

News About Servers (SUSE, Ubuntu, Red Hat and More)

By *Roy Schestowitz*

Created 03/12/2019 - 3:45pm

Submitted by Roy Schestowitz on Tuesday 3rd of December 2019 03:45:44 PM Filed under [Server](#) [1] [SUSE](#) [2]

- [What is Cloud Native? \[3\]](#)

Cloud native is more than just a buzzword, though. It's an approach used by some of the largest organizations on the planet, including Walmart, Visa, JP Morgan Chase, China Mobile, Verizon and Target, among others. Cloud native is an approach that enable developers and organization to be more agile, providing workload portability and scalability.

- [What is Kata Containers and why should I care? \[4\]](#)

Kata Containers can significantly improve the security and isolation of your container workloads. It combines the benefits of using a hypervisor, such as enhanced security, and container orchestration capabilities provided by Kubernetes.

Together with Eric Erns from Intel, we have recently performed a webinar in which we presented the benefits of using Kata Containers in a Charmed Kubernetes environment. In this blog, we aim to highlight the key outcomes from this webinar.

- [An idiot's guide to Kubernetes, low-code developers, and other industry trends \[5\]](#)

As part of my role as a senior product marketing manager at an enterprise software company

with an open source development model, I publish a regular update about open source community, market, and industry trends for product marketers, managers, and other influencers. Here are five of my and their favorite articles from that update.

- [**A blueprint for OpenStack and bare metal \[6\]**](#)

The bare metal cloud is an abstraction layer for the pools of dedicated servers with different capabilities (processing, networking or storage) that can be provisioned and consumed with cloud-like ease and speed. It embraces the orchestration and automation of the cloud and applies them to bare metal workload use cases.

The benefit to end users is that they get access to the direct hardware processing power of individual servers and are able to provision workloads without the overhead of the virtualization layer?providing the ability to provision environments in an Infrastructure-as-code methodology with separation of tenants and projects.

- [**Software Development, Microservices & Container Management ? Part III ? Why Kubernetes? A Deep Dive into Kubernetes world \[7\]**](#)

Together with my colleague Bettina Bassermann and SUSE partners, we will be running a series of blogs and webinars from SUSE (Software Development, Microservices & Container Management, a SUSE webinar series on modern Application Development), and try to address the former questions and doubts about K8s and Cloud Native development and how it is not compromising quality and control.

- [**Epic Performance with New Tuning Guide ? SUSE Linux Enterprise Server on AMD EPYC* 7002 Series Processors \[8\]**](#)

EPYC is AMD's flagship mainstream server microprocessors and supports 1-way and 2-way multiprocessing. The first generation was originally announced back in May 2017 and replaced the previous Opteron server family with the introduction of the Zen microarchitecture for the mainstream market.

- [**Content Lifecycle Management in SUSE Manager \[9\]**](#)

Content Lifecycle management is managing how patches flows through your infra in a staged manner. In ideal infra, latest patches will always be applied on development servers. If everything is good there then those patches will be applied to QA servers and lastly to

production servers. This enables sysadmins to catch issues if any and hence preventing patching of prod system which may create downtime of live environments.

SUSE Manager gives you this control via content lifecycle. In this, you create custom channels in SUSE Manager for example dev, qa and prod. Then you register your systems to those channels according to their criticality. Now whenever channels gets the new patches it will be available to respective systems (registered to those channels) to install. So if you control channels you control the patch availability to systems.

In content lifecycle management, suse manager enables you to push patches to channels manually. Like on first deploy all latest patches will be available to dev channels and hence dev systems. At this stage, if you run update commands (zypper up, yum update) they will show latest patches only on dev servers. QA and prod servers wont show any new patches.

- [The Early History of Usenet, Part VII: Usenet Growth and B-News \[10\]](#)

For quite a while, it looked like my prediction ? one to two articles per day ? was overly optimistic. By summer, there were only four new sites: Reed College, University of Oklahoma (at least, I think that that's what uucp node uok is), vax135, another Bell Labs machine ? and, crucially, U.C. Berkeley, which had a uucp connection to Bell Labs Research and was on the ARPANET.

In principle, even a slow rate of exponential growth can eventually take over the world. But that assumes that there are no "deaths" that will drive the growth rate negative. That isn't a reasonable assumption, though. If nothing else, Jim Ellis, Tom Truscott, Steve Daniel, and I all planned to graduate. (We all succeeded in that goal.) If Usenet hadn't shown its worth to our successors by then, they'd have let it wither. For that matter, university faculty or Bell Labs management could have pulled the plug, too. Usenet could easily have died aborning. But the right person at Berkeley did the right thing.

Mary Horton was then a PhD student there. (After she graduated, she joined Bell Labs; she and I were two of the primary people who brought TCP/IP to the Labs, where it was sometimes known as the "datagram heresy". The phone network was, of course, circuit-switched?) Known to her but unknown to us, there were two non-technical ARPANET mailing lists that would be of great interest to many potential Usenet users, HUMAN-NETS and SF-LOVERS. She set up a gateway that relayed these mailing lists into Usenet groups; these were at some point moved to the fa ("From ARPANET") hierarchy. (For a more detailed telling of this part of the story, see Ronda Hauben's writings.) With an actual traffic source, it was easy to sell folks on the benefits of Usenet. People would have preferred a real ARPANET connection but that was rarely feasible and never something that a student could set up: ARPANET connections were restricted to places that had research contracts with DARPA. The gateway at Berkeley was, eventually, bidirectional for both Usenet and email; this enabled Usenet-style communication between the networks.

Source URL: <http://www.tuxmachines.org/node/131245>

Links:

- [1] <http://www.tuxmachines.org/taxonomy/term/147>
- [2] <http://www.tuxmachines.org/taxonomy/term/117>
- [3] <https://www.datamation.com/cloud-computing/what-is-cloud-native.html>
- [4] <https://ubuntu.com/blog/what-is-kata-containers>
- [5] <https://opensource.com/article/19/12/technology-advice-and-other-industry-trends>
- [6] <https://www.redhat.com/en/blog/blueprint-openstack-and-bare-metal>
- [7] <https://www.suse.com/c/why-kubernetes-a-deep-dive-into-kubernetes-world/>
- [8] <https://www.suse.com/c/epic-performance-with-new-tuning-guide-suse-linux-enterprise-server-on-amd-epyc-7002-series-processors/>
- [9] <https://kerneltalks.com/tools/content-lifecycle-management-in-suse-manager/>
- [10] http://www.circleid.com/posts/20191202_the_early_history_of_usenet_part_vii_usenet_growth_and_b_news/