

# IBM/Red Hat/Fedora Leftovers

By *Roy Schestowitz*

Created 30/06/2020 - 8:29pm

Submitted by Roy Schestowitz on Tuesday 30th of June 2020 08:29:40 PM Filed under [Red Hat](#) [1]

- [Ask the experts during Red Hat Summit Virtual Experience: Open House](#) [2]

One of the most popular activities during the Red Hat Summit Virtual Experience was the Ask the Experts sessions, where attendees could engage with Red Hat experts and leadership in real time, so we're bringing it back for our Open House in July.

- [Making open source more inclusive by eradicating problematic language](#) [3]

Open source has always been about differing voices coming together to share ideas, iterate, challenge the status quo, solve problems, and innovate quickly. That ethos is rooted in inclusion and the opportunity for everyone to meaningfully contribute, and open source technology is better because of the diverse perspectives and experiences that are represented in its communities. Red Hat is fortunate to be able to see the impact of this collaboration daily, and this is why our business has also always been rooted in these values.

Like so many others, Red Hatters have been coming together the last few weeks to talk about ongoing systemic injustice and racism. I'm personally thankful to Red Hat's D+I communities for creating awareness and opportunities for Red Hatters to listen in order to learn, and I'm grateful that so many Red Hatters are taking those opportunities to seek understanding.

- [The latest updates to Red Hat Runtimes](#) [4]

Today, we are happy to announce that the latest release of Red Hat Runtimes is now available. This release includes updates that build upon the work the team has done over the past year for

building modern, cloud-native applications.

Red Hat Runtimes, part of the Red Hat Application Services portfolio, is a set of products, tools and components for developing and maintaining cloud-native applications. It offers lightweight runtimes and frameworks for highly-distributed cloud architectures, such as microservices or serverless applications. We continuously make updates and improvements to meet the changing needs of our customers, and to help developers better build business-critical applications. Read on for the latest.

- 

#### [Kourier: A lightweight Knative Serving ingress](#) [5]

Until recently, Knative Serving used Istio as its default networking component for handling external cluster traffic and service-to-service communication. Istio is a great service mesh solution, but it can add unwanted complexity and resource use to your cluster if you don't need it.

That's why we created Kourier: To simplify the ingress side of Knative Serving. Knative recently adopted Kourier, so it is now a part of the Knative family! This article introduces Kourier and gets you started with using it as a simpler, more lightweight way to expose Knative applications to an external network.

Let's begin with a brief overview of Knative and Knative Serving.

- 

#### [CodeTheCurve: A blockchain-based supply chain solution to address PPE shortages](#) [6]

This past April, creative techies from all over the world gathered online for CodeTheCurve, a five-day virtual hackathon organized by the United Nations Educational, Scientific, and Cultural Organization (UNESCO) in partnership with IBM and SAP. Participants all worked toward the goal of creating digital solutions to address the global pandemic.

Our team focused on the goal of improving the efficiency of the personal protective equipment (PPE) supply chain in order to prevent shortages for health care workers. With the rise of the current global pandemic, supplies of medical equipment have become more critical, particularly PPE for medical workers. In many places, PPE shortages have been a serious problem. To address this challenge, we proposed that a blockchain-based supply chain could help make this process faster and more reliable, thereby connecting health ministries, hospitals, producers, and banks, and making it easier to track and report information on supplies.

- 

#### [Analyze your Spark application using explain](#) [7]

It is important that you have some understanding of Spark execution plan when you are optimizing your Spark applications. Spark provides an explain API to look at the Spark execution plan for your Spark SQL query. In this blog, I will show you how to get the Spark query plan using the explain API so you can debug and analyze your Apache Spark application. The explain API is available on the Dataset API. You can use it to know what execution plan Spark will use for your Spark query without actually running it. Spark also provides a Spark UI where you can view the execution plan and other details when the job is running. For Spark jobs that have finished running, you can view the Spark plan that was used if you have the Spark history server set up and enabled on your cluster. This is useful when tuning your Spark jobs for performance optimizations.

- [What's new in Apache Spark 3.0 \[8\]](#)

The Apache Spark community announced the release of Spark 3.0 on June 18 and is the first major release of the 3.x series. The release contains many new features and improvements. It is a result of more than 3,400 fixes and improvements from more than 440 contributors worldwide. IBM Center of Open Source for Data and AI Technology (CODAIT) focuses on a number of selective open source technologies on machine learning, AI workflow, trusted AI, metadata, and big data process platform, etc. has delivered approximate hundreds of commits, including a couple of key features in this release.

- [GSoC Progress Report: Dashboard for Packit \[9\]](#)

Hi, I am Anchit, a 19 y.o. from Chandigarh, India. I love programming, self-hosting, gaming, reading comic books, and watching comic-book based movies/tv.

The first version of Fedora I tried was 21 when I came across it during my distro-hopping spree. I used it for a couple of months and then moved on to other distros. I came back to Fedora in 2017 after a couple of people on Telegram recommended it and have been using it ever since. A big reason why I stuck with Fedora this time is the community. Shout out to @fedora on Telegram. They're nice, wholesome and helpful. They also got me into self-hosting and basic sys-admin stuff.

- [Fedora Looking To Offer Better Upstream Solution For Hiding/Showing GRUB Menu \[10\]](#)

Fedora for the past few releases doesn't show the GRUB boot-loader menu by default when only Fedora is installed on the system as there is little purpose for most users and it just interrupts the boot flow. But for those wanting to access the GRUB bootloader menu on reboot, they offer integration in GNOME to easily reboot into this menu. The other exception is the menu will be shown if the previous boot failed. This functionality has relied on

downstream patches but now they are working towards a better upstream solution.

Hans de Goede of Red Hat who led the original GRUB hidden boot menu functionality is looking to clean up this feature for Fedora 33. The hope is to get the relevant bits upstream into GNOME and systemd for avoiding the downstream patches they have been carrying. This reduces their technical debt and also makes it easier for other distributions to provide similar functionality.

•

#### [Fedora Developers Discussing Possibility Of Dropping Legacy BIOS Support](#)<sup>[11]</sup>

Fedora stakeholders are debating the merits of potentially ending legacy BIOS support for the Linux distribution and to only support UEFI-based installations.

Given Fedora 33 GRUB changes planned and things being easier if they were to just switch to the UEFI-based systemd sd-boot as well as Intel planning to end legacy BIOS support in 2020 and UEFI being very common to x86\_64 systems for many years now, Fedora developers are discussing whether it's a good time yet for their bleeding-edge platform to also begin phasing out legacy BIOS support.

[Red Hat](#)

---

**Source URL:** <http://www.tuxmachines.org/node/139346>

#### **Links:**

- [1] <http://www.tuxmachines.org/taxonomy/term/142>
- [2] <https://www.redhat.com/en/blog/ask-experts-during-red-hat-summit-virtual-experience-open-house>
- [3] <https://www.redhat.com/en/blog/making-open-source-more-inclusive-eradicating-problematic-language>
- [4] <https://www.redhat.com/en/blog/latest-updates-red-hat-runtimes>
- [5] <https://developers.redhat.com/blog/2020/06/30/kourier-a-lightweight-knative-serving-ingress/>
- [6] <https://developer.ibm.com/blogs/codethecurve-hackathon-ppe-supply-blockchain-project/>
- [7] <https://developer.ibm.com/blogs/how-to-understanddebug-your-spark-application-using-explain/>
- [8] <https://developer.ibm.com/blogs/ibm-contributions-to-spark-30/>
- [9] <https://communityblog.fedoraproject.org/gsoc-progress-report-dashboard-for-packit/>
- [10] [https://www.phoronix.com/scan.php?page=news\\_item&px=Fedora-Better-Up-GRUB-Menu](https://www.phoronix.com/scan.php?page=news_item&px=Fedora-Better-Up-GRUB-Menu)
- [11] [https://www.phoronix.com/scan.php?page=news\\_item&px=Fedora-Discuss-Legacy-BIOS-Drop](https://www.phoronix.com/scan.php?page=news_item&px=Fedora-Discuss-Legacy-BIOS-Drop)