MORENO GARCIA, David (CERN)

Session Classification: Byte-sized talks from CERN engineers - Main Auditorium

Contribution ID: 19

Type: **Quickie**

Graph Network in High Energy Physics

Wednesday 1 May 2019 12:25 (20 minutes)

Collisions at the CERN Large Hadron Collider (LHC) produce showers of particles that are detected by heterogeneous detectors composed of hundreds of millions of individual sensors, laid out under complex geometry. An event can be seen as a tree of detectable particles branching from the unstable particles (e.g., the Higgs boson) produced in the collisions. Once detected, events are collected as arrays of isolated hits, which are then collectively processed to reconstruct the trajectory and energy of the particles that created them. In this contribution, we describe how the reconstruction and identification of these particles can be performed using graph networks. Given their capability of learning sparse representations, graph networks are ideal tools to create a fixed-geometry representation of an event, abstracting from the irregular geometry of the detectors used at colliders. As a first processing step of raw data, they provide an interface between particle detection and more rigid deep learning techniques, e.g., convolutional neural networks. In this respect, they represent a step forward to realistic deep learning applications for collider physics. As examples, we consider the task of reconstructing the trajectory of charged particles bending in the magnetic field of the detectors (tracking), denoising from parasitic collision (pile-up mitigation) and the identification of heavy particles (e.g., Higgs bosons) from the spray of particles that they produce (jets), where state-of-the-art performances are achieved.

Author: Dr VLIMANT, Jean-Roch (California Institute of Technology (US))

Presenter: Dr VLIMANT, Jean-Roch (California Institute of Technology (US))

Session Classification: Byte-sized talks from CERN engineers - Council Chamber

Contribution ID: 20

Type: **Quickie**

Docker image testing in GitLab CI

Wednesday 1 May 2019 12:45 (20 minutes)

The IT-DB-DAR section at CERN is moving Java web applications from VMs to containers running in Kubernetes. For this a Docker image is being developed which needs to integrate well with existing CERN services. In this session I want to show you how we test the image using GitLab CI. This includes our usage of Docker-in-Docker, how we isolate ourselves from the central CERN SSO and LDAP services, writing pipelines in a way that allows local execution through GitLab CI's "exec" feature, scaling the pipeline to handle multiple built versions, lightweight desired state testing with Goss, as well as an overview of the features used.

Author: LOEKKEBORG, Thomas Holene (Norwegian University of Science and Technology (NTNU) (NO))

Presenter: LOEKKEBORG, Thomas Holene (Norwegian University of Science and Technology (NTNU) (NO))

Session Classification: Byte-sized talks from CERN engineers - Council Chamber

Contribution ID: 21

Type: **Quickie**

Providing Continuous Integration at CERN

Wednesday 1 May 2019 13:05 (20 minutes)

Continuous practices like Continuous Integration, Deployment, Testing and Delivery play a big role in modern software development life cycles that allow organizations to frequently and reliably release new services and provide updates and CERN users rely on it to a large extent. Because CERN's current CI infrastructure and design have evolved largely over the last years, this talk aims to describe its evolution and past problems and solutions. This talk reviews our current tools and features that CERN provides to facilitate Continuous Integration and Continuous Deployment, as well as extra components that make it as accessible as possible to most of the organisation use cases, such as Jenkins, Nexus, distributed S3 caching and Docker registry.

Author: JUAREZ GONZALEZ, Daniel (Universidad de Oviedo (ES))

Co-author: LOSSENT, Alexandre (CERN)

Presenter: JUAREZ GONZALEZ, Daniel (Universidad de Oviedo (ES))

Session Classification: Byte-sized talks from CERN engineers - Council Chamber