

# “WE’VE HAD HUGE COMMUNITY SUPPORT. WE’RE DOING GREAT THINGS...”

We meet the LibreOffice developer to find out what next for the suite, and what it’s like to start a business supporting open source development.

**A**lmost three years ago, OpenOffice – one of the most iconic pieces of open source software – was forked and a new office suite was born: LibreOffice.

Michael Meeks was the driving force behind the technical advancement of the fledgling software, and oversaw the project’s rebirth into a cutting-edge suite. This has led to LibreOffice becoming standard on almost all Linux distros.

In 2013 he started a new business, Collabora Productivity, to support the development of LibreOffice. This new

venture – focused purely on LibreOffice – is moving the office suite forward at a phenomenal pace, and working with the UK Cabinet Office to ensure document freedom for the UK government. Michael is also the most enthusiastic person in free software, even when the subject is something seemingly mundane like bugfixing.

We caught up with him in Collabora’s Cambridge office to find out about the future of LibreOffice, the crazy things people do with spreadsheets and why the GPU’s days are numbered.



## What’s happened in the last three years since LibreOffice split from OpenOffice?

**Michael Meeks:** In three years, quite a lot has happened. Initially we set up LibreOffice and we got a lot of people suddenly appearing and wanting to get involved, which was gratifying. And that’s grown over time. We’ve now got [around] 100 people a month committing, and that’s trending upwards. That’s individuals giving us code. There’s a whole load more translators and other people involved: documentation, QA – there’s a huge amount of work going on. We had about 3,000 commits last month, which is a pretty good ratio: 100 commits per day. And it’s cool. Lots of features, lots of fixing, lots of clean-up, lots of core work, lots of features that people have really needed for a long time. For example, comment printing went into 4.3. It’s nice

to see these things finally coming in... finally seeing the spreadsheet performing, and having a reasonable design internally. It’s an exciting time with plenty of new stuff.

The project is going really well. We created a non-profit in Germany. We took about \$1m worth of donations and advisory board fees last year. We’ve had huge community support, and we’re doing great things with that.

## In every LibreOffice release notes, there’s a comment about cleaning up the internal code. How’s that going?

**MM:** Pft! Haha! There’s a huge amount to do there. We’re doing pretty well. One of the things you see in 4.2 for the first time complete is that we used to have four string classes and now we have two. So this means that in 4.2 we have long paragraph support.

You wouldn’t believe it, but in Spain, for some reason, when they take legal notes they’ve all got to be in one paragraph. Previously you could only go two pages or so with one paragraph. Now, the world is your oyster. You can go, really, a very long way. There are some fixes like that that you can start to see. German comments: we’ve come down from something like a hundred-and-something thousand lines to something like 25, and we’re starting to shrink that further.

We’ve been re-doing all the dialogs. They used this horrible old way of creating dialogs, and now they use Glade, so they’re all easy to tweak and anyone can work on them – that’s now 85% complete. Hopefully by our next release, that will be done.

If you look at the Coverity stats (Coverity does this static analysis and checking of the code). There were



**“Towards the end of this year you’ll see an Android viewer coming out, and probably an Adroid editor – we’re still looking into that.”**

10,000 or so of these potential issues, and again, we’ve fixed vast numbers of these and the number is going rapidly down. The new build infrastructure is now complete and that’s been there for two releases. But there’s still a huge amount of clean-up – these percentages hide problems – and of course when we do that, we create more problems.

Uncovering new bugs before the users see them is important. One thing Marcus Mohrhard is doing is taking every bug document we can find, so we scrape our bugzilla, KDE’s, Gnome’s, Free Desktop’s, you know, anywhere we

**“It’s a never-finished problem, but we’re trying to make it a little better everywhere.”**

can get documents from, and we’ve ended up with something like 45,000 problem documents that have been nasty to deal with in the past, and then we run them through LibreOffice. We also compile with all of our debugging assertions on, so we’re in paranoid mode, and we run this in a loop. It takes a week or so each time.

That gives us a whole load of bugs, which we’ve fixed, but from time to time new ones come. We’re starting to export them to every format we support, so not just loading 45,000 thousand documents, but re-exporting them as PDF, XLS, ODF, XLSX etc. That takes longer, but also finds lots of errors. We’re really starting to build the infrastructure to tackle these things and make that visible to people.

It’s a never-finished problem – that’s the punchline – but we’re trying to make it a little better everywhere.

**LV Looking forward over the next few years, what are the big things coming up for LibreOffice?**

**MM:** Towards the end of this year, you’ll see an Android viewer coming out, and probably an Android editor – we’re still looking into that – but certainly a viewer. You’ll see an iOS version before then. There’s a company working on that.

In terms of online versions and integrations with cloud-y things, those are probably further out, but it’s all a matter of money. These things aren’t terribly expensive, but it’s finding people to pay for them. What else? We’ve seen OpenCL enablement of the Calc core – you can use the GPU to calculate spreadsheets, which is pretty exciting. If you look at the industry as a whole, we’re moving to GPUs. Discrete GPUs will probably die.

Previously, the GPU was stuck on a card somewhere else, so if you wanted



LibreOffice is now part of the PCMark benchmarking suite – that's a major win.

to send a texture over to this and get it back, you have to put it in a special bit of memory, and this is specially DMA'd over and you do something and it specially comes back later, and this is lot of kernel transitions, round trips, pain and aggravation. You have to copy everything: put it in a physical bit of memory somewhere that is in the DMA aperture. The next evolution is these Intel integrated graphics chips. They share the same memory, but again you have to move it into an aperture. Because you're talking physical memory not virtual, any pointer you have in this space is meaningless to the graphics chip. It will see it as a number, but it can't see where this is in physical memory. Traditionally, they could only map a small part of physical memory.

HAS [Heterogeneous System Architecture] is a thing from AMD, Imagination, ARM, etc. The basic idea is to give the GPU virtual memory address space, so the GPU and the CPU are both able to work on the same chunk of memory, follow pointers, follow linked lists, do stuff, and this means that you can move work across to the GPU much more easily: zero copy, not even having to manipulate your structures, so you can start to get the GPU to act on something – like spreadsheets for example – so you can get an advantage with a very small amount of data.

**LV** It's surprising that people use spreadsheets to the extent that they need GPU acceleration.

**MM:** It's surprising what people do with spreadsheets, and they typically extend

them to the point that they're too slow, and only then do they typically do something about it.

**LV** But then they get OpenCL and they can make them even more hideously complicated!

**MM:** Yeah. Then they can do stupid stuff even more, and it actually make the thing glow. Even HR people, you'll see them doing `vlookup()`'s on salaries and trying to find out how many hours a week they spent in each area for each of these people in each day, and the days keep growing, and it's building a huge table behind the scenes to say: bill this much for this month, etc.

**LV** How do you decide which of these things are important?

**MM:** AMD came to us and encouraged us to work on this. It makes sense, because it's power efficient – the CPU is

clocked at gigahertz, and the GPU is clocked much lower and it gives a much wider power profile.

That's GPGPU. There's loads more we can do there – faster JPEG loading, faster rendering into the graphics domain, GL stuff. It's the future of rendering. Even Microsoft is using it now. The beauty of WebGL is that it means you have to have GL everywhere, so finally there's grunt behind this standard that means it'll be ubiquitous. It's just a shame that there's some divisions in the GPGPU space: there's OpenCL, but there isn't WebCL, there's Google's thing which is a shame. GL now works on Windows, Linux and Mac, so it's brilliant to have that rendering kit.

In LibreOffice on Windows, you have a scene that's rendered in GDI (which is the native toolkit), and we want to composite things on top of this. So you have the GDI rendering – which is hardware accelerated anyway – going into a texture, but you can't get that texture to composite another GL thing on top of it, so you end up with a square window which is really horrible.

**LV** What else is coming?

**MM:** Better integration, document management systems. There's this thing called LibreOffice Kit that we're producing that enables you to reuse LibreOffice on the phones. They're better document indexing and conversion in the background.

**LV** It sounds like you're doing much more support than you were five or six years ago.

**MM:** This is really a Collabora thing. Collabora Productivity is really doing a



Michael left SUSE in 2013, taking most of its LibreOffice developers with him to start Collabora.



**“We have to sell a lot, so I spend a lot of my time selling consultancy – but I love to code”.**

lot of that, and we love people to use LibreOffice. If they can't because there's no support, we'd love to support them. If they can't because it's missing a feature, we can close that gap and off they run. I think we're lubricating – getting LibreOffice into everywhere.

**LV You've been involved in LibreOffice and OpenOffice for a long time. Have you seen it become more popular in the commercial world?**

**MM:** I think there are certainly areas where it's very popular and it's growing quite rapidly. I think a lot of people are using it without paying anything. If you want to see a 150,000-seat deployment, you can go to Spain, and that's great. What we really want to make sure is that when people deploy it, they have

**“People feel that they are a part of LibreOffice in a way that I never saw with OOo.”**

services and support so that we can reinvest in making the product better. We want to make a virtuous cycle so that we can grow. Lots of business are doing that. If you look at the companies around the ecosystem, many years ago it was Sun, Red Hat, Suse – that's about it. These days, there's not Sun or Oracle,

but there is Red Hat and Suse, and there's a company called CloudOn which is investing a lot of money. There's Igalia, which is a consultancy doing a whole load of stuff. Collabora of course, Synerzip, Atomic, Ericsson – lots of companies are starting to get around this.

**LV Is that something that's part of LibreOffice, or something that's come over from OpenOffice?**

**MM:** I think that's part of LibreOffice. It is a vendor-neutral space where everyone is equal and we want people to get involved. No one is privileged – Collabora Productivity may employ the largest group of the most skilled developers (we have something like 12 certified developers) and nine out of the 20 top committers, but certainly not 100%, and we're eager for competitors to come in. With my board hat on, I spend time building business for our competitors, so there's no one who dominates it.

**LV Apache is still developing OpenOffice. How do you see that going?**

**MM:** We're diverging pretty rapidly. I think virtually everything I've told you is distinctive to LibreOffice thus far. We get commits from them that we include where they're suitable. Not every commit gets included. Two thirds do.

People feel that they own and are part of LibreOffice in a way that I never saw before with OpenOffice. It is something that people can be loyal to, passionate about and spend their evenings going that bit extra. They're unhappy when they get pulled off onto some other consulting project because they really want to make this absolutely beautiful for the project. It's a great feeling.

**LV Do you spend most of your time now managing people and talking to the press?**

**MM:** Gosh! I don't talk to the press nearly enough! Almost no one knows about Collabora Productivity, which is a shame, but there's quite a lot to do.

I love to code – that's the problem – but I have to sell, so I spend a lot of my time building pitches, investigating what can be done, going to companies and saying “hey, would you like to invest in this”, and chasing that sort of thing.

**LV Is there anything you've learned from putting together pitches that could be learned by other open source projects?**

**MM:** Let me give you the LibreOffice pitch. “It's easy for me to sit down and show our financial controller a feature she doesn't know in Excel or LibreOffice, and she works in it all day. There are features that could make her more productive, but since people don't really use most of the features, two questions: Why not? And why are we paying for them? You can actually solve both of these in a single blow using LibreOffice. We'll provide you with a subscription for about a twentieth of the cost of Microsoft Office and secondly you can use the money you've saved to train your staff who think they know how to use the software. You're probably going to have to do that anyway to move to LibreOffice. What you discover when you start training people is that they're like ‘wow, it can do that’, or ‘oh! It does this’. Of course, there are some rough edges, but we can fix those. People can become more productive, and save money and move to LibreOffice so you win in every direction.”

That's the pitch. I think we can save people a significant amount of money, we can move them to open formats. There are lots of reasons people want to do that, cost saving is one, but common sense is better. There's security – we live in an era that's very concerned about documents and leakage and so on. In the UK, the cabinet office are doing a fantastic job at the moment. They're doing a lot of good things

**LV They certainly seem to say a lot of good things...**

**MM:** I think those have an impact, but it always takes time. We've been involved with the ODF consultation. If you want competition, and you want lots of implementations you can choose from, ODF is a great choice. To privilege the incumbent (with 90%+ market share) by using a format they created as an exact cardboard cut out of their suite is a really foolish way of encouraging competition. Lucky, the Cabinet Office has done the right thing and chosen ODF, and they seem to be still going through the process of making that recommendation, but I think that's a sign of deep clue-fullness. LV