

## openQA Infrastructure - action #105594

coordination # 80142 (Blocked): [saga][epic] Scale out: Redundant/load-balancing deployments of openQA, easy containers, containers on kubernetes

### Two new machines for OSD and o3, meant for bare-metal virtualization size:M

2022-01-27 08:13 - okurz

<b>Status:</b>	Workable	<b>Start date:</b>	2022-06-16
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	nicksinger	<b>% Done:</b>	100%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	Ready		
<b>Description</b>			
<b>Current situation</b>			
<p>Two new x86_64 machines have been ordered by SUSE QE, meant for QE virtualization purposes. One within o3, the other one for OSD. Previously o3 hardware needed to be in SRV1 where there is no more space. We might be able to move osd hardware physically to a different location, e.g. SRV2, QA labs, etc. to use the place for o3 hardware. There is no problem to connect from any location over HTTPS to o3 like is already done for external ARM cloud workers. The challenge is how to prevent access to the rest of SUSE network from this machine. <a href="https://openqa.opensuse.org">uno.openqa.opensuse.org</a> might be a candidate to remove and make place for new machines. However I see the ROI very little to conduct such changes in the productive environment within SRV1. We should discuss with EngInfra what options we have to setup maybe a new dedicated network that has access to openqa.opensuse.org:443, i.e. public internet, but no access to the rest of the internal SUSE network.</p>			
<b>Subtasks:</b>			
action # 112553: [osd][amd][zen3][network][sriov] New AMD Zen3 machine on OSD lost its ...			<b>Resolved</b>
<b>Related issues:</b>			
Related to openQA Project - action #110227: Stop showing ipmi passwords in au...		<b>New</b>	<b>2022-04-24</b>

### History

#### #1 - 2022-01-27 08:15 - okurz

- Status changed from New to Feedback
- Assignee set to okurz

[mgriessmeier](#) I added the ticket to our backlog as discussed and assign it to myself to wait for your feedback after initial clarification.

#### #2 - 2022-01-27 08:44 - nicksinger

- Description updated

I've added OSD-Admins to the corresponding jira ticket (<https://sd.suse.com/servicedesk/customer/portal/1/SD-74616>). I think we should discuss our possible solutions before approaching infra.

#### #3 - 2022-01-27 13:12 - okurz

- Status changed from Feedback to Blocked

<https://sd.suse.com/servicedesk/customer/portal/1/SD-74616> reads like everything was clarified already and EngInfra plans to do it next week so I suggest we just wait for feedback in the SD ticket if there are problems.

#### #4 - 2022-01-27 15:34 - nicksinger

- Subject changed from Two new aarch64 machines for o3, meant for bare-metal virtualization to Two new machines for OSD and o3, meant for bare-metal virtualization
- Description updated

#### #5 - 2022-01-27 15:35 - nicksinger

We already have an o3 worker in SRV2. Therefore I added that remark in Jira SD.

#### #6 - 2022-02-25 11:19 - okurz

@Julie\_CAO I would like to help with getting you setup on remote administration side regarding o3. What do you mean with "I have readonly permission with my account jcao@ariel in O3" that you asked in <https://sd.suse.com/servicedesk/customer/portal/1/SD-74616> ?

#### #7 - 2022-02-27 01:50 - Julie\_CAO

okurz wrote:

@Julie\_CAO I would like to help with getting you setup on remote administration side regarding o3. What do you mean with "I have readonly permission with my account jcao@ariel in O3" that you asked in <https://sd.suse.com/servicedesk/customer/portal/1/SD-74616> ?

I found only the machine IP was added to ariel, while the IP address of IPMI was not. I failed to add it to /etc/hosts and /etc/dnsmasq.d/openqa.conf because they are readonly for jcao@ariel.

```
dhcp-host=ec:2a:72:0c:23:c0,amd-zen2-gpu-sut1-ipmi //the MAC of the IPMI
dhcp-host=ec:2a:72:02:83:c4,amd-zen2-gpu-sut1
```

#### #8 - 2022-02-27 06:52 - okurz

Right. This is just a standard Linux system so normal users don't have write permissions. If I have created the user account for you then I have added you to the corresponding group that has sudo permissions. Please be aware that o3 is a critical production machine. Stay responsive in the #opensuse-factory libera.chat IRC room when you do chances on the machine

#### #9 - 2022-02-28 05:31 - Julie\_CAO

Thank you, Oliver. I did not notice that I can sudo. I just added the ipmi of the new machine zen2, amd-zen2-gpu-sut1-ipmi, to those two files, but I did not touch the dnsmasq.service in case of breaking anything.

#### #10 - 2022-03-02 01:59 - waynechen55

Currently there are two issues found on machine amd-zen3-gpu-sut1-1:

- PXE boot is not configured for the machine. It can not find PXE boot file. The normal PXE boot for machines in qa network in openqa.suse.de should look like this: [https://openqa.suse.de/tests/8208090#step/boot\\_from\\_pxe/6](https://openqa.suse.de/tests/8208090#step/boot_from_pxe/6) Pressing 'esc' will bring up the 'boot:' prompt.
- I installed SLES 15-SP4 PublicBeta Build101.1 onto this machine by using ISO image successfully. The machine is up and running now. But its one ip address 10.162.32.106 is not the one configured here: <https://gitlab.suse.de/qa-sle/qanet-configs/-/commit/4f744851316fd981a520de5d39f24e43413ec96e>

```
1: lo: mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
inet 127.0.0.1/8 scope host lo
valid_lft forever preferred_lft forever
inet6 ::1/128 scope host
valid_lft forever preferred_lft forever
2: p3p1: mtu 1500 qdisc mq master br0 state UP group default qlen 1000
link/ether b4:96:91:9c:5a:d4 brd ff:ff:ff:ff:ff:ff
altname enp65s0f0
3: em1: mtu 1500 qdisc mq master br1 state UP group default qlen 1000
link/ether ec:2a:72:02:84:20 brd ff:ff:ff:ff:ff:ff
altname eno8303
altname enp225s0f0
4: em2: mtu 1500 qdisc noop state DOWN group default qlen 1000
link/ether ec:2a:72:02:84:21 brd ff:ff:ff:ff:ff:ff
altname eno8403
altname enp225s0f1
5: p3p2: mtu 1500 qdisc noop state DOWN group default qlen 1000
link/ether b4:96:91:9c:5a:d5 brd ff:ff:ff:ff:ff:ff
altname enp65s0f1
6: br0: mtu 1500 qdisc noqueue state UP group default qlen 1000
link/ether b4:96:91:9c:5a:d4 brd ff:ff:ff:ff:ff:ff
inet 10.162.32.106/18 brd 10.162.63.255 scope global br0
valid_lft forever preferred_lft forever
inet6 2620:113:80c0:80a0:10:162:29:e843/64 scope global dynamic noprefixroute
valid_lft 2517529sec preferred_lft 1545529sec
inet6 fe80::b696:91ff:fe9c:5ad4/64 scope link
valid_lft forever preferred_lft forever
7: br1: mtu 1500 qdisc noqueue state UP group default qlen 1000
link/ether ec:2a:72:02:84:20 brd ff:ff:ff:ff:ff:ff
inet 10.162.2.132/18 brd 10.162.63.255 scope global br1
valid_lft forever preferred_lft forever
inet6 2620:113:80c0:80a0:10:162:29:5183/64 scope global dynamic noprefixroute
valid_lft 2517535sec preferred_lft 1545535sec
inet6 fe80::ee2a:72ff:fe02:8420/64 scope link
valid_lft forever preferred_lft forever
```

```
amd-zen3-gpu-sut1-1:~ # cat /etc/sysconfig/network/ifcfg-br0
BOOTPROTO='dhcp'
STARTMODE='auto'
BRIDGE='yes'
BRIDGE_PORTS='p3p1'
BRIDGE_STP='off'
BRIDGE_FORWARDDELAY='15'
ZONE=public
```

```
amd-zen3-gpu-sut1-1:~ # cat /etc/sysconfig/network/ifcfg-br1
BOOTPROTO='dhcp'
STARTMODE='auto'
BRIDGE='yes'
BRIDGE_PORTS='em1'
BRIDGE_STP='off'
BRIDGE_FORWARDDELAY='15'
ZONE=public
```

- The configured mac address and ip address are
  - hardware ethernet ec:2a:72:02:84:20; fixed-address 10.162.2.132
  - hardware ethernet ec:2a:72:02:84:21; fixed-address 10.162.2.133
- But on the machine:
  - br0 is bound with p3p1 whose mac address is b4:96:91:9c:5a:d4. So br0 ip address is 10.162.32.106/18 instead of 10.162.2.133.
  - br1 is bound with em1 whose mac address is ec:2a:72:02:84:20. So br1 ip address is 10.162.2.132.
- Ethernet interface em2 is down and has no cable plugged in, but p3p1 is UP. So it seems that the cable is connected to the wrong interface:

```
2: p3p1: mtu 1500 qdisc mq master br0 state UP group default qlen 1000
link/ether b4:96:91:9c:5a:d4 brd ff:ff:ff:ff:ff:ff
altname enp65s0f0
```

```
4: em2: mtu 1500 qdisc noop state DOWN group default qlen 1000
link/ether ec:2a:72:02:84:21 brd ff:ff:ff:ff:ff:ff
altname eno8403
altname enp225s0f1
```

Please help fix the issue. Thanks.

#### #11 - 2022-03-09 03:37 - Julie\_CAO

Hi [okurz](#) and [nicksinger](#), could you help make this change in <https://gitlab.suse.de/qa-sle/qanet-configs/-/commit/4f744851316fd981a520de5d39f24e43413ec96e?>

change the MAC in

```
host amd-zen3-gpu-sut1-2 hardware ethernet ec:2a:72:02:84:21; fixed-address 10.162.2.133; option host-name "amd-zen3-gpu-sut1-2";
```

to

```
B4:96:91:9C:5A:D4
```

#### #12 - 2022-03-09 09:09 - okurz

Could you prepare a MR yourself? Simply checkout <https://gitlab.suse.de/qa-sle/qanet-configs/> and prepare a merge request and then we can apply the change on the machine

#### #13 - 2022-03-10 02:21 - Julie\_CAO

The MR is submitted, could you please help review and merge? [https://gitlab.suse.de/qa-sle/qanet-configs/-/merge\\_requests/37](https://gitlab.suse.de/qa-sle/qanet-configs/-/merge_requests/37)

#### #14 - 2022-03-11 09:54 - nicksinger

Merged and deployed since yesterday :)

#### #15 - 2022-03-15 05:31 - Julie\_CAO

Hi [nicksinger](#) and [okurz](#), about the zen2 machine in O3. I added the MAC of the IPMI (according to your commit in the infra ticket, <https://gitlab.suse.de/qa-sle/qanet-configs/-/commit/8a9f96d0b5d5dd5ea5a630de873b1b8f3b255317>) to

```
/etc/dnsmasq.d/openqa.conf
dhcp-host=ec:2a:72:0c:23:c0,amd-zen2-gpu-sut1-ipmi
```

```
/etc/hosts:
192.168.112.16 amd-zen2-gpu-sut1-ipmi.openqanet.opensuse.org amd-zen2-gpu-sut1-ipmi
```

and restart dnsmasq.service, but its ip is not ping'able

```
jcao@ariel:> ping 192.168.112.16
PING 192.168.112.16 (192.168.112.16) 56(84) bytes of data.
^C
--- 192.168.112.16 ping statistics ---
3 packets transmitted, 0 received, 100% packet loss, time 2043ms
```

Could you help check if 3 network cables are connected to zen2 correctly? Are the MAC addresses of them correct?

```
1st cable is expected to be connected to the IPMI: ec:2a:72:0c:23:c0
2nd is expected to be connected to the onboard network card port one: ec:2a:72:02:83:c4
3rd is expected to be connected to the extra network card port one: don't know the MAC
```

I am unable to access the ipmi or iDRAC of the machine, so I can do nothing to it now.

#### #16 - 2022-03-16 05:12 - waynechen55

For the zen3 machine on OSD, I found that it sets itself as 'amd-zen3-gpu-sut1-2' instead of our preferred 'amd-zen3-gpu-sut1-1'. Do you know how to let the host always set itself as 'amd-zen3-gpu-sut1-1' ? It should survive reboot, fresh installation and upgrade. Thanks. [nicksinger](#)

```
waynechen-opensuse:~ # ping -c5 amd-zen3-gpu-sut1-2.qa.suse.de
PING amd-zen3-gpu-sut1-2.qa.suse.de (10.162.2.133) 56(84) bytes of data.
64 bytes from amd-zen3-gpu-sut1-2.qa.suse.de (10.162.2.133): icmp_seq=1 ttl=59 time=196 ms
64 bytes from amd-zen3-gpu-sut1-2.qa.suse.de (10.162.2.133): icmp_seq=2 ttl=59 time=192 ms
64 bytes from amd-zen3-gpu-sut1-2.qa.suse.de (10.162.2.133): icmp_seq=3 ttl=59 time=195 ms
64 bytes from amd-zen3-gpu-sut1-2.qa.suse.de (10.162.2.133): icmp_seq=4 ttl=59 time=197 ms
64 bytes from amd-zen3-gpu-sut1-2.qa.suse.de (10.162.2.133): icmp_seq=5 ttl=59 time=194 ms

--- amd-zen3-gpu-sut1-2.qa.suse.de ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4005ms
rtt min/avg/max/mdev = 192.034/194.845/197.007/1.643 ms
waynechen-opensuse:~ # ping -c5 amd-zen3-gpu-sut1-1.qa.suse.de
PING amd-zen3-gpu-sut1-1.qa.suse.de (10.162.2.132) 56(84) bytes of data.
64 bytes from amd-zen3-gpu-sut1-1.qa.suse.de (10.162.2.132): icmp_seq=1 ttl=59 time=194 ms
64 bytes from amd-zen3-gpu-sut1-1.qa.suse.de (10.162.2.132): icmp_seq=2 ttl=59 time=193 ms
64 bytes from amd-zen3-gpu-sut1-1.qa.suse.de (10.162.2.132): icmp_seq=3 ttl=59 time=193 ms
64 bytes from amd-zen3-gpu-sut1-1.qa.suse.de (10.162.2.132): icmp_seq=4 ttl=59 time=192 ms
64 bytes from amd-zen3-gpu-sut1-1.qa.suse.de (10.162.2.132): icmp_seq=5 ttl=59 time=192 ms

--- amd-zen3-gpu-sut1-1.qa.suse.de ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4006ms
rtt min/avg/max/mdev = 191.627/192.723/194.238/0.848 ms
waynechen-opensuse:~ #
waynechen-opensuse:~ # ssh amd-zen3-gpu-sut1-1.qa.suse.de
(root@amd-zen3-gpu-sut1-1.qa.suse.de) Password:
Last login: Wed Mar 16 17:01:34 2022 from 10.67.19.99
amd-zen3-gpu-sut1-2:~ # hostnamectl
  Static hostname: n/a
Transient hostname: amd-zen3-gpu-sut1-2
      Icon name: computer-server
    Chassis: server
  Machine ID: d13c5ald24c047ala4f5c1f56392edfd
    Boot ID: b11920d9d4cc40b5b4ff8ef94ab2cf3f
  Operating System: SUSE Linux Enterprise Server 15 SP4
    CPE OS Name: cpe:/o:suse:sles:15:sp4
      Kernel: Linux 5.14.21-150400.11-default
  Architecture: x86_64
  Hardware Vendor: Dell Inc.
  Hardware Model: PowerEdge R7525
```

#### #17 - 2022-03-18 19:25 - okurz

- Status changed from Blocked to New  
- Assignee deleted (okurz)

After <https://sd.suse.com/servicedesk/customer/portal/1/SD-74616> was resolved this is to be continued in our backlog.

#18 - 2022-03-23 11:43 - okurz

- Due date set to 2022-04-06
- Status changed from New to Feedback
- Assignee set to okurz

Just collecting some information to be on the same page.

- The machine "zen2" shows up in racktables as [https://racktables.nue.suse.com/index.php?page=object&tab=edit&object\\_id=16386](https://racktables.nue.suse.com/index.php?page=object&tab=edit&object_id=16386)
- <https://gitlab.suse.de/qa-sle/qanet-configs/-/blob/master/etc/dhcpd.conf#L317> shows the DHCP config including the "host-name" options (which does not need to be the same as how the system refers to itself as hostname and also it does not need to be the same as DNS). DHCP seems to be properly configured. We might want to change the host-name entries or completely remove them to prevent confusion.
- <https://gitlab.suse.de/qa-sle/qanet-configs/-/blob/master/var/lib/named/master/forward/qa.suse.de#L372> shows the DNS config. That looks to be fine, in line with DHCP. Be aware that these entries are interface specific, so it's two interfaces for one host and one DNS entry for each interface. For the service processor we know for sure that it's connected to "VLAN 12", i.e. QA-network.

Julie\_CAO wrote:

Could you help check if 3 network cables are connected to zen2 correctly?

According to nsinger's notes:

- amd-zen2-gpu-sut1-sp is connected to qanet10nue:gi9. on the switch show interfaces switchport gi9 confirms it's connected to VLAN 12 (QA), link is up (show interfaces status)
- one physical interface is connected to qanet10nue:gi10. on the switch show interfaces switchport gi10 confirms it's connected to VLAN 662 (o3), link is up (show interfaces status)
- another physical interface is connected to qanet10nue:gi11. on the switch show interfaces switchport gi11 confirms it's connected to VLAN 662 (o3), link is up (show interfaces status)

With access to the BMC - we don't know username and password, likely you changed it? - we could likely crosscheck which of the physical interfaces is connected to which port.

Are the MAC addresses of them correct?

1st cable is expected to be connected to the IPMI: ec\*:c0

yes. on the switch show mac address-table interface gi9 confirms

2nd is expected to be connected to the onboard network card port one: ec\*:c4

Right now show mac address-table interface gi10 indeed shows that mac address, so yes, correct.

3rd is expected to be connected to the extra network card port one: don't know the MAC

<https://gitlab.suse.de/qa-sle/qanet-configs/-/blob/master/etc/dhcpd.conf#L319> says it should be ec::c5 (most likely looked up from nsinger from BMC) but qanet10nue says 'b4::82' which according to is an intel card

I am unable to access the ipmi or iDRAC of the machine, so I can do nothing to it now.

nsinger had access to the BMC but does not have access anymore so someone likely changed the password. I assume it was one of you or your team. Please find out the password and crosscheck the above config (or share the password to us OVER A SECURE CHANNEL, not in the ticket)

waynechen55 wrote:

For the zen3 machine on OSD, I found that it sets itself as 'amd-zen3-gpu-sut1-2' instead of our preferred 'amd-zen3-gpu-sut1-1'. Do you know how to let the host always set itself as 'amd-zen3-gpu-sut1-1' ?

To have a static consistent hostname just set it using hostnamectl, see the man page of hostnamectl (or <https://linuxhint.com/set-hostname-using-hostnamectl-command/>)

For the sake of completeness I checked the interfaces on zen3 from the currently running installation. I could login using ssh\_nt root@amd-zen3-gpu-sut1-1.qa.suse.de and call ip link and found:

- p3p1 (same as br0), UP: b4\*:d4
- em1 (same as br1), UP: ec\*:20
- em2 DOWN: ec\*:21
- p3p2 DOWN: b4\*:d5

meaning that the information for zen3 in <https://gitlab.suse.de/qa-sle/qanet-configs/-/blob/master/etc/dhcpd.conf#L322> is correct but I don't know if zen2 is correct. Maybe there the "second" network card is also an Intel one so the mac address would not be what is written in <https://gitlab.suse.de/qa-sle/qanet-configs/-/blob/master/etc/dhcpd.conf#L319>

With the help of nsinger I updated the racktable entries for both machines so now we have up to date link, port and ip information in racktables as well.

#### #19 - 2022-03-24 03:01 - waynechen55

Additionally, I do not think work on zen3 is done. PXE boot is not configured for zen3 machine in OSD network. Would you please help arrange and get the work done ? My previous experience told me PXE boot in OSD network just looks like this: [https://openqa.suse.de/tests/8350667#step/boot\\_from\\_pxe/6](https://openqa.suse.de/tests/8350667#step/boot_from_pxe/6)

#### #20 - 2022-03-24 06:27 - Julie\_CAO

- Status changed from Feedback to In Progress

Thank you, @okurz.

about zen2 in O3:

- amd-zen2-gpu-sut1-sp is connected to qanet10nue:gi9. on the switch show interfaces switchport gi9 confirms it's connected to VLAN 12 (QA), link is up (show interfaces status)

Yes, I just checked the ipmi of zen2 is connected to VLAN12(QA). We need it in the VLAN 662 (o3) as our test need ipmitool to this machine. I'll open an infra ticket to handle this issue.

3rd is expected to be connected to the extra network card port one: don't know the MAC

<https://gitlab.suse.de/qa-sle/qanet-configs/-/blob/master/etc/dhcpd.conf#L319> says it should be `ec::c5` (most likely looked up from nsinger from BMC) but `qanet10nue` says `b4::82` which according to is an intel card

`ec::c5` listed in `dhcpd.conf` is not correct as it was the other port of the same onboard network card. `b4::82` sounds be more reseanable but I am not sure yet.

nsinger had access to the BMC but does not have access anymore so someone likely changed the password. I assume it was one of you or your team. Please find out the password and crosscheck the above config (or share the password to us OVER A SECURE CHANNEL, not in the ticket)

My team and I did not change the ipmi password because we have not access the iDRAC successfully yet. I just tried `amd-zen2-gpu-sut1-sp.qa.suse.de`, the default root password did not work for me as well. I have to ask infra for help in that ticket.

#### #21 - 2022-03-24 11:36 - okurz

waynechen55 wrote:

Additionally, I do not think work on zen3 is done. PXE boot is not configured for zen3 machine in OSD network. Would you please help arrange and get the work done ? My previous experience told me PXE boot in OSD network just looks like this: [https://openqa.suse.de/tests/8350667#step/boot\\_from\\_pxe/6](https://openqa.suse.de/tests/8350667#step/boot_from_pxe/6)

PXE config can be configured as part of DHCP config in <https://gitlab.suse.de/qa-sle/qanet-configs/-/blob/master/etc/dhcpd.conf> so you can provide merge requests there based on what is needed. I will crosscheck with nsinger if everything is accessible to you to be able to solve this.

Julie\_CAO wrote:

about zen2 in O3:

- amd-zen2-gpu-sut1-sp is connected to qanet10nue:gi9. on the switch show interfaces switchport gi9 confirms it's connected to VLAN 12 (QA), link is up (show interfaces status)

Yes, I just checked the ipmi of zen2 is connected to VLAN12(QA). We need it in the VLAN 662 (o3) as our test need ipmitool to this machine. I'll open an infra ticket to handle this issue.

An infra ticket does not help. The QA switches are managed by us. We will do that.

3rd is expected to be connected to the extra network card port one: don't know the MAC

<https://gitlab.suse.de/qa-sle/qanet-configs/-/blob/master/etc/dhcpd.conf#L319> says it should be ec::c5 (most likely looked up from nsinger from BMC) but qanet10nue says `b4::82` which according to is an intel card

ec::c5 listed in dhcpd.conf is not correct as it was the other port of the same onboard network card. b4::82 sounds be more reseanable but I am not sure yet.

nsinger had access to the BMC but does not have access anymore so someone likely changed the password. I assume it was one of you or your team. Please find out the password and crosscheck the above config (or share the password to us OVER A SECURE CHANNEL, not in the ticket)

My team and I did not change the ipmi password because we have not access the iDRAC successfully yet. I just tried amd-zen2-gpu-sut1-sp.qa.suse.de, the default root password did not work for me as well. I have to ask infra for help in that ticket.

ok, do that and please involve us or come back to us with what you learned.

## #22 - 2022-03-24 11:48 - Julie\_CAO

okurz wrote:

about zen2 in O3:

Yes, I just checked the ipmi of zen2 is connected to VLAN12(QA). We need it in the VLAN 662 (o3) as our test need ipmitool to this machine. I'll open an infra ticket to handle this issue.

An infra ticket does not help. The QA switches are managed by us. We will do that.

Thank you. I just canceled my infra request.

My team and I did not change the ipmi password because we have not access the iDRAC successfully yet. I just tried amd-zen2-gpu-sut1-sp.qa.suse.de, the default root password did not work for me as well. I have to ask infra for help in that ticket.

ok, do that and please involve us or come back to us with what you learned.

Infra helped to found out the credentials for me. As here is the public space I'll not paste it here. you can find the user/password in SD-81238. Or email, rocketchat?

## #23 - 2022-03-24 12:38 - okurz

okurz wrote:

waynechen55 wrote:

Additionally, I do not think work on zen3 is done. PXE boot is not configured for zen3 machine in OSD network. Would you please help arrange and get the work done ? My previous experience told me PXE boot in OSD network just looks like this:

[https://openqa.suse.de/tests/8350667#step/boot\\_from\\_pxe/6](https://openqa.suse.de/tests/8350667#step/boot_from_pxe/6)

PXE config can be configured as part of DHCP config in <https://gitlab.suse.de/qa-sle/qanet-configs/-/blob/master/etc/dhcpd.conf> so you can provide merge requests there based on what is needed. I will crosscheck with nsinger if everything is accessible to you to be able to solve this.

Ok, so you just need to add "pxelinux.0" for the dhcp entry, like e.g. done in <https://gitlab.suse.de/qa-sle/qanet-configs/-/blob/master/etc/dhcpd.conf#L315>

about zen2 in O3:

- amd-zen2-gpu-sut1-sp is connected to qanet10nue:gi9. on the switch show interfaces switchport gi9 confirms it's connected to VLAN 12 (QA), link is up (show interfaces status)

Yes, I just checked the ipmi of zen2 is connected to VLAN12(QA). We need it in the VLAN 662 (o3) as our test need ipmitool to this machine. I'll open an infra ticket to handle this issue.

An infra ticket does not help. The QA switches are managed by us. We will do that.

We have to reconsider. What nsinger and gschlotter have brought up as well: With making ipmi accessible in the o3 network we basically just have ariel as only line of defence against the public internet. Given how much you can do over IPMI (control the whole machine, install firmware and such)

this is really dangerous and we should consider if we really want such scenarios. I think we should avoid that. So sorry, I can currently not do that. Maybe you have good ideas what we can do as a more secure solution.

**#26 - 2022-03-29 07:51 - xlai**

okurz wrote:

about zen2 in O3:

- amd-zen2-gpu-sut1-sp is connected to qanet10nue:gi9. on the switch show interfaces switchport gi9 confirms it's connected to VLAN 12 (QA), link is up (show interfaces status)

Yes, I just checked the ipmi of zen2 is connected to VLAN12(QA). We need it in the VLAN 662 (o3) as our test need ipmitool to this machine. I'll open an infra ticket to handle this issue.

An infra ticket does not help. The QA switches are managed by us. We will do that.

We have to reconsider. What nsinger and gschlotter have brought up as well: With making ipmi accessible in the o3 network we basically just have ariel as only line of defence against the public internet. Given how much you can do over IPMI (control the whole machine, install firmware and such) this is really dangerous and we should consider if we really want such scenarios. I think we should avoid that. So sorry, I can currently not do that. Maybe you have good ideas what we can do as a more secure solution.

[okurzmgiessmeier](#) Thanks for your consistent support on this ticket. We fully agree that security is very very important. There should be solution for this before the zen2 ipmi machine is added in O3 network.

This zen2 machine is planned to support the tumbleweed virtualization testing in O3. If there is no way to add it, we will have to reject the "factory first policy" for virtualization testing. This is serious. We need to be cautious.

We are not expert in security and infra. Would you please give us some suggestions? Is there no solution at all? Or who else do you think we should involve to seek for potential solutions?

**#27 - 2022-03-30 10:17 - okurz**

- Due date deleted (2022-04-06)

- Status changed from In Progress to Feedback

[gschlotternicksinger](#) can you comment on the above regarding IPMI access from openQA tests within o3?

**#28 - 2022-03-31 08:06 - waynechen55**

It seems that ipmi sol connection to amd-zen3-gpu-sut1-sp.qa.suse.de is broken:

```
host:~ # ipmitool -H amd-zen3-gpu-sut1-sp.qa.suse.de -I lanplus -U xxxx -P xxxx chassis power status
```

```
Error: Unable to establish IPMI v2 / RMCP+ session
```

I tried many times with different ipmitool subcommands. Could anyone have a look ?

**#29 - 2022-03-31 09:10 - waynechen55**

waynechen55 wrote:

It seems that ipmi sol connection to amd-zen3-gpu-sut1-sp.qa.suse.de is broken:

```
host:~ # ipmitool -H amd-zen3-gpu-sut1-sp.qa.suse.de -I lanplus -U xxxx -P xxxx chassis power status
```

```
Error: Unable to establish IPMI v2 / RMCP+ session
```

I tried many times with different ipmitool subcommands. Could anyone have a look ?

It seems that BIOS settings changed somehow. I changed it back. Now ipmi sol is enabled and active.

**#30 - 2022-04-06 09:55 - waynechen55**

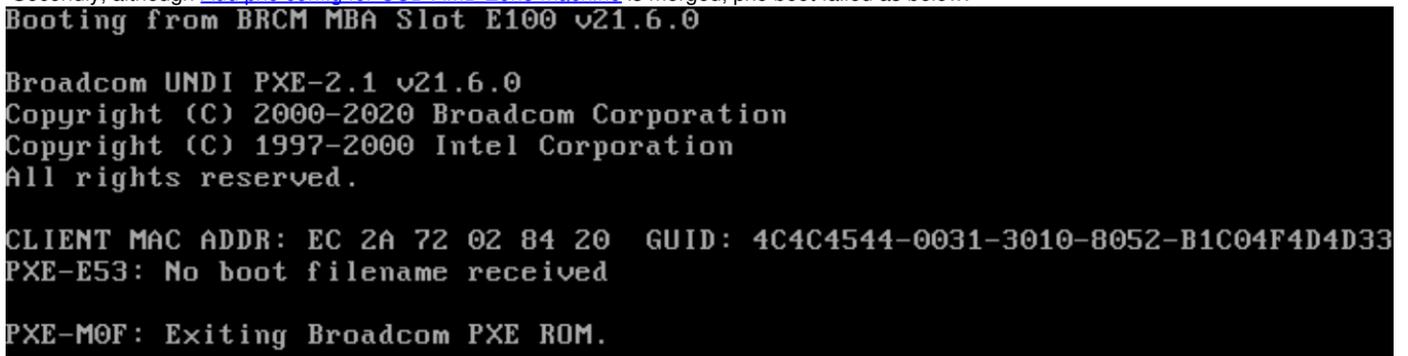
- File `pxe_boot_failure.png` added

I found two issues with amd-zen3-gpu-sut1-1.qa.suse.de:

- Firstly, it only has one ip address now. It seems that it secondary:

```
host amd-zen3-gpu-sut1-2 { hardware ethernet b4:96:91:9c:5a:d4; fixed-address 10.162.2.133; option host-name "amd-zen3-gpu-sut1-2"; filename "pxelinux.0"; }
amd-zen3-gpu-sut1-1:~ # ip addr show
1: lo: mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
inet 127.0.0.1/8 scope host lo
valid_lft forever preferred_lft forever
inet6 ::1/128 scope host
valid_lft forever preferred_lft forever
2: em1: mtu 1500 qdisc mq master br0 state UP group default qlen 1000
link/ether ec:2a:72:02:84:20 brd ff:ff:ff:ff:ff:ff
altname eno8303
altname enp225s0f0
3: p3p1: mtu 1500 qdisc noop state DOWN group default qlen 1000
link/ether b4:96:91:9c:5a:d4 brd ff:ff:ff:ff:ff:ff
altname enp65s0f0
4: em2: mtu 1500 qdisc noop state DOWN group default qlen 1000
link/ether ec:2a:72:02:84:21 brd ff:ff:ff:ff:ff:ff
altname eno8403
altname enp225s0f1
5: p3p2: mtu 1500 qdisc noop state DOWN group default qlen 1000
link/ether b4:96:91:9c:5a:d5 brd ff:ff:ff:ff:ff:ff
altname enp65s0f1
6: br0: mtu 1500 qdisc noqueue state UP group default qlen 1000
link/ether ec:2a:72:02:84:20 brd ff:ff:ff:ff:ff:ff
inet 10.162.2.132/18 brd 10.162.63.255 scope global br0
valid_lft forever preferred_lft forever
inet6 2620:113:80c0:80a0:10:162:29:37ac/64 scope global dynamic noprefixroute
valid_lft 2575753sec preferred_lft 1603753sec
inet6 fe80::ee2a:72ff:fe02:8420/64 scope link
valid_lft forever preferred_lft forever
```

- Secondly, although [Add pxe config for OSD AMD Zen3 machine](#) is merged, pxe boot failed as below:



[nicksinger](#) Any idea ?

**#31 - 2022-04-07 04:22 - cachem**

@wayne, after config p3p1 port to dhcp in system, amd-zen3-gpu-sut1-2 is up. network connection and dhcp setting works.

**#32 - 2022-04-07 05:12 - waynechen55**

- File `pxe_boot_failure_1.png` added

cachen wrote:

@wayne, after config p3p1 port to dhcp in system, amd-zen3-gpu-sut1-2 is up. network connection and dhcp setting works.

Now both ip addresses and fqdns work. But pxe boot does not work.  
I enabled pxe boot on both em1 and p3p1.

```
Booting from BRCM MBA Slot E100 v21.6.0

Broadcom UNDI PXE-2.1 v21.6.0
Copyright (C) 2000-2020 Broadcom Corporation
Copyright (C) 1997-2000 Intel Corporation
All rights reserved.

CLIENT MAC ADDR: EC 2A 72 02 84 20  GUID: 4C4C4544-0031-3010-8052-B1C04F4D4D33
PXE-E53: No boot filename received

PXE-M0F: Exiting Broadcom PXE ROM.

Booting from IBA 40G Slot 4100 v1130

Intel(R) Boot Agent 40G v1.1.30
Copyright (C) 1997-2020, Intel Corporation

CLIENT MAC ADDR: B4 96 91 9C 5A D4  GUID: 4C4C4544 0031 3010 8052 B1C04F4D4D33
PXE-E53: No boot filename received

PXE-M0F: Exiting Intel Boot Agent.
```

**#33 - 2022-04-07 05:16 - cachem**

@wayne, your 2nd pxe boot issue: I assume dhcp service still need to be restarted by manually to enable the PXE for this machine after your PR merged? please @nick help to check.

I just restarted dhcpd - can you please check again :)

**#34 - 2022-04-08 03:43 - cachem**

cachem wrote:

@wayne, your 2nd pxe boot issue: I assume dhcp service still need to be restarted by manually to enable the PXE for this machine after your PR merged? please @nick help to check.

I just restarted dhcpd - can you please check again :)

Thank you for the help, it is confirmed OSD PXE works for this Zen3 machine now, restart dhcpd manually was needed :)

Let's keep tracking this ticket for Zen2 machine to o3,

**#35 - 2022-04-09 12:40 - okurz**

- Parent task set to #109743

**#36 - 2022-04-14 15:14 - viktors.trubovics**

xlai wrote:

okurz wrote:

about zen2 in O3:

- amd-zen2-gpu-sut1-sp is connected to qanet10nue:gi9. on the switch show interfaces switchport gi9 confirms it's connected to VLAN 12 (QA), link is up (show interfaces status)

Yes, I just checked the ipmi of zen2 is connected to VLAN12(QA). We need it in the VLAN 662 (o3) as our test need ipmitool to this machine. I'll open an infra ticket to handle this issue.

An infra ticket does not help. The QA switches are managed by us. We will do that.

We have to reconsider. What nsinger and gschlotter have brought up as well: With making ipmi accessible in the o3 network we basically just have ariel as only line of defence against the public internet. Given how much you can do over IPMI (control the whole machine, install firmware and such) this is really dangerous and we should consider if we really want such scenarios. I think we should avoid that. So sorry, I can currently not do that. Maybe you have good ideas what we can do as a more secure solution.

[okurzmgruessmeier](#) Thanks for your consistent support on this ticket. We fully agree that security is very very important. There should be solution for this before the zen2 ipmi machine is added in O3 network.

This zen2 machine is planned to support the tumbleweed virtualization testing in O3. If there is no way to add it, we will have to reject the "factory first policy" for virtualization testing. This is serious. We need to be cautious.

We are not expert in security and infra. Would you please give us some suggestions? Is there no solution at all? Or who else do you think we should involve to seek for potential solutions?

The only way I see in this case, where IPMI must be exposed to internet - the server must be not able to connect to internal SUSE networks and 20 character strong unique password must be used for IPMI. In case the server will be hacked - SUSE network must stay secure.

**#37 - 2022-04-15 06:22 - Julie\_CAO**

The only way I see in this case, where IPMI must be exposed to internet - the server must be not able to connect to internal SUSE networks and 20 character strong unique password must be used for IPMI. In case the server will be hacked - SUSE network must stay secure.

Thanks, [viktors.trubovics](#)

Our test does NOT require to connect to SUSE internal network, because the install media and repositories for Tumbleweed are from [download.opensuse.org](http://download.opensuse.org) over http.

20 character strong unique password is ok for us. But is it acceptable if the IPMI password would be possibly exposed in openqa test log in the case of failing ipmi connection? or [okurz](#), would it be feasible to keep the ipmi user/password secret in autoinst.txt by opening an openqa ticket?

**#38 - 2022-04-15 07:02 - Julie\_CAO**

I missed 'NOT' in my previous comment and I just corrected it, but I'd like to paste a new comment as the mailsystem might not have notice about my update.

"Our test does require to connect to SUSE internal network" => "Our test does NOT require to connect to SUSE internal network"

**#39 - 2022-04-20 06:35 - xlai**

[viktors.trubovics](#) Thanks for the suggestions. [nicksingergschlotterokurzmgruessmeier](#) Hello guys, as @Julie\_CAO confirmed, the tumbleweed virtualization tests to be put on this new zen2 machine won't need to access SUSE internal network, so we can accept whatever infra solution to ban that. Would you please let us know whether this ticket can be continued?

**#40 - 2022-04-21 15:40 - okurz**

[nicksingergschlotter](#) do you think it would be possible to create a new dedicated VLAN for that purpose?

**#41 - 2022-04-22 09:16 - nicksinger**

viktors.trubovics wrote:

The only way I see in this case, where IPMI must be exposed to internet - the server must be not able to connect to internal SUSE networks and 20 character strong unique password must be used for IPMI. In case the server will be hacked - SUSE network must stay secure.

The opensuse network is strictly separated from the SUSE network. My biggest concern is the fact that over IPMI a potential attacker could really dig into the system because it can control the whole machine completely. But this might be the case with a hacked linux too - not sure.

okurz wrote:

[nicksingergschlotter](#) do you think it would be possible to create a new dedicated VLAN for that purpose?

Should be possible. But we would need another jump host and I wonder if this would really change anything compared to the current VLAN where we also need a jump host (ariel) to gain access from the outside.

**#42 - 2022-04-22 09:26 - nicksinger**

- Due date set to 2022-04-25

- Assignee changed from okurz to nicksinger

I think with the stated requirements:

- 20 character password

- Not connected to the SUSE network

we should be fine with just connecting it to the current opensuse network. I will talk to Johannes Segitz on Monday (he's on FTO currently) to make sure we don't overlook anything. Assigning to me and setting due date as remember for me.

**#43 - 2022-04-24 11:55 - Julie\_CAO**

- Related to action #110227: Stop showing ipmi passwords in autoinst.txt from a ipmi backend job in O3 added

**#44 - 2022-04-26 09:18 - cdywan**

nicksinger wrote:

I think with the stated requirements:

- 20 character password
- Not connected to the SUSE network

we should be fine with just connecting it to the current opensuse network. I will talk to Johannes Segitz on Monday (he's on FTO currently) to make sure we don't overlook anything. Assigning to me and setting due date as remember for me.

Did you have a chance to talk to Johannes?

**#45 - 2022-04-27 07:11 - cdywan**

- Due date changed from 2022-04-25 to 2022-05-02

Let's wait a bit, given more urgent tickets

**#46 - 2022-04-27 11:56 - jstehlik**

I asked Johannes about this issue 14.April, now he is back and I told him to contact Nick directly. Victor also gave his opinion, so it seems to me we have enough information to decide and connect those machines as long as the proposed security measures are in place.

**#47 - 2022-04-27 12:00 - nicksinger**

jstehlik wrote:

I asked Johannes about this issue 14.April, now he is back and I told him to contact Nick directly. Victor also gave his opinion, so it seems to me we have enough information to decide and connect those machines as long as the proposed security measures are in place.

Yes, I talked to Johannes directly yesterday. We also came to the conclusions that the most important part is to never connect these machines to the SUSE network which isn't the case for o3 testing anyway. But he also recommended me to get in touch with Petr Spirik and Team as they're doing IT security in the company. [jstehlik](#) WDYT about this?

**#48 - 2022-05-02 10:54 - jstehlik**

Thank you [nicksinger](#) for making progress on this. I see no harm in asking Petr's team. The technical solution is getting clear and on top of that we might think of a process to ensure the machine is connected properly. For example the cable could be labelled, so we know it needs to stay out of internal network.

**#49 - 2022-05-02 13:38 - okurz**

- Subject changed from *Two new machines for OSD and o3, meant for bare-metal virtualization* to *Two new machines for OSD and o3, meant for bare-metal virtualization size:M*

**#50 - 2022-05-03 08:33 - okurz**

- Due date changed from 2022-05-02 to 2022-05-13

[nicksinger](#) as discussed please discuss security relevant implications and then at best continue as decided to put both the BMC and main machine ethernet interface into the openSUSE VLAN

**#51 - 2022-05-11 09:46 - cdywan**

Discussed briefly in the Unblock. This is still pending Nick talking to Petr for now.

**#52 - 2022-05-13 13:57 - cdywan**

- Due date changed from 2022-05-13 to 2022-05-20

cdywan wrote:

Discussed briefly in the Unblock. This is still pending Nick talking to Petr for now.

Email conversation on-going

**#53 - 2022-05-20 12:41 - cdywan**

- Due date changed from 2022-05-20 to 2022-05-27

No concrete update for now. Discussed briefly that Nick could probably go ahead at the next opportunity and consider the lack of objection sufficient.

**#54 - 2022-05-29 17:58 - okurz**

- Due date changed from 2022-05-27 to 2022-06-03

- Priority changed from Normal to High

[nicksinger](#) is the change something we can do ourselves within QA switches or EngInfra?

**#55 - 2022-05-30 08:42 - waynechen55**

The second link to new zen3 machine on OSD is down:

```
dhcpcd.conf
host amd-zen3-gpu-sut1-2          { hardware ethernet b4:96:91:9c:5a:d4; fixed-address 10.162.2.133; option host-name
"amd-zen3-gpu-sut1-2"; filename "pxelinux.0"; }
```

```
ping -c5 10.162.2.133
PING 10.162.2.133 (10.162.2.133) 56(84) bytes of data.
```

```
--- 10.162.2.133 ping statistics ---
5 packets transmitted, 0 received, 100% packet loss, time 4082ms
```

```
amd-zen3-gpu-sut1-1:~ # ip addr show
2: em1: mtu 1500 qdisc mq master br0 state UP group default qlen 1000
link/ether ec:2a:72:02:84:20 brd ff:ff:ff:ff:ff:ff
altname eno8303
altname enp225s0f0
3: em2: mtu 1500 qdisc noop state DOWN group default qlen 1000
link/ether ec:2a:72:02:84:21 brd ff:ff:ff:ff:ff:ff
altname eno8403
altname enp225s0f1
4: p3p1: mtu 1500 qdisc noop state DOWN group default qlen 1000
link/ether b4:96:91:9c:5a:d4 brd ff:ff:ff:ff:ff:ff
altname enp65s0f0
5: p3p2: mtu 1500 qdisc noop state DOWN group default qlen 1000
link/ether b4:96:91:9c:5a:d5 brd ff:ff:ff:ff:ff:ff
altname enp65s0f1
```

Could you help fix this [nicksinger](#) ?

**#56 - 2022-06-07 14:15 - cdywan**

- Due date changed from 2022-06-03 to 2022-06-10

Bumping because of availability / other urgent tickets

**#57 - 2022-06-08 10:20 - okurz**

- Due date changed from 2022-06-10 to 2022-06-17

waynechen55 wrote:

The second link to new zen3 machine on OSD is down: [...]

[waynechen55](#) please handle that in a separate ticket if you need this and help by others to resolve. This ticket is getting too big to tackle.

**#58 - 2022-06-15 09:18 - okurz**

- Due date changed from 2022-06-17 to 2022-07-01

nicksinger unavailable right now

**#59 - 2022-06-16 02:45 - waynechen55**

- Assignee deleted (*nicksinger*)
- Target version deleted (*Ready*)

okurz wrote:

waynechen55 wrote:

The second link to new zen3 machine on OSD is down: [...]

[waynechen55](#) please handle that in a separate ticket if you need this and help by others to resolve. This ticket is getting too big to tackle.

New ticket <https://progress.opensuse.org/issues/112553> created.

**#60 - 2022-06-16 05:28 - xlai**

- Status changed from *Feedback* to *Workable*
- Assignee set to *nicksinger*
- Target version set to *Ready*

**Files**

---

pxe_boot_failure.png	85.6 KB	2022-04-06	waynechen55
pxe_boot_failure_1.png	152 KB	2022-04-07	waynechen55