

Continuous Integration for IaaS

- Open vSwitch/DPDK, QEMU/libvirt, OpenStack, and Applications.

2017/05/10

Yusuke TATSUMI

(ytatsumi@yahoo-corp.jp)

\$ whoami

- Yusuke TATSUMI (立見 祐介)
 - @Yahoo! JAPAN
 - Network engineer for several years
 - building/operating/monitoring production network
 - OpenStacker from last year
 - with OvS/DPDK in our OpenStack

Agenda

- Yahoo! JAPAN and OvS/DPDK/OpenStack
- Start to use OvS/DPDK as vSwitch in compute nodes
- Recent issue
 - Adopting latest components
 - hard work for integration test
- CI for IaaS
- Actual examples
- Summary

Yahoo! JAPAN Portal Site Overview

The screenshot displays the Yahoo! JAPAN portal homepage. At the top, there is a navigation bar with the Yahoo! JAPAN logo and various service icons like 'My Yahoo!', 'ボックス', and 'メール'. Below this is a search bar with a '検索' button. The main content area is divided into several sections:

- お気に入り** (Favorites): A vertical menu on the left with icons for 地図 (Map), 路線 (Routes), 主なサービス (Main Services), ショッピング (Shopping), ヤフオク! (Yahoo! Auctions), LOHACO, 旅行、ホテル予約 (Travel, Hotel Reservations), ニュース (News), 天気 (Weather), スポーツ (Sports), ファイナンス (Finance), テレビ (TV), GYAO!, ゲーム (Games), Y!モバイル (Y! Mobile), 地図 (Map), 路線 (Routes), 食ベログ (Food Blogs), 求人、アルバイト (Job Openings, Part-time Jobs), 不動産 (Real Estate), 自動車 (Cars), 掲示板 (Bulletin Board), ブログ (Blogs), 美容、ダイエット (Beauty, Diet), 出会い、お見合い (Meetings, Matchmaking), and クックパッド (Cookpad).
- ニュース** (News): A central section with a list of news items under the 'ニュース' tab, including headlines like '60日ルール考えず 自公幹事長' and '薬地市場の建物から出火'.
- 天気** (Weather): A section on the right showing the current date (2015年6月24日) and weather forecasts for today and tomorrow in Tokyo, including temperature ranges and heat index.
- ランニング** (Running): A section with tips for running, such as '暑い日はちょっと食欲が減退気味' and '本格タイ料理を自宅で'.
- あなたにおすすめ** (Recommended for you): A section with product recommendations like 'アウトレット' and 'ロハコセール'.
- ログイン** (Login): A section at the bottom right for user login and account management.

Yahoo! JAPAN Portal Site Overview

Monthly Page Views

69+ Billion

Apps Total Download

380+ Million

Monthly Page Views
(Only Smartphone)

39+ Billion

Number of services

100+

OpenStack/OvS/DPDK supports our many services!

Start to use OvS/DPDK as vSwitich in Compute nodes

- High-performance/functioning needs for our web services.
- <https://youtu.be/NSeWzdulnTg>
 - Effective for networking
 - stable L7 performance even under noisy condition

Recent issue


- **To ensure our NW infrastructure stability**
 - Adopting latest version is good way
 - new function implements
 - ex. OvS client-mode
 - performance improvement/bug fixed
- **But func/perf integration test is hard work**
 - combination problems
 - many different release date

many components and many release

Every month integration test?

Every combinations can't be tested! ($3*4*4*6$)

=> Continuous Integration Tool-chain



CI for IaaS

- Package preparation by Jenkins
- Jenkins start tests
 - deployment
 - BareMetal
 - OpenStack, applications
 - function/performance testing
 - return the results

- iperf
- ab
- sysbench
- etc...

- Env 1.
 - liberty
 - OvS 2.5.90/DPDK 16.07
 - QEMU 2.6.0
- Env 2.
 - mitaka
 - OvS 2.6.1/DPDK 16.11
 - QEMU 2.7.1
- Env 3
 -

Our confirmation points

- Packet tx-drop in a permissible range
- Ovs-vsitchd return in a few seconds by client-mode
- PPS performance is (or not) improved with different combination
- etc...

Actual example (1)

- QEMU vhostforce-option mismatch by differences of QEMU version
 - leading performance degradation
 - <https://review.openstack.org/#/c/404606/>
 - merged :)

```
10 ■■■■ devstack/libs/ovs-dpdk
@@ -112,6 +112,16 @@ while [ $# -gt 0 ]; do
112 112          shift
113 113          fi
114 114          ;;
115 115  +   -netdev)
116 116  +   args[i]="$1"
117 117  +   (( i++ ))
118 118  +   shift
119 119  +   if [[ "$1" =~ "vhost-user" ]]; then
120 120  +       args[i]=${1},${VHOST_FORCE}
121 121  +       (( i++))
122 122  +       shift
123 123  +   fi
124 124  +   ;;
115 125  *)
116 126  args[i]="$1"
117 127  (( i++ ))
```

Actual example (2)

- Restarting “nova-compute” process recreate I/F of Ovs
 - an instant disconnection happen
- fixed internally
 - few reported
 - preparing a report

```
2 ■■■ nova/network/linux_net.py
@@ -1368,7 +1368,7 @@ def _ovs_vsctl(args):
1368 | 1368 |
1369 | 1369 |     def _create_ovs_vif_cmd(bridge, dev, iface_id, mac,
1370 | 1370 |                             instance_id, interface_type=None):
1371 | 1371 | -     cmd = ['--', '--if-exists', 'del-port', dev, '--',
1372 | 1372 | +     cmd = ['--may-exist',
1373 | 1373 |             'add-port', bridge, dev,
1374 | 1374 |             '--', 'set', 'Interface', dev,
1375 | 1375 |             'external-ids:iface-id=%s' % iface_id,
```

Summary

- Continue to adopt latest release is good way to improve stability
- It is hard to ensure func/perf with integrated infrastructure
- Whole integration test with CI-tools
 - Finding combination problems
 - It makes NW infrastructure much stable

EOP

Thank you for your kind attention.