

THE COMPLETE GUIDE TO **OPENSTACK**



CREDIT: ISTOCK

Produced by  IDG

Open for business



Open source has gone from strength to strength in the preceding decades and is now widely considered a core component of enterprise IT. It now powers much of the world, whether that is from Linux supporting the Internet, to the way software is developed for both consumer and business applications. With the advent of cloud, OpenStack came into the fore, a collection of open source cloud projects that hoped to transform the underlying infrastructure in the public and private cloud landscape.

The platform – which now supports everything from bare-metal provisioning at on-premise data centres in the Ironic project, to virtualized GPUs in Nova – has undergone a series of iterations over the years. Its latest fits in with a new parent umbrella organization, ‘Open Infrastructure’, that advocates open infrastructure projects whatever they might be.

There have been ebbs and flows of vendor support, and in its early days, OpenStack was derided as being overly complex, difficult to configure, and difficult to manage. But recent cycles of the platform have seen the community pushing to address that. There can be no question today that OpenStack works in the real-world: it’s got backers in heavyweight rivals Red Hat and Canonical, as well as in telecommunications providers such as AT&T. And it powers one of the world’s most important scientific research facilities, CERN.

In this month’s Complete Guide we run through everything you need to know about the state of OpenStack today, the key players in the ecosystem,

how to become an OpenStack engineer and some case studies of organizations committing to the open source infrastructure project. [Tamlin Magee](#)

Contents

- 4** OpenStack Stein offers range of enhancements
- 9** OpenStack Foundation branches out to 'open source infrastructure'
- 15** OpenStack moves one step closer to the edge
- 19** How to become an OpenStack engineer
- 24** Ubuntu/Red Hat drama erupts at OpenStack
- 30** Red Hat confirms its commitment to OpenStack for 'at least' 10 years
- 36** Amadeus picks VMware Integrated OpenStack for private cloud



Credit: iStock

OpenStack Stein offers range of enhancements

Improvements, including updates to container functionality, edge computing and networking upgrades, are all part of the latest release

The recent release of open source infrastructure platform OpenStack, called 'Stein', was released last month with updates to container functionality, edge computing and networking upgrades, as well as improved bare metal provisioning and tighter integration with popular container orchestration platform Kubernetes – led by super-user science facility CERN.

It also marks roughly a year since the OpenStack Foundation pivoted towards creating a more all-encompassing brand that covers under-the-bonnet open source in general, with a new umbrella organization called the Open Infrastructure Foundation. OpenStack itself had more than 65,000 code commits in 2018, with an average of 155 per day during the Stein cycle.

Magnum, the certified Kubernetes installer, can now launch clusters far quicker than previous iterations – down from roughly 10 to 12 minutes per node to five minutes. Integration has been improved due to increased collaboration between the two open source heavyweight platforms, both from the OpenStack development side and the Kubernetes side – such as with the ‘OpenStack Cloud Provider’ service found within Kubernetes.

“That’s happened over the last year through involvement upstream in the Kubernetes community,” executive director of the OpenStack Foundation Jonathan Bryce tells us. “On the Magnum side, it’s taking that work and also parallelizing the way the provisioning actions are actually run – so that instead of saying you want a five-node Kubernetes cluster, you want a 50-node Kubernetes cluster, and it can basically spin up those machines in parallel and bring your environment up much more quickly.”

CERN, which has long been running an OpenStack environment at massive scale, had become a top contributor to Magnum, including leading the project. According to Bryce this user-led involvement means that the tighter integrations are in line with the needs of major users of the service.

Another user-led improvement for the Stein release is the clustering service Senlin, which has been driven

by gaming giant Blizzard. APIs in this service now issue synchronous failures in case of problems with clusters or nodes, cool-downs, or conflicts at the service level.

“They’ve had the project lead on that project for a couple of cycles and obviously have been a big contributor to Senlin over the years, too,” says Bryce. “They had an environment with multiple data centres and multiple OpenStack clouds, and they wanted to have the ability to spread out and automate their operations across those.”

Improvements have also been included to the resource reservation service Blazar, with a new resource allocation API that provides operators with visibility into their reserved cloud resources.

And a new project, Placement, which originates from the Nova compute service, promises to simplify specifying a host for workload migration. The project, according to the foundation, increases API performance “by 50 percent” for common scheduling operations.

Identity service Kolla has been refined for better multi-factor authentication receipts, while Kolla has added support for full and incremental backups of MariaDB.

In a press statement, CEO of OpenStack cloud business Vexxhost, Mohammed Naser, says that all these improvements together speak to the open infrastructure platform being much more simple to deploy and run, especially from the operator perspective.

Naser tells us: “The level of stability has significantly increased... This has been really helpful because as we try to continue to run the latest release as it comes out, obviously one of the bigger concerns is making sure we don’t break our customers that are relying on those APIs. I think that what’s happening is that the

operator community has been much more involved. The indirect result of that, is there isn't this weird gap between the developers and the operators, and so a lot of the additions being made are very conscious, to make sure that we don't break anything down the line."

Bryce draws attention to the improvements in the IronC bare metal, which have been incremental over the past few releases and is now getting to a state where users can manage physical servers in the same way they have managed virtual servers. In other words, they're much faster and more user-friendly than might be expected for physical server management, with more automation and scalability capabilities, too.

"Physical servers aren't the right answer for every workload, and virtual servers aren't the right answer for every workload, so for OpenStack to be able to provide an elegant way to provision and manage both types of resources is a super valuable feature," he reveals.

The 'Open Infrastructure Summit' will take place this April in Denver, Colorado. Commenting, Thierry Carrez, VP of engineering at OpenStack explains that the Open Infrastructure branding is already encouraging a better environment of collaboration with other users and projects outside of its umbrella.

"We have run a few Open Infra days already, the new format for what used to be called OpenStack days, and as a result we've seen a lot of involvement from adjacent communities that have joined the discussion," he says.

"Through those events I felt very much they were here to collaborate and discuss integration points rather than be in a defensive mode, where they would be at an OpenStack conference and talking about their technologies... so we can already see the positive

effects in terms of branding, and we are really consonant that at the upcoming Open Infrastructure event in Denver, it will have the same kind of vibe to it, around facilitating open development of open infrastructure.

“It’s not about one project or the other project, but about making sure open infrastructure actually works for users.” **Tamlin Magee**



Credit: iStock

OpenStack Foundation branches out to ‘open source infrastructure’

OpenStack has reached maturity to a degree that the Foundation is moving towards the goal of delivering open source infrastructure

Changes are underway at the OpenStack Foundation, with the community branching out from the open source cloud platform to the umbrella of ‘open infrastructure’.

OpenStack is the catch-all term for a series of open source cloud infrastructure components, ranging from

bare metal provisioning (IroniC) to networking (Neutron) and compute (Nova), to name just a few.

Originally emerging from a joint project between Rackspace and NASA, OpenStack has travelled through various development models – such as the ‘big tent’ [link] approach, which essentially said that anything from any vendor could ‘be’ OpenStack.

Along the way vendor contributions have waxed and waned, but now the core software is reaching a level of maturity that has given the Foundation enough confidence to branch out, and try to foster integration between all of the open source projects that are out there, not just the ones associated with OpenStack.

The ‘OpenDev’ conferences now runs concurrently to OpenStack, and another project focusing on open infrastructure – rather than OpenStack per se – is the Edge Computing Group. This has been around since an OpenDev edge event in September 2017, and was formed into an official mailing list at the November summit in Sydney – and continues weekly calls every Tuesday.

“The open infrastructure theme is coherent,” says Canonical founder Mark Shuttleworth. “In the past I have been critical of the inability of the OpenStack community to be able to draw a line between the stuff that was really important to build an OpenStack, and the stuff that was sci-fi and possibly a bit of a sideshow.

“I think this is different. It is essentially generalizing from OpenStack to other classes of infrastructure, and that’s interesting... I think the open infrastructure meme rather than an OpenStack meme is interesting. I think they will need to attract a set of projects which share some values, so the same audience can dip

into those different things. They need to be mindful of the cost of complexity,” he adds.

The OpenLab initiative meanwhile is dedicated to getting OpenStack integrated with the broader infrastructure ecosystem.

Governance board member Melvin Hillsman tells us: “OpenStack for some time has been very specific about testing and integration for OpenStack services, focusing only on the projects started at OpenStack.

“It’s been working very well, but for me as a person in the user community, my getting involved in OpenStack was more on the operator-user side. So it would be great to have some of these integrations, features, ideas, for the community above where the OpenStack APIs end.”

Kata Containers

The new projects don’t require that user to run OpenStack. At the first of last year’s summits in Vancouver, Kata Containers 1.0 was released – with the backing of major chipmakers AMD, Intel and Arm, as well as 40-plus other partners.

It’s designed to be architecture agnostic, billing itself as a project that delivers the speed and performance of containers with a boosted layer of security, because each ‘container’ is actually wrapped in lightweight virtual machines. Kata is based on Intel Clear Containers and runV from Hyper.sh, compatible with the Open Container Initiative and the container runtime for Kubernetes.

Zuul

The Foundation, which has a board ranging from all kinds of vendors and organizations and from all corners of the world, also announced ‘Zuul’, an open source

continuous integration/continuous development platform designed to help users ‘stop merging broken code’. It specializes in gating changes across systems and applications. A version of Zuul had already been in production and at scale for users running OpenStack since as early as 2012, but the Foundation decided to ‘decouple’ it from OpenStack-specific systems. Now it can be run on platforms like GitHub, and has contributors from BMW, GoDaddy, Huawei, Red Hat and SUSE.

Maturity

When OpenStack first came onto the scene it was met with excitement by the open source community and a wide gamut of vendors from various industries, all keen to collaborate so they could also steer the ship, as it were. From its inception, the question that was repeatedly aired at the biannual Summits put on by the Foundation was whether it was mature enough for production.

Fast forward to 2019 and that is no longer the major question: the proof is large-scale deployments from some of the biggest organizations in the world, from telecommunications networks through to enormous research projects and technology vendors themselves. Criticisms continue to be about reducing complexity, as running OpenStack still requires a certain level of technical expertise that can be daunting for businesses without the resources to learn it.

General manager for OpenStack at Red Hat Radhesh Balakrishnan tells us: “What we are seeing fundamentally is the maturity of adoption, which has spread across multiple verticals. Clearly, telcos is by far the number one vertical we are seeing OpenStack being adopted

– that constitutes about 39 percent for us in terms of make-up of customers.”

There is also movement in financial, public sector, media and technology “in that order”.

“If you look at the use cases outside of telcos, what we’re seeing is a lot of big data, with which I will include Apache Spark, Hadoop, data lakes and all that – that is one of the key use cases we have seen recurrent across all the verticals we’re working on.”

On the open infrastructure side, Balakrishnan says that for him and Red Hat, the “obvious focus is bringing the diverse communities to make open infrastructure happen faster”.


“That starts with Linux, to KVM, to OpenStack, to Kubernetes, Ansible, the Open Containers Initiative, CRI-O – ODL for networking – bring all of that together, so the core open infrastructure that is needed can be brought together. There’s a huge opportunity for that.”

But, says principle analyst at Hurwitz & Associates Jean Bozman, recent initiatives are aimed at making OpenStack more consumable and “easier for AppDev, DevOps, and system administrators to use in their enterprise infrastructure”.

“It’s clear that the OpenStack technology is maturing, especially for the core projects like Nova, Cinder, and Neutron, among others,” she continues. “OpenStack software is being adopted by large customers, including Walmart, CERN, eBay, China Mobile and Verizon for their clouds – and by small/medium organizations, often through partnerships for projects, for example, Red Hat, Google, Rackspace, Mirantis.

“The biggest challenges for OpenStack include wider adoption of the technology especially by cloud

service providers and by enterprises that are building private clouds. OpenStack is making progress regarding wider adoption, especially among large enterprises, government agencies – in the US, EU, and Asia – and scientific/university uses.” [Tamlin Magee](#)



Credit: iStock

OpenStack moves one step closer to the edge

We round up everything from the OpenStack Summit Berlin 2018, including its moves further into edge computing

The OpenStack Summit Berlin 2018 was the last of its name as it will rebrand as the Open Infrastructure Summit in 2019, a move that seems largely in line with the evolution of the open source cloud platform as it shifts further into edge and builds out a series of related pilot projects with OpenStack as the core proposition.

Many of the keynotes at the event showed the progress that the community had made in building out the pilot projects announced at the Vancouver Summit earlier last year. One in particular, the first release of StarlingX, might well help cement the open infrastructure platform in edge.

StarlingX is branded as an open source edge platform, with telecom and IoT use cases in mind. According to the Foundation, it “leverages components of Ceph, OpenStack and Kubernetes, and complements them with new services including configuration and fault management”, in particular to address technology challenges around high availability and ultra-low latency compute.

Then the Airship project is designed to help users deploy and manage containers, virtual machines, and bare metal environments – it’s got the backing of AT&T and SK Telecom for 5G to NFV, VDI, and big data processing.

Although edge is a nebulous term, it can loosely be described as distributed compute nearer the edge device. For example, in autonomous cars where milliseconds count for a lot, it makes more sense to have compute closer to the device rather than bouncing back to data centres potentially hundreds of miles away.

There is a firm perception that OpenStack is for telcos, and with 5G en route in cities around the world some time soon it’s not exactly a bad place to be. As AT&T demonstrated its early 5G edge projects deployed with OpenStack, telcos are set to converge in the compute space, with the new wireless networks enabling low-latency cloud computing nearer the customer – whether that’s in enterprise, IoT, or content delivery.

Updates were also announced to the Intel and Huawei-backed Kata Containers, which since its launch in December 2017 now support architectures, including AMD64, ARM and IBM p-series. Kata are lightweight VMs that ‘feel and perform like containers’, but with the security advantages of virtual machines, for example, workload isolation.

One of the Foundation’s occasionally outspoken community members, Canonical founder Mark Shuttleworth, reaffirmed that he believes the open infrastructure piece is a good fit for the organization, as he agreed six months earlier at the Vancouver Summit.

He said: “I think the idea of open infrastructure is exactly right... the future is already here, it’s just not evenly distributed. If you look inside a Google, it’s open infrastructure, so that’s the future – that’s the future for everybody, it’s just not evenly distributed.”

Although he added that if he had one concern it would be the potential for a loss on the core OpenStack projects cloud.

Lauren Sell, VP for marketing and community services at the Foundation rebuffed this during a media round table. She said that the way OpenStack fit in to the open infrastructure model was “pretty analogous” to the way Kubernetes fits into the work of the Cloud Native Computing Foundation – that is to say, a star project with others running alongside it that are related and still part of the ecosystem.

“We’re really focused at that infrastructure layer and the projects we’re bringing on are really connected and relevant to OpenStack,” Sell explained. “Even with StarlingX, for example, that’s probably the one that has some of the greatest overlap with OpenStack

– literally using OpenStack services and they’re contributing upstream to OpenStack, and they have these additional modules, but it’s really helping to take OpenStack into some of these new industries like manufacturing. It’s making OpenStack more accessible for some of these industries.”

If one of the main perceptions around OpenStack is that it is a telco product, the other is that it has a tendency to be overly complex to deploy, maintain, and integrate. But some of the developments with these complementary projects are helping to not only simplify OpenStack but build bridges into other communities, for example, Kubernetes. [Tamlin Magee](#)



Credit: iStock

How to become an OpenStack engineer

OpenStack has hit the mainstream. Is it time you learned your way around the platform? The community offers some advice

There are some preconceptions about open source cloud platform OpenStack: that it is overly complex, or that there is so much noise from so many vendors it's hard to know where to begin. Another challenge, tied intrinsically to these, is that it can be hard to find the right people, so it's arguably a shrewd move to get acquainted with the technology as a career move.

SUSE's Alan Clark, director for industry initiatives, emerging standards and open source, is also chairman of the board of directors at the OpenStack Foundation. And according to Clark, becoming an OpenStack engineer is a "great career move" for a couple of reasons.

"Centring your career on open source is a smart career move because most employers today are scrambling to embrace open source in research and product development," he says. "With open source adoption comes the need for skilled engineers, so open source experience on a resume is a great asset in today's market."

Clark adds that it doesn't take much more than a cursory glance of job boards to see the demand for not only open source, but OpenStack talent.

"The OpenStack ecosystem continues to grow and expand – demand will continue to grow for several more years," he explains.

"OpenStack engineering opportunities will continue to grow and evolve with the advent of additional innovation and technologies," Clark adds. "That spells opportunity for OpenStack engineers to tailor and grow their knowledge, experience and job focus."

One that OpenStack Foundation executive director Jonathan Bryce concedes is tough can be finding qualified engineers in the first place.

"Many organizations are already using OpenStack to automate their infrastructure, but finding qualified engineers can be a challenge," Bryce tells us. "For someone who has built a career administering servers and networks, there's a huge opportunity right now to extend their skills into the cloud world. OpenStack is available directly from all the major Linux distributions,

as well as other commercial and free versions – the barrier to entry to start learning is very low.”

Anyone with knowledge of administering Linux servers or virtualized environments will already have a decent baseline of knowledge for moving over to OpenStack, Bryce argues. Once that interest is piqued, there are resources online and in person where you can better acquaint yourself with the ins and outs.

Bryce suggests finding local user groups to attend meetups – there are regular meetings in many of the world’s major cities. The Foundation has made books and documentation available for free online too, such as the Operations Guide, available at fave.co/2PjpeJb.

He recommends more in-depth training, such as the Academy workshops running during the OpenStack Summit, and outside of this there are in-person boot camps available or courses from other official OpenStack training providers, listed here: fave.co/2verMix.

Practically speaking, like Bryce, SUSE’s Alan Clark points to the community itself for engineers starting out. “There is a large array of detailed technical information and road maps available through the community web interfaces,” he reveals. “This is a great place to start learning, but also there is open access to the mailing lists and community events, like the Summit we just had in Barcelona.”

He adds: “Once an engineer has a basis of information, I recommend they reach out to projects of interest and begin to involve themselves. It can be quite educational, and career enhancing, to contribute to a project – there’s a wide range of opportunities available.

“And I’d also strongly recommend demonstrating gained knowledge and expertise through a successful

completion of the Certified OpenStack Administrator programme,” he says. “This provides real evidence of proficiency to a future employer – tying this with a history of contribution conveys a unique skill set and experience that will be a huge differentiator in today’s market.”

Jonathan Bryce tells us: “You can build your own OpenStack cloud using packages from free community distributions like Ubuntu, Debian, OpenSUSE, or Fedora. Or commercial versions of the most widely deployed projects, Keystone, Nova, Cinder, Glance, Neutron, Swift, and Horizon, and improve your knowledge of running OpenStack services in a more real-world configuration.”

Bronsilav Kantor is a solution architect at Finnish IT services business Tieto, and he has real experience of getting to grips with OpenStack through a practical deployment with professional training, and his own initiative, running parallel to this.

Tieto began working on a cloud project a few years ago, including a portal for users to deploy and manage applications – Tieto deployed OpenStack internally and then began building an OpenStack cloud in its lab, too.

“I was looking for some information, some skilled people, but there weren’t so many,” Kantor recalls. “So I tried to find some good training and found Mirantis – it was good, although it didn’t cover all of my expectations – but I found very skilled people and participants. It was interesting to share the experience between us.”

Kantor does warn that the OpenStack landscape can be a little daunting at first. “There are so many services in OpenStack that it’s hard to find the services you should select,” he says. “Fragmentation of OpenStack is too big – I don’t have enough time to follow all the news.”

But his experience of the training was overall positive. “There were some nice lectures about debugging OpenStack, which is needed,” he enthuses. “If you have any problems it can be a stressful situation, so we learned a lot about how to solve your problems quickly on this training.”

“I think practising is important,” he says. “You should have courage, just install it and use it. I think some engineers don’t have enough patience, and motivation is important. If you really want to achieve something, you can learn it from OpenStack documentation – I did it like that. The training was more to confirm that I understand it well enough.” **Tamlin Magee**



Credit: OpenStack Foundation

Ubuntu/Red Hat drama erupts at OpenStack

Canonical founder Mark Shuttleworth opened a war of words with Red Hat on stage at OpenStack Summit Vancouver 2018

Canonical founder and part-time hobby astronaut Mark Shuttleworth began his OpenStack keynote at last year's Vancouver event by criticising Red Hat for being too expensive.

Among the nervous laughter someone at the front yelled "fact check", but Shuttleworth insisted to us that no one has refuted his numbers.

"Good luck to it," he commented when asked about the fact-check-heckler. "Everything I put out there is

supported by customers. There have been a number of analyst surveys of cloud economics that put Canonical on top of the list in terms of doing it efficiently.

“I put that message out there because I think this community – the engineering side, the community side, the vendor side – has to realize that in a public cloud-oriented world, private data centres can’t be steaming piles of nonsense.”

Directly after the keynote we heard some Ubuntu staff backstage were whooping following the presentation (with one apparently saying that Shuttleworth ‘smashed it’) and we also heard at least a handful of people refer to the keynote as a trolling attempt.

But Shuttleworth seemed to be deadly serious.

He said that large enterprise data centres tend to be full of “steaming piles of yesterday’s decisions” and that Canonical/Ubuntu is well-positioned to help businesses get their first OpenStack clouds up and running.

“They have to be efficient, they have to be automated, they have to be cost-effective,” he said.

When asked whether the presentation was community-spirited, he said that he “believes in open source and I believe in communities”.

“The Ubuntu community is arguably the biggest and most productive open source community and we do that by not pandering to people, by keeping it real, by having open conversations, and by inviting competing factions to work together despite the fact they’re competing.

“So I think I understand community. While I have had people express surprise because it’s not the sort of keynote anybody has seen here before... I think a lot of those people came around to the view that it is important for them to understand why OpenStack exists.

“You may be here because you love the technology, but the technology is only going to continue to get investment if it generates a return. It’s super helpful for people to understand it.”

Giving Red Hat right of reply, the company’s GM for OpenStack Radhesh Balakrishnan said that “different people make up a village, if you will... Our focus and remit is very straightforward. We continue to be very bullish about OpenStack, given that we have hundreds of customers in production. The last number we publicly talked about was 500-plus in production for fiscal year 18, in multiple verticals.

“And it’s not just OpenStack. It’s the full stack of open solutions that customers are consuming, to get to an open hybrid cloud.

“So we are excited about how far we have come in terms of the relationships we have established with these organizations, from being the Linux provider, to becoming cloud advisor for hybrid cloud adoption. We have stars in our eyes and we continue to get excited about our prospects, immediately as well as in the future.”

Talk about taking it on the chin.

Shuttleworth agreed that Red Hat is doing alright for itself. He said that there are only two vendors that he tends to see for building the architecture of a new data centre, and those are VMWare and Red Hat.

“You’ll find SUSE somewhere,” he added. “There are other folks who have versions of OpenStack on top of those distribution,s but it’s increasingly clear now that it doesn’t make sense to separate something like OpenStack from the underlying Linux, because it becomes impossible to know where to

go to get an issue resolved. All the folks who have OpenStack on top of somebody else's Linux have failed, effectively. That leaves you with Red Hat and Ubuntu.”

“VMware is obviously still very prevalent in enterprise data centres... It was the king of virtualization. But it's also a significant driver of cost in the data centre, cost which I think OpenStack is well placed to take out when it's done right.”

And what does he think of the direction the Foundation is steering the ship? This Summit, the Foundation made a point of talking about Kata, the open container project based on Intel and Hyper tech, as well as Zuul (CI/CD for testing) and OpenLab (for getting OpenStack up and running in potentially challenging outside environments). The messaging this summit was about delivering a wider ‘open infrastructure’, with all these projects affiliated to one another, but not OpenStack per se.

“I think the Foundation is being thoughtful about how to attract an audience that has a common set of challenges and interests,” Shuttleworth” argued. “The open infrastructure theme is coherent. In the past I have been critical of the inability of the OpenStack community to be able to draw a line between the stuff that was really important to build an OpenStack, and the stuff that was sci-fi and possibly a bit of a side show.

“I think this is different. It is essentially generalizing from OpenStack to other classes of infrastructure, and that's interesting. I think they could still do a sharper job of being pointed about what's OpenStack and what's not OpenStack.... We do that, that's one of the reasons we've been successful, because we didn't get drawn into an endless list of shiny possibilities.

“We deliver virtual desktop for a fixed price on a predictable timeline, and it works at scale. I think that’s what enterprises need and want from OpenStack. All the unicorns and fairies are interesting possibilities in the future, but they need to know that OpenStack can be delivered in two weeks for a fixed price, and our ability to stand up on stage and commit to that is unique.”

A couple of years ago Shuttleworth suggested that the ‘big tent’ approach to OpenStack (millions of vendors doing various things) would collapse and leave the businesses that were delivering the core. And two years on, he says he feels validated in that.

“I think the open infrastructure meme rather than an OpenStack meme is interesting,” he added. “I think they will need to attract a set of projects which do share some values so the same audience can dip into those different things. They need to be mindful of the cost of complexity.

“Complexity has a price attached to it. The person having fun introducing the complexity doesn’t necessarily appreciate that. I spoke forcefully in the keynote about cost because I think everybody should be aware of cost... It’s not just a licensing question at the end of the game, it’s about how the thing was designed. Was it designed to run efficiently, or was it designed to be interesting from a technical point of view?

“Those are things everybody needs to be mindful of, the community, the engineering people, the business people, and the Foundation themselves.

He added that he thinks the Foundation has gone through a “tough time” recently but that it has “rallied”.

“They responded, they deserve a certain amount of support for trying,” he said, before comparing the Foundation’s early days to a ‘hot debutante’.

“It’s pretty clear that you go through phases of a project where you’re the hot debutante and everybody wants to be your friend,” he explained. “Then there’s a new hot debutante somewhere else: whether it’s AI or containers, it doesn’t really matter. There’s always going to be something newer and shinier.

“I think the Foundation now is coming to terms with the fact it’s a mature project.

“As with everybody, once you get to middle age, dressing like a teenager doesn’t work. It’s not authentic, it doesn’t pull off, it doesn’t resonate.

“And so I think that drives you to a set of changes in how you programme an event and how you articulate a value proposition of your product and so on. That’s an adjustment for the Foundation. There’s fewer people here than there were two years ago. I think that’s good. There’s less hot air here, there’s less tire kicking here.”

Of course, he acknowledges the shiny stuff is also important. But the “truth is the world has to move on”.

“If OpenStack was new and exciting and shiny all the time then something would be wrong. And I’m excited by the new shiny too. But I am here to build clouds for people who need them and I’m proven at being able to do that.” [Tamlin Magee](#)



Credit: Red Hat

Red Hat confirms its commitment to OpenStack for ‘at least’ 10 years

Speaking at the OpenStack Summit in Berlin, two Red Hatters outlined plans present and future for the open source giant

At last year’s Berlin OpenStack summit, director of product management James Labocki and senior director for product management at OpenStack Nick Barcet confirmed Red Hat’s commitment to OpenStack for at least the next “10 years”.

The open source giant announced OpenStack Platform 14, which Barcet said aimed to make OpenStack

a better platform to run container orchestration system Kubernetes on, while also helping to better manage the deployment of Red Hat's container platform OpenShift on bare metal, as well as easing the integration of OpenShift and OpenStack at the networking and storage layer. This, explained Barcet, is a landmark move for Red Hat because it is part of a new strategy to focus on its whole portfolio as a single entity rather than individual products.

"The biggest problem with OpenStack is it's solving lots of problems, and different problems depending on how you configure it, which means that your initial understanding of OpenStack is going to be defined by the use case you're trying to solve," said Barcet.

"That can take quite a bit of time for people, and we failed as a community to make that part simpler. It would be really nice to have pre-scripted scenarios per use case you're trying to solve. This is something we have to improve over time."

But Barcet added that the company is seeing "lots of success" with OpenStack, ranging from general purpose infrastructure, big data deployments, and calculation environments such as for AI or standard high performance computing. "We are seeing a lot of people also using OpenStack not for the 'above the cloud' aspect but the ability to deploy workloads on bare metal, and that is something that is increasing really fast," he revealed. "More and more people are addressing their application deployment with a cloud API but directly on bare metal because they need to talk directly to some kind of hardware at the physical layer."

That might seem to go against the grain compared to many of the digital-native big brands that often

shout the loudest about their infrastructure. However, Labocki added that the company has announced new tooling to help mass migrate virtual machine-based workloads from proprietary virtualization directly onto OpenStack instead.

“Basically being able to discover them, then seamlessly move them over en masse,” he said.

“What we are starting to see is a lot of customers are recognizing there are diminishing returns in their virtualization infrastructure, and they’re starting to say: ‘hey, I want to invest in other technologies: Kubernetes, automation, multi-cloud management, you name it’.

“We’re seeing more differentiation and business value being driven out of that, as opposed to the standard virtualization layer, which is not necessarily delivering returns any more,” Labocki added. “We’re seeing a large amount of them starting to say, I want to reduce the spend here and migrate, whether it’s OpenStack or another KVM-based virtualization product.”

Barcet added: “I don’t know any Fortune 1000 company that has really migrated everything to public cloud. Some have made big announcements that they have the intention but the reality is some of their workload has to remain in-house for a variety of reasons, whether it is regulation, whether it is capacity.... It’s just not happening. We are also seeing customers swinging back because the cost of op-ex from running on public cloud for a few years just doesn’t make it.”

When it was announced in October 2018 that IBM was intending to acquire Red Hat there was widely reported speculation that it was primarily to strengthen the century-old corporation with a better standing in open source and crucially, hybrid cloud.

Labocki stressed the importance of open source past, present and future. “AI, machine learning, big data, microservices, containers, IoT, what are those built on? They’re fundamentally built on Linux and open source technologies. So at the end of the day organizations, if they are going to want to digitally transform, will have to embrace Linux and open source technologies.

“The reality is that if your staff is not getting trained in that, you’re going to be behind the curve on how able you are to digitally transform. And the nice thing about it is when you look at all those fundamental technologies of virtualization and infrastructure, all those, they’re built on Linux. The past is built on Linux, and the future is built on Linux, so I think it’s a pretty obvious choice if you’re a leader in a company you have to invest in Linux and open source skills to be able to do that.”

Barcet added that in his view Red Hat has balanced its development model and business model to achieve success, and being one of the relatively few number of highly successful open source companies.

“The development model is upstream first, it is massive investment into the community, everything is open, the business model is about delivering stability through a subscription model – and answering those complex problems that you have as soon as the software touches the hardware, because regardless of how fast we can make software evolve, hardware, the cost of racking something, hardware is always a drag. You’re going to have to stabilize the platform for at least three years, maybe five years, maybe 10 years, depending on your investment model.

“Red Hat has understood that and delivered while investing at a very fast pace in the community, a

subscription model that allows customers to deploy at their own pace, and maintain within the right range.”

When asked about how forward looking Red Hat is able to be about the coming months and years, Barcet said that it is committed to OpenStack “for at least the next 10 years” although this will probably be part of a wider “open infrastructure” project, and stresses the importance of 5G and edge computing.


“It is going to be used for a long time. Will OpenStack transform? Obviously it will, but it has been transforming already – public cloud, private cloud for telcos... It’s just a continued advancement. My only thing is I think we need to innovate faster nowadays than when we started the project, that’s why I’m suggesting the development cycle be shortened.”

“One of the big benefits of the open source model is the ability to pivot,” added Labocki. “If you go back to the Zen virtualization-KVM virtualization pivot that Red Hat did – no other traditional software company that had a traditional development model with proprietary development would have said we’re just going to switch our virtualization technology out in the next release. The idea that we can openly say we’re going to make a move and change: you can see the same thing on the container landscape, we embraced Docker and we embraced Kubernetes, we went all in on these technologies when we already had a container technology at Red Hat.

“So when you’re sitting in the audience and hear open infrastructure and you go, oh no, something is wrong – it’s actually no, that’s the way the open source development model works. You change quickly and you move.”

What else is in the future? Barcet explained that the company is heavily involved in edge computing pilots with a number of telcos, and the ability to manage multiple data centres that are nearer the customer as well as looking into the way 5G will impact manufacturing.

“One of the big problems when you’re manufacturing is you need to collect a tremendous amount of data to understand what’s happening in your production chains,” he said. “And sending this data over the wire may not make sense, so you also need to have a distributed model, distributed compute nodes or putting small footprint deployments even closer to your end user.” [Tamlin Magee](#)



Credit: iStock

Amadeus picks VMware Integrated OpenStack for private cloud

Global travel IT services company Amadeus wanted a solution that was flexible, scalable and useful for agile development

If you've travelled at all it's likely you will have crossed paths with products from Amadeus IT Group – whether that's the software that manages flight searches and hotel bookings or boarding a plane. Now the 30-year-old company has picked VMware to deliver a private OpenStack-based cloud, to support VMware vSphere workloads and as part of a wider digital transformation.

Amadeus built its own private cloud to better address the agility, stability and scale needs of its own internal operations team and those of its customers.

“[Travel] is a very specific segment and a very demanding environment,” says Udo Sebald, who runs service portfolio management at Amadeus. “The systems that we operate have to scale to a very large number of transactions – if you imagine we process 3.8 billion transactions per day in peak situations it gives you an impression. Our systems need to be very scalable and flexible to satisfy peak demands.”

The company had a long-standing partnership running with VMware in other areas of the business, and found success in network, infrastructure, storage and server virtualization.

“Now we not only provide the infrastructure as a platform, we also build the applications ourselves to a very large extent,” Sebald explains. “So we looked at what would be the next step after virtualization, how we can get to that infrastructure as a service environment, so we would be able to use not only our own infrastructure in our own data centre environments but also rely on third-party infrastructure, which we could use as if it was our own infrastructure.”

VMware integrated OpenStack

Amadeus’s infrastructure as a service platform runs on VMware Integrated OpenStack, NSX, vSAN and vSphere. The company says it specifically picked the OpenStack distribution for its integration with vSphere and NSX, as well as because it is conducive to agile development.

The company is not yet running a majority of its workload in this environment, but has made solid

progress in virtualization and using OpenStack. However, as customer demand increases Sebald believes more of this will be migrated over.

“These investments we are making are, for us as a company, pretty significant,” Sebald says, speaking about why the company plucked for VMware. “That’s why we want to make sure whatever we do, we do it in an environment where we can be reasonably sure those investments are there to stay and we’re not betting on the wrong horse.”

The company will measure the success of the deployment by examining how many of its processes it can automate, which “in turn relies on things being virtualized and things being on infrastructure as a service”, Sebald explains. “We’re basically measuring what level of automation we have achieved in specific processes, for example, in change and incident management. We also measure what level of virtualization we have introduced in the data centre.”

But ultimately, the company will be counting customer satisfaction as the main metric for success.

“If you look into the whole situation from the customer perspective, customers demand from us a certain level of service stability, they expect a certain level of agility, to be able to bring in changes, and for a certain cost ratio,” he says.

“Those are basically the parameters that steer our activities. If we cannot reach, for example, our system stability targets, we would not go for deploying a technology that would jeopardize this target. It’s all linked to those commitments we make to customers and how a technology such as OpenStack can support it. It’s all broken down from there, and so it needs to be seen

in conjunction with those contractual commitments.” There will be some challenges ahead – Amadeus is running a lot of software that will need to be migrated to newer technology stacks. It also faces industry-specific demands in terms of system monitoring. “You can imagine off-the-shelf tools more often than not do not satisfy the requirements we have, so we have to build a lot of this stuff ourselves,” he says. [Tamlin Magee](#)



©2019 International Data Group