Feature profiles for Fedora are created by the community-run Marketing Team; full versions of the excerpts below can be found at https://fedoraproject.org/wiki/F13_feature_profiles.

Tools for Python developers: David Malcom
Interviewed by Mel Chua

There are 3 Python-language related features coming out in F13. Let's start with parallel installation of Python 2 and Python 3. What is this? Why is it cool?

Python 3 fixes a lot of long-standing issues in the language, but the cost of doing that is [having to ask] "does the Python 2 or Python 3 universe give me the modules with the functionality I need?" And Fedora's answer is "well, we'll have both." You can use both Python 2 and Python 3 simultaneously - you don't have to pick one or the other. For the Python 3 one, we've tried to provide RPM packages of Python code known to work with Python 3. We've gone through various modules... and tested them. And we're doing because we need the Python 3 stuff ourselves. Much of Fedora's web infrastructure is written in Python.

Tell us about the "Easier Python Debugging" feature. What does that mean?

One of the great things about Python is how easy it is to wrap external libraries (e.g. written in C). The downside of this is that if one of these libraries has a bug, then that bug takes out the whole of the Python process, without giving you a nice Exception/traceback. What I did was use this to write Python code that knows about the insides of Python itself, so you now have Python code running inside the gdb process, which knows how to scrape data out of another dying process. So rather than being told the hexadecimal address of where the object is stored in RAM, gdb should tell you that e.g. you have a [1, 2, 3].

Now, this feature is something that was originally made for Fedora - this is the first place it's come out?

Yes. I also recently got this code into upstream, into Python's SVN repository, and it's likely to be in Python 2.7 when that comes out, though it works fine with 2.6.

In other words, the Python community liked the work you were doing so much they decided to make it part of the Python language itself.

I have a very keen, not-at-all-vested interest in making Python easier to debug!

More of David's interview, including coverage of how SystemTap allows novice programmers to profile their running code, is available at https://fedoraproject.org/wiki/Python_in_Fedora_13.

Feature profiles for Fedora are created by the community-run Marketing Team; full versions of the excerpts below can be found at https://fedoraproject.org/wiki/F13_feature_profiles.

boot.fedoraproject.org and you : Mike McGrath
Interviewed by Paul Frields

I had the privilege of talking to Mike McGrath, who works for Red Hat and leads the community Infrastructure team, about the new services on boot.fedoraproject.org.

So boot.fedoraproject.org got its start from the boot.kernel.org website. What does boot.kernel.org do exactly and who created it?

BKO (boot.kernel.org) provides a network based install system to users who have downloaded a pre-configured image which comes in ISO, disk, USB, etc formats. It was started as a Google Summer of Code project done by Pravin Shinde and sponsored by John 'Warthog9' Hawley. Users can go to http://boot.fedoraproject.org and download images from the download section.

What attracted you to the idea of making a version of this site for use in the Fedora Project? What are your goals for it?

During the last FUDCon, I attended a talk John put together on BKO. The technology was something that had been in the back of my mind for some time. Well, there it was sitting in front of me, done and functional -- I immediately liked it.

John and I both share a general hatred of CD/DVD media. Network speeds are generally getting faster in many locations making network based booting more feasible. Also, when you download a DVD or ISO image, you're likely downloading lots of bits and packages you don't need or want. Doing a network boot allows users to just download the packages they need. It was decided the best course of action was to put a site together for Fedora to help showcase the technology, as well as letting Fedora easily control settings and config defaults.

Who do you see using the site primarily?

I suspect people with fast, reliable network connections will use BFO the most. People will like being able to carry only a small image around with them and be able to boot any current Fedora release. You'll never have to burn DVD media again.

More of Mike's interview, including behind-the-scenes details of setting up boot.fedoraproject.org, is available at https://fedoraproject.org/wiki/Boot.fedoraproject.org_and_you.

Feature profiles for Fedora are created by the community-run Marketing Team; full versions of the excerpts below can be found at https://fedoraproject.org/wiki/F13_feature_profiles.

Easy Print : Tim Waugh
Interviewed by Robyn Bergeron

Give me an example. Let's say I go and pick up a new printer at the Giant Electronics Store - what is the process I would go through to install a printer? How technical do I have to be to make this work?

The idea is that you don't have to be technical at all. You've come home, switched on your computer and logged in, unpacked the box, and plugged in and connected the printer. The very next thing you should see is a dialog box offering to download and install a driver. After that, you'll see a pop-up notification telling you that it is ready to print.

If the driver is already installed, of course, you'll just see the notification that the printer is ready. The Fedora 13 feature is to make sure we have a driver installed that will work.

Color Management : Richard Hughes
Interviewed by Robyn Bergeron

What are the compelling reasons for me, an average user, to use Color Management?

Sure, I would say I'm probably the average user in this regard (which may surprise you) -- I don't understand all the color transformation maths, I just want to make things work with as few clickety-buttons as possible. This whole GNOME Color Manager adventure started with my hobby, digital photography. I spent a few months taking photos which looked great on the viewfinder, which then looked "wrong" on the screen, and even more "wrong" when I printed them out. Tweaking the colours (I'm British, really!) one way or another, which seemed really unscientific and a bit of a black art. I just wanted the screen colors to match the camera colors, and the screen colors to match the printed colors. The existing tools kinda-worked, but were all command line, and were really hard to understand and use. There was no integration in the desktop and they all had to be setup by hand...

I started the project as most projects do: scratching an itch. Things didn't work for me, and I figured they didn't work for other people either.

More from Tim and Richard, as well as Ben Skeggs on the open-source Nouveau driver for NVIDIA GPUs, is available at https://fedoraproject.org/wiki/Hardware_enablements_in_Fedora_13.

Feature profiles for Fedora are created by the community-run Marketing Team; full versions of the excerpts below can be found at https://fedoraproject.org/wiki/F13_feature_profiles.

btrfs : Josef Bacik
Interviewed by Hannah Kowen

What is Btrfs exactly?

It is a new filesystem thats trying to bring in some more modern features to Linux. For example, it has snapshotting and subvolume support built into the filesystem, instead of having to rely on another layer, such as LVM, to give you those features. It has some unique data integrity features, such as checksumming for all data and metadata and we use a method call copy-on-write to deal with general data integrity, instead of the traditional journaling filesystems like Ext4 and Xfs.

And to break it down a bit for those non computer scientists out there...

Snapshotting means that you can create an exact copy of the volume at a given point in time. This is useful for backup programs that don't want to have to deal with a changing filesystem. You can take a snapshot of the filesystem, and then backup from the snapshot. For Fedora we've used this feature to give us extra safety when doing yum transactions. So right before doing a big yum update, our plugin will take a snapshot of the filesystem, now you have two copies of the filesystem, and then yum does its update onto your original volume. Then if something goes horribly wrong, you can "rollback" to the snapshot that was taken before you did the yum update. You do this by making the snapshot the default volume to be mounted, so next time you reboot your system will be just like it was before you ran a yum update. Then you can mount the broken volume somewhere else and try and fix things manually

Btrfs is not the default for Fedora 13?

Oh no, its not ready for primetime yet. It's still very much an experimental fs that is under heavy development. A lot of the key features are there, but a lot of stabilizing and such needs to be done still

Fedora 13 users can opt to use it?

Yes, if you give the installer the "btrfs" option on boot it will allow you to install onto a Btrfs fs. It's been an option since Fedora 11. The box I'm currently working on is Fedora 11 with a Btrfs root. Of course it is much faster and much more stable in Fedora 13.

More of Josef's interview, including an overview of btrfs and RAID and instructions on how to test the feature, is available at https://fedoraproject.org/wiki/Btrfs_in_Fedora_13.
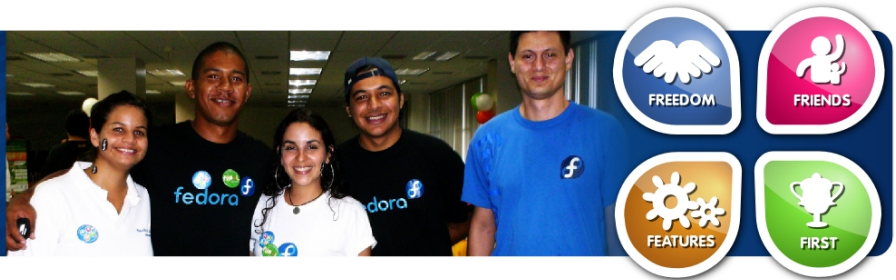
Feature profiles for Fedora are created by the community-run Marketing Team; full versions of the excerpts below can be found at https://fedoraproject.org/wiki/F13_feature_profiles.

NetworkManager : Dan Williams
Interviewed by Nelson Marques

Dan Williams is a developer who works on NetworkManager, a service that manages and reports network changes combined with a graphical desktop applet that allow users to manipulate network connections. Nelson Marques asked Dan about his thoughts on being a Free and Open Source Software (FOSS) contributor.

From your scope of view, what are the barriers that potential developers face when they try to join FOSS Projects?

Probably #1 is lack of upstream guidance; when a potential developer shows up, often upstream developers don't have a ton of time to help nurture that developer through some of the initial stages, because they have so much to do. It's a very fine line; you can spend 75% of your time helping new potential developers and then you don't hit your feature targets because you didn't spend enough time coding.

What helps that is documentation; if you put beginning documentation up somewhere then you can direct new developers to that instead of repeating over and over each time one shows up. Or, if there are other community members that can handle this duty while you concentrate on the tougher questions. Filling that role (called "developer relations" in larger companies) is a great job for somebody who can't commit tons of time to the project, but who is still willing to periodically build the source (even if they don't fix core bugs) and answer emails.

What would be your personal advice for potential developers who would willing to help with NetworkManager?

Keep asking questions; no question is too stupid or embarrassing to ask. Then, when you understand the solution to your question, tell us how to improve the documentation so other people benefit too. Since we've been involved in the project for a long time, we don't always know what the initial roadblocks are to building the source, understanding the architecture, etc.

More of the interview with Dan, including an overview of NetworkManager functionality and his thoughts on the major achievements achieved in the last 6 months on NetworkManager development, are available at https://fedoraproject.org/wiki/NetworkManager_in_Fedora_13.

Feature profiles for Fedora are created by the community-run Marketing Team; full versions of the excerpts below can be found at https://fedoraproject.org/wiki/F13_feature_profiles.

SSSD : Stephen Gallagher
Interviewed by Robyn Bergeron

Can you give us a brief description of SSSD - what it is, and what it does?

The modern SSSD is actually not a single daemon, but a collection of services that provides a common interface for user identity and authentication. The SSSD provides user information through the standard NSS (name-service switch) interface used by traditional identity services like nss_ldap and nss_nis. We also provide a PAM (pluggable authentication module) module to perform authentication.

My site's system administrator doesn't know about SSSD. What are the reasons I should give to him or her about why SSSD is going to be great for my coworkers?

There are a few features here that administrators are going to love. One of the big ones is that it will reduce load on their identity servers. In terms of authentication, administrators also have the ability to configure the SSSD such that it will only allow users to log in using cached credentials for a fixed time. The idea behind this is to ensure that a laptop has to check in with the corporate network every so often, to guarantee that the user is still allowed access. The user will see a message at login time when logging in with cached credentials telling them how many days they have left before they can no longer log in offline.

I understand that SSSD had a test day in March. Can you tell me how that worked? What kinds of problems were you able to solve by having a lot of community members participating in testing all at once?

It worked fairly well. We were actually surprised (and pleased!) to discover very few bugs in the SSSD itself. Most of the issues we encountered came from the authconfig UI (and have since been fixed). The real benefit of having the Test Day was that it allowed us an opportunity to get a feel for a wider variety of deployments. We were very pleased to see that, for the majority of testers, we were meeting their needs.

This is how good software is written; it's not just about having clever coders, it's about the people who care enough about what you're doing to step in and help you cross that final mile.

More with Stephen on SSSD, including the streamlined UI design for administrative setup and a more in-depth discussion on exactly how this feature reduces the load on identity servers, is available at https://fedoraproject.org/wiki/SSSD_in_Fedora_13.