


# The edge is open: Why scale-out computing doesn't exist without open hybrid cloud



By Paul Cormier Executive Vice President and President of Products and Technologies, Red Hat, Inc.

The past year has seen the rise of applications that push enterprise IT to the (literal) edge, from using autonomous vehicles guided by artificial intelligence (AI) to vast sensor networks that rely on 5G for instant connectivity and emergency reaction times. Whether it's the Internet-of-Things (IoT), fog computing or edge computing, the intent is to bring computing resources like processing power and storage closer to the end user or data source to improve the ability to scale, responsiveness and the overall service experience.

We can look at the edge as the newest IT footprint, becoming an extension of the data center just like bare-metal, virtual environments, private cloud and public cloud. In a sense, edge computing is a summation of the other four footprints, blending pieces from each to create infrastructure aimed at tackling specific customer demands that traditional IT models cannot address. But unlike the other footprints, edge computing has two key delineating factors:

- It simply does not exist without the hybrid cloud.
- The foundation of edge computing must be open or it WILL fail.

Bold statements? Sure, but from my point of view, they are wholly accurate. If edge computing is going to be a realistic future for enterprise IT, it needs the hybrid cloud AND open source to thrive.

Why edge computing is “hybrid or die”

Edge computing turns the concept of cloud computing on its head. Where “traditional” cloud deployments are about centralizing on a single infrastructure that can scale up as business needs dictate, edge is focused on “scaling out” geographically.

This could be small-form servers on cell towers, sensors monitoring a global energy network or next-generation factory automation systems that anticipate maintenance needs. Whatever the specific workload for edge, the need is the same: Faster responses for more timely services, whatever they may be. eBay, for example, is adopting edge computing by decentralizing its datacenters with the intent to create a faster, more consistent user experience by moving data and online services closer to users.

With the disparate nature of edge computing, consistency is key – an edge deployment could theoretically be hundreds of thousands of tiny sensors connected to a data aggregation tier which help to provide real-time feedback to what the sensors are actually monitoring. It's fundamentally impossible to manage each of these deployments if they don't share a more secure control plane via automation, management and orchestration.

This consistency is offered through the hybrid cloud – from edge devices to the network to the centralized datacenter, a hybrid cloud deployment provides sanity to what would otherwise be sheer madness across a technology ecosystem. It gives all of these varied components a common foundation upon which to rest, whether it's Linux, Kubernetes or Ansible, enabling IT teams to manage ten thousand networked devices just as they would their centralized IT.

Open not closed: True in the datacenter, true at the edge

Innovation can be easily crippled by fragmentation and, worse, the introduction of “proprietary” or open core models. You can look back 25 years to UNIX to see how painful this challenge was, where every hardware vendor had their own bespoke operating system that wouldn't work except on their own boxes. The introduction of enterprise-grade Linux broke this chain, and kickstarted the innovation cycle, leading to virtualization, cloud computing, containers, Kubernetes and, now the edge.

If this innovation is chained via “edge edition” technologies rather than being driven by common, industry-wide open standards, fragmentation rears its head again. This is where Red Hat, along with our industry partners, are driving the edge – through the creation of these standards in upstream communities and industry working groups, like the LF Edge, Kubernetes Edge & IoT Working Group, the OpenStack Edge Working Group, Akraio Edge Stack, OPNFV and more.

This isn't to say that every edge solution is going to be based on open source technologies, despite what glossy marketing materials from the vendor world might say. A commitment to open source doesn't only apply when convenient or when something new and shiny, like edge computing, comes along. At the same time, layering proprietary features or capabilities onto an “open core” is antithetical to the standards of open source – half open is fully closed, whether you're talking about an operating system or the edge.

Red Hat's vision for the edge: Innovative, hybrid and open

Red Hat helped drive the early standards around enterprise Linux, which has established the platform as the datacenter workhorse for innovation and production workloads. We've helped fuel the adoption of KVM as the open virtualization platform of choice for business computing, and we've been instrumental in the build-out of Kubernetes into the juggernaut of container orchestration that it is.

Now, as we and the collective business world turn our eyes to the network's edge, we see another wave of raw innovation. The promise and power of edge computing is very real and Red Hat, along with the open source communities and our ecosystem of partners, intends to help the business world capitalize on these benefits without fear of fragmentation or lock in.

The edge is open. The edge is hybrid. And the edge is powered by Red Hat.