bhyve graphics

Peter Grehan
grehan@freebsd.org
BSD-TW 2017

What is bhyve ?

(from EuroBSDCon 2013, with additions)

- A “minimally viable x86 hypervisor”
- serial console, PCI virtio block/net, 64-bit host, 64/32-bit guests and extra dev emulations
- Requires Intel VT-x/EPT or AMD SVM/RVI CPU support
- In base-system FreeBSD as of 10.0

What is “bhyve graphics” ?

- The code bhyve that emulates a mouse, keyboard and a display

- Provides a workstation-style user experience

Hey, wasn't serial good enough ?

- A serial port emulation will support almost all guest o/s's
- However, the user-experience leaves a lot to be desired
  - A point re-iterated many, many times by users
  - Very different than existing hypervisor experiences

ubuntu 17.10 server install

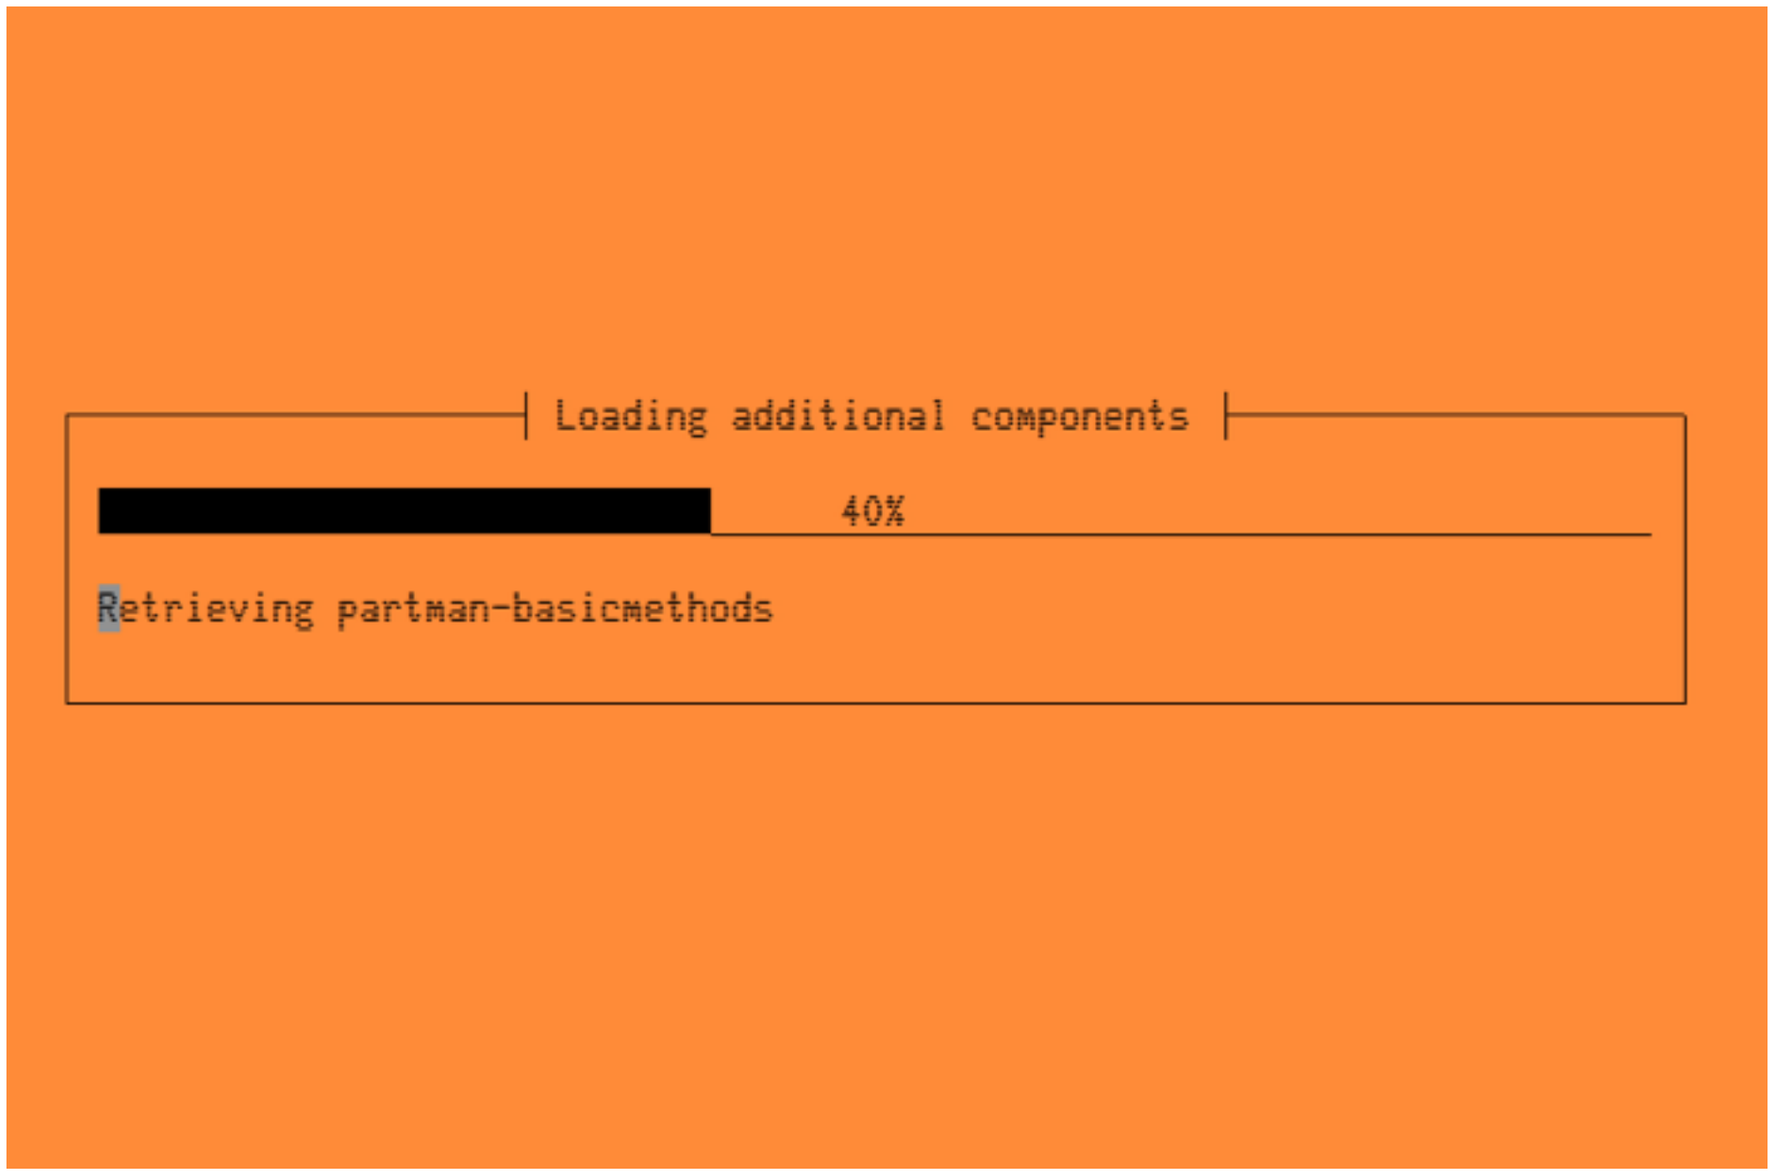

ubuntu 17.10 server login

```
[  OK  ] Reached target Remote File Systems.
         Starting LSB: automatic crash report generation...
         Starting Permit User Sessions...
[  OK  ] Started Permit User Sessions.
         Starting Hold until boot process finishes up...
         Starting Terminate Plymouth Boot Screen...
[  OK  ] Started Hold until boot process finishes up.
         Starting Set console scheme...
[  OK  ] Started Serial Getty on ttyS0.
[  OK  ] Started Terminate Plymouth Boot Screen.
[  OK  ] Started Set console scheme.
[  OK  ] Created slice system-getty.slice.
[  OK  ] Started Getty on tty1.
[  OK  ] Reached target Login Prompts.
[  OK  ] Started OpenBSD Secure Shell server.
[  OK  ] Started LSB: automatic crash report generation.
[  OK  ] Reached target Multi-User System.
[  OK  ] Reached target Graphical Interface.
         Starting Update UTMP about System Runlevel Changes...
[  OK  ] Started Update UTMP about System Runlevel Changes.

Ubuntu 17.10 ubuntu ttyS0

ubuntu login:
```

ubuntu 17.10 desktop ???

```
Successfully activated service 'org.gtk.vfs.Metadata'
A connection to the bus can't be made
Using CD-ROM mount point /cdrom/
Identifying... [e5af3c70da95efb4a84e894f28d2e469-2]
Scanning disc for index files...
Found 2 package indexes, 0 source indexes, 0 translation indexes and 1 signature
s
Found label 'Ubuntu 17.10 _Artful Aardvark_ - Release amd64 (20171018)'
This disc is called:
'Ubuntu 17.10 _Artful Aardvark_ - Release amd64 (20171018)'
Copying package lists...gpgv: Signature made Wed Oct 18 18:53:08 2017 UTC
gpgv:                using RSA key
gpgv: Good signature from "Ubuntu CD Image Automatic Signing Key (2012) <cdimage
@ubuntu.com>"
Reading Package Indexes... Done
Writing new source list
Source list entries for this disc are:
deb cdrom:[Ubuntu 17.10 _Artful Aardvark_ - Release amd64 (20171018)]/ artful ma
in restricted
Repeat this process for the rest of the CDs in your set.

Ubuntu 17.10 ubuntu ttyS0

ubuntu login:
```

What about Windows ?

- All 64-bit versions of Windows starting with Vista don't require a graphics adapter
- Server versions support "System Administrator Console" aka SAC; a tmux-like interface on the serial port.
- If an ACPI SPCR table is present, WinPE (1st phase install) will output to the serial port, and Windows server will instantiate SAC on this with VT100 emulation.
- Unattended install required
    - XML script, extremely version sensitive
    - Requires re-pack of UDF-formatted DVD, with virtio net driver "slipstreamed" in to allow RDP access post-install
    - A daunting install experience; black screen of death for desktop versions.
- This is how bhyve first booted Win 2k12 (UEFI already existed.

w2k12r2 install experience

```
Computer is booting, SAC started and initialized.

Use the "ch -?" command for information about using channels.
Use the "?" command for general help.


SAC>
EVENT: The CMD command is now available.
SAC>
EVENT:   A new channel has been created.  Use "ch -?" for channel help.
Channel: SACSetupAct
SAC>
EVENT:   A new channel has been created.  Use "ch -?" for channel help.
Channel: SACSetupErr
SAC>
```

w2k12r2 install error

```
2017-11-09 22:52:19, Info       [0x0605b3] IBS    Callback_WinPE_SetSourceMedia
Info_Unattend:Show flag for early F6 UI page is being to set to [Hide]
2017-11-09 22:52:19, Info       [0x064047] IBSLIB MarkUnattendSettingAsProcesse
d: Marking unattend setting [Diagnostics] as processed.
2017-11-09 22:52:19, Info       [0x064047] IBSLIB MarkUnattendSettingAsProcesse
d: Marking unattend setting [ProductKey\ProductKey] as processed.
2017-11-09 22:52:19, Info       [0x060114] IBS    Callback_Productkey_Validate_
Unattend:User specified an unattend xml file.
2017-11-09 22:52:19, Info       [0x060115] IBS    Callback_Productkey_Validate_
Unattend:Using ProductKey WillShowUI value of [OnError]; retrieving key from un
attend file...
2017-11-09 22:52:19, Info                  IBS    Callback_Productkey_Validate_
Unattend:Look for ei.cfg.
2017-11-09 22:52:19, Info                  IBS    Callback_Productkey_Validate_
Unattend:Look for pid.txt file.
2017-11-09 22:52:19, Info       [0x060118] IBS    Callback_Productkey_Validate_
Unattend:Validating Product key
2017-11-09 22:52:19, Error      [0x060120] IBS    Callback_Productkey_Validate_
Unattend: An error occurred preventing setup from being able to validate the pr
oduct key; hr = 0x80070002[gle=0x00000002]
2017-11-09 22:52:19, Info       [0x0640ae] IBSLIB PublishMessage: Publishing me
ssage [Windows cannot read the <ProductKey> setting from the unattend answer fi
le.]
```

w2k12r2 SAC

```
SAC>?
ch                      Channel management commands.  Use ch -? for more help.
cmd                     Create a Command Prompt channel.
d                       Dump the current kernel log.
f                       Toggle detailed or abbreviated tlist info.
? or help               Display this list.
i                       List all IP network numbers and their IP addresses.
i <#> <ip> <subnet> <gateway> Set IPv4 addr., subnet and gateway.
id                      Display the computer identification information.
k <pid>                 Kill the given process.
l <pid>                 Lower the priority of a process to the lowest possible.
lock                    Lock access to Command Prompt channels.
m <pid> <MB-allow>      Limit the memory usage of a process to <MB-allow>.
p                       Toggle paging the display.
r <pid>                 Raise the priority of a process by one.
s                       Display the current time and date (24 hour clock used).
s mm/dd/yyyy hh:mm      Set the current time and date (24 hour clock used).
t                       Tlist.
restart                 Restart the system immediately.
shutdown                Shutdown the system immediately.
crashdump               Crash the system. You must have crash dump enabled.
SAC>
SAC>
SAC>
```

The UEFI frame buffer

- The OVMF build of UEFI supported the “Graphics Output Protocol” interface
  - Qemu had a number of SVGA+ emulations; S3, etc
  - However, UEFI only requires a linear frame buffer
  - A random experiment in providing this at a fixed address showed Windows writing to this, even with no PCI adapter for it.
- Could this be the solution for installation ?
  - Would it work with other o/s’s ?
    - Yes, it totally did.

Getting bits to users

- bhyve is a FreeBSD base-system component
    - so, can’t link against Xorg/SDL libs
- VNC looked the obvious next choice
    - IETF spec
    - Many free clients available
    - Protocol didn’t appear too onerous
    - Also provided keyboard/mouse input

Prototyping VNC

- Started out writing a simple program using Cairo rendering chars to a bitmap on keyboard input
- Used the GPL'd libvnc to get something going.
  - Very useful to tcpdump to see what really goes on
- Then, started implementing a from-scratch BSD-licensed version
  - Harmed many VNC viewers in the process
  - This gave the skeleton of a VNC server

The “fbuf” device emulation

- bhyve implements a proprietary PCI frame buffer emulation
  - The opposite of the bhyve only-emulate-well-supported-devices philosophy
    - But, with UEFI, guests accept what the GOP protocol reports
      - and since bhyve provides UEFI, a custom driver is provided for this
- 8MB of frame buffer memory provided with 32-bit pixels. Resolution can be changed on the fly (though usually only in UEFI on a GOP requests)

“fbuf” #2

- Guest frame buffers accesses are NOT emulated, but passed through
    - No instruction emulation required (most likely exotic instr’s)
- Host memory is inserted into the guest EPT map to create the frame buffer
        - Marked as non-executable
- Guest rendering runs at memory bandwidth
        - Not through slow vesa/scfb bus-attached mem
        - But, non-accelerated

fbuf: meet VNC

- fbuf memory has to be passed to a user
- The age-old technique of “screen scraping” is used
    - The screen is sampled every 1/30 seconds
- Sending the entire frame buffer in this interval requires a lot of b/w
    - So, run a CRC over each 32x32 pixel “tile”, and don’t send the region if the CRC of the current screen is equal to the previous.
        - If more than XXX % of cells have changed, give up and send the whole thing instead of tiny rectangles
    - A simple compression technique that works very well, and doesn’t use a lot of CPU when idle.
    - Some older VNC clients really don’t like it.
    - Also: use zlib compression if the client supports it.
- No CPU used when VNC not connected

The mouse

- VNC provides absolute mouse coordinates
- bhyve emulates a PS2 keyboard/mouse controller which hooks up to the VNC server
- Unfortunately, the PS2 mouse only supplies relative coordinates.
  - Most guests implement s/w mouse acceleration
  - fbuf provides no h/w cursor rendering (avoiding guest drivers), so guests will s/w render
  - This results in the mouse “running away” from the VNC cursor

Fixing the mouse

- Tablet devices provide absolute coordinates
    - Good match for VNC
- Hard to find a generic tablet device
- Solution was to implement a USB XHCI controller
    - XHCI supports MSI, and much friendlier on the system than an EHCI controller emulation (legacy intr's only, polling required)
        - Provides a future path for USB emulations/passthru/etc
    - Tablet device attachment hooked up to VNC to provide absolute coords
    - Downside: XHCI not supported in older guests (Win7)
    - Bigger downside: still not supported by FreeBSD :(

The keyboard

- The bhyve ps2 keyboard emulation was hooked up to VNC
- Every single x86 guest in existence supports this h/w
  - A USB keyboard attached to XHCI doesn't add much value
- VNC provides Xorg scan-codes
  - A very reasonable and proven format
  - Unfortunately, the mix of VNC clients and non-US keyboards creates a torrid mix
  - A large percent of bhyve graphics issues relate to this issue

w10-ct install

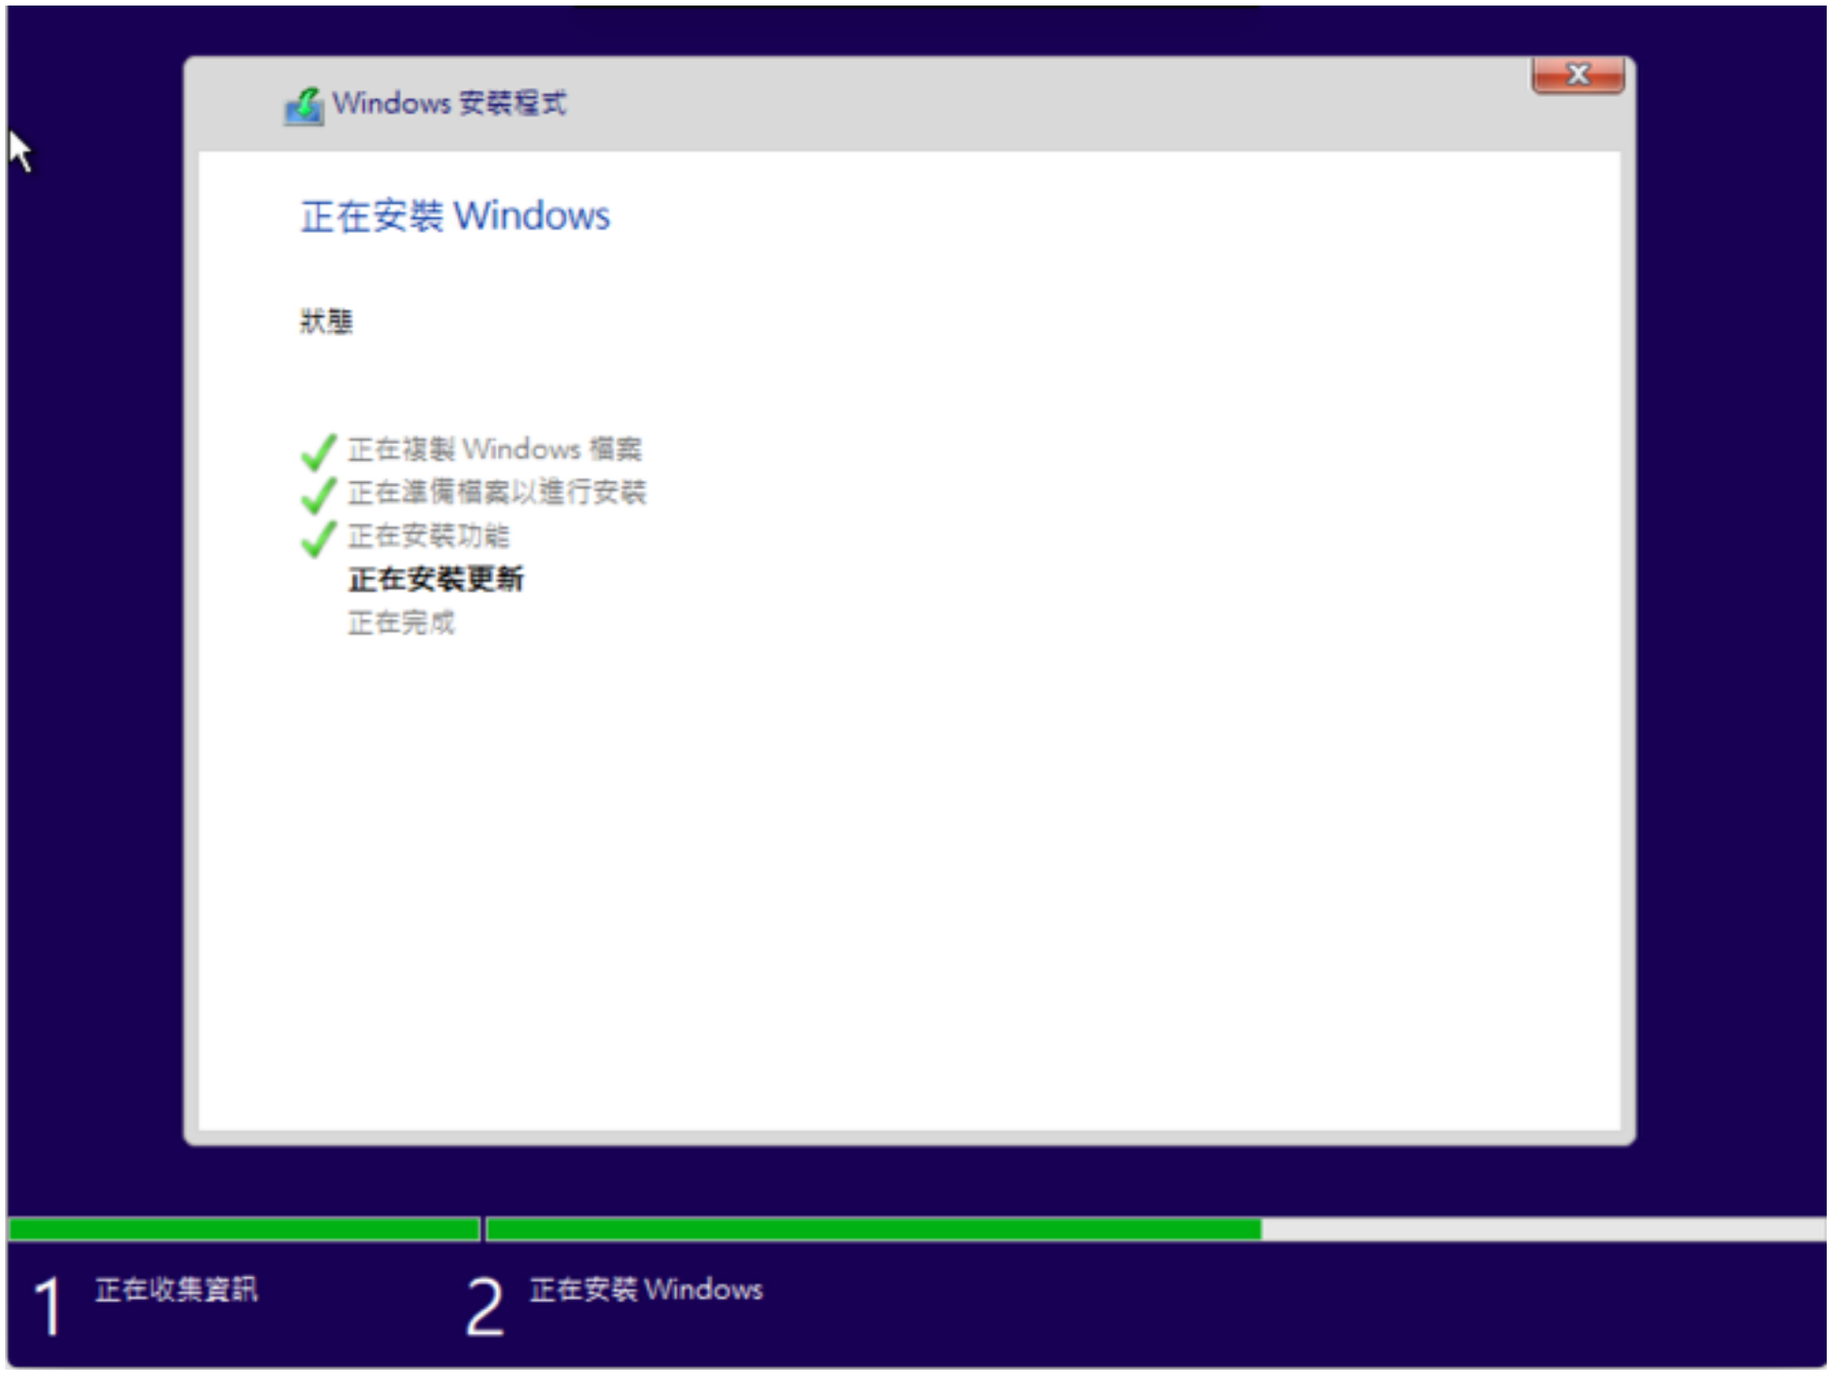

w10-ct setup


服務
選擇裝置的隱私設定
Microsoft 讓您掌控隱私。請選擇您的設定，再選取 [接受] 儲存它們。您可以隨時變更這些設定。
診斷
在基本層級，您傳送給 Microsoft 協助修正錯誤的資料較少。
基本
相關廣告
您看到的廣告數量不會變更，但可能與您較無關聯。
關閉
以診斷資料量身打造的體驗
您獲得的提示會較為一般，且建議也可能與您較無關聯。
關閉
請選取 [深入了解] 以取得上述設定的相關資訊，以及 Windows Defender SmartScreen 運作及相關資料傳輸與使用的方式。
深入了解
接受

ubuntu 17.10 desktop

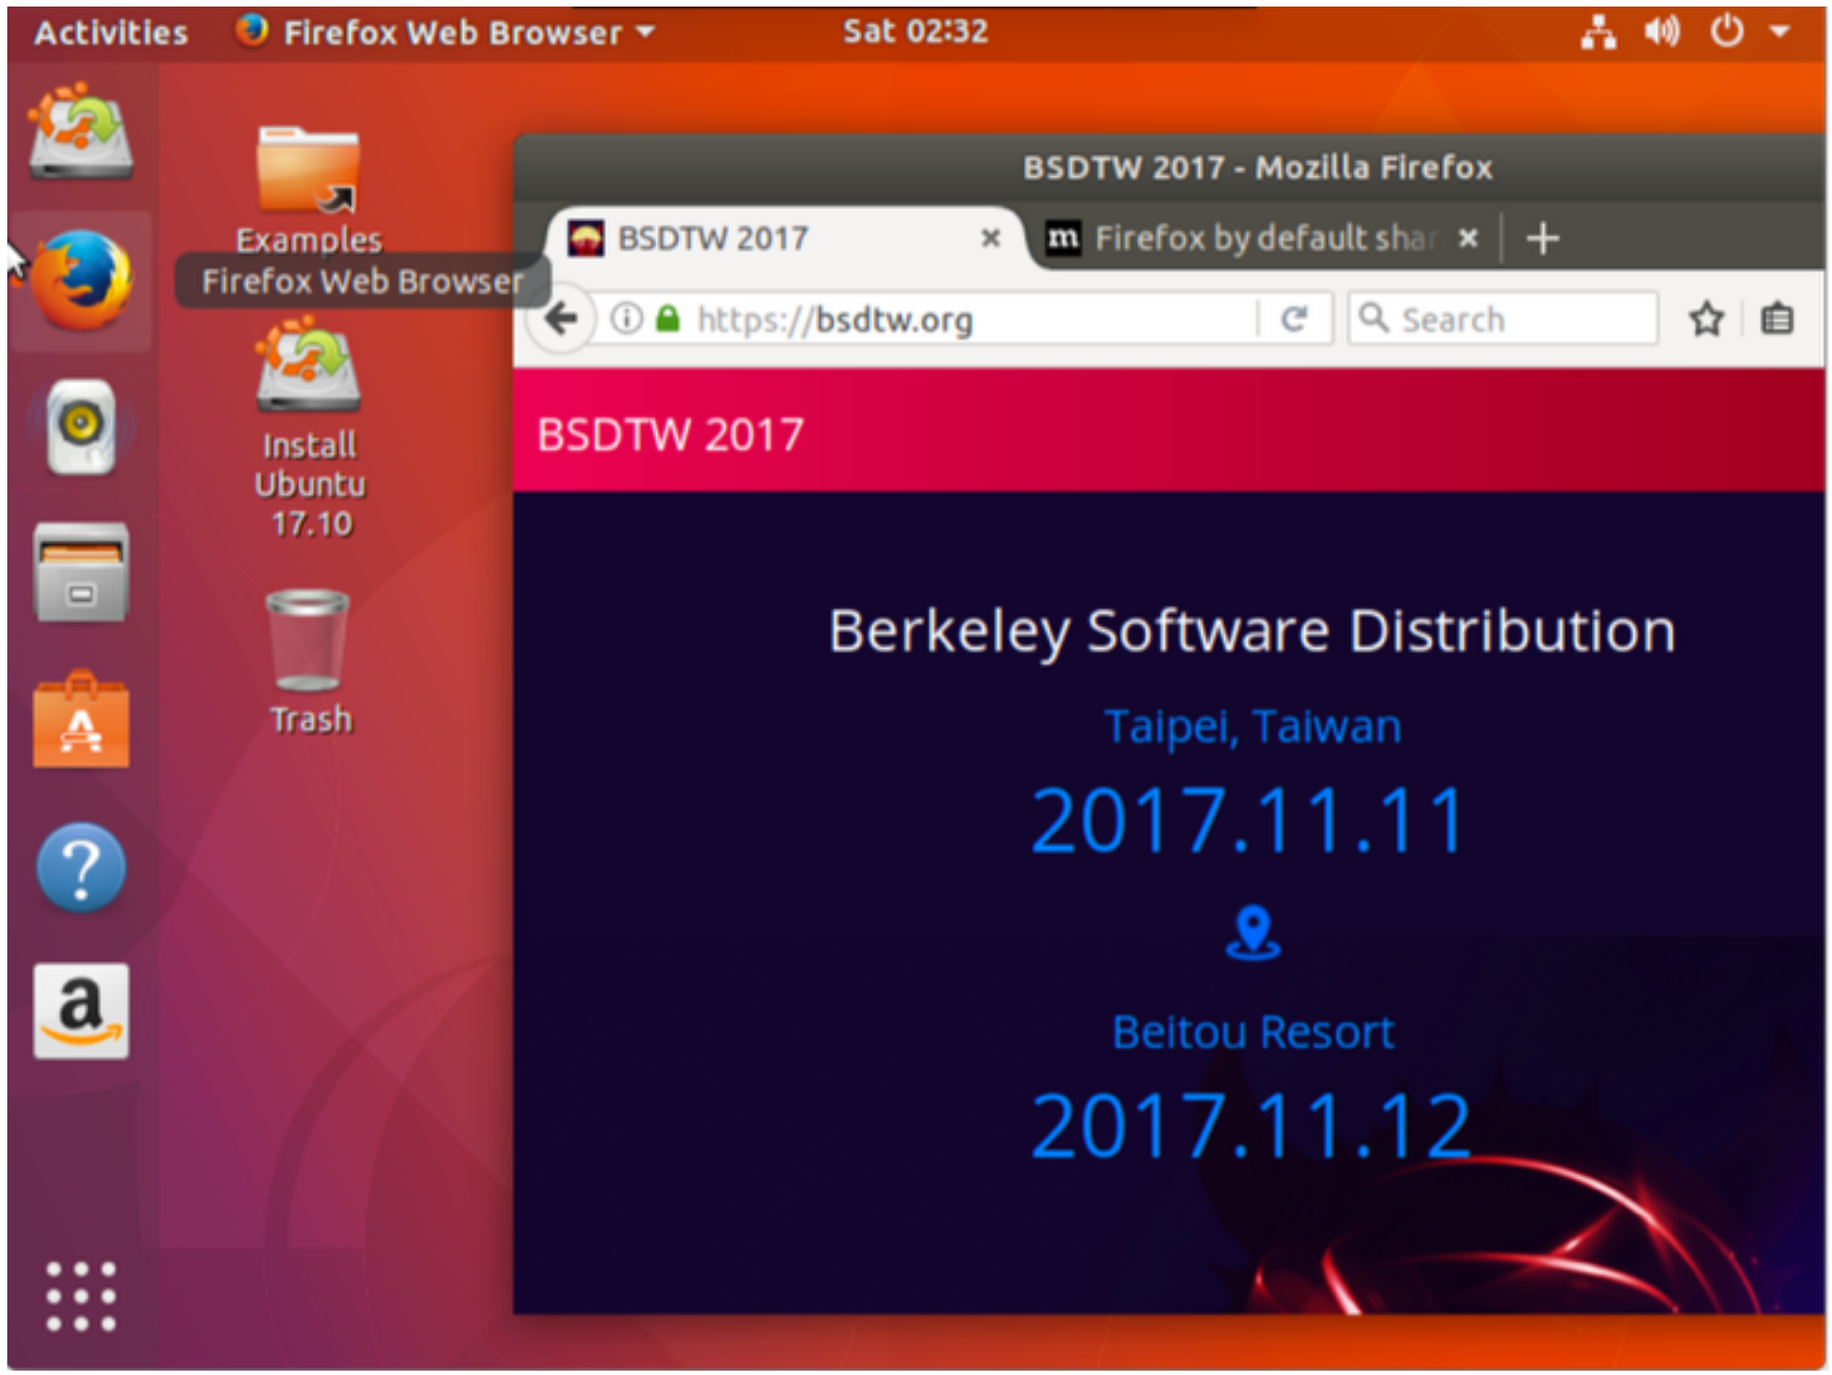

What about VGA ?

- VGA (more likely SVGA+) emulation was considered early on
- However, it is extremely complicated to support all the various modes
    - Only supports a limited resolution.
- An emulation has been written, though not fully enabled
    - Renders the various modes into the linear frame-buffer to allow export with VNC
    - Requires trapping all accesses for fidelity e.g. planar modes: extremely slow
    - UEFI CSM has a BIOS INT10 interface to support this
- Existing keyboard/mouse

The collision of VGA and UEFI

- Windows Vista, 7, and Server 2k8 require both UEFI GOP and BIOS INT10h
- VGA registers are accessed
- Fixed by:
    - an INT10h 16-bit asm stub in non-CSM UEFI that reports the required VESA BIOS info
    - partial VGA register implementation
- Unfortunately this breaks OpenBSD UEFI since it thinks the system is VGA, so yet another option required to allow that (“vga=off”)
- Currently forces resolution to 1024x768 regardless of config, though that is fixable

Futures

- PCI passthru of graphics adapters
  - Allows full-speed 3D rendering
- USB keyboard support
- Expand VNC client support
  - Fix language keyboard issues
- External API for non-VNC viewers (FreeRDP, Spice)
- Worth supporting virtio graphics emulation ?

Thanks to

- Leon Dang for most of this work
- tychon@freebsd.org for VGA/VNC/ps2 work
- neel@freebsd.org for guest EPT mem support
- peter.fang@tidalscale.com for Win10/CT support

w10-ct

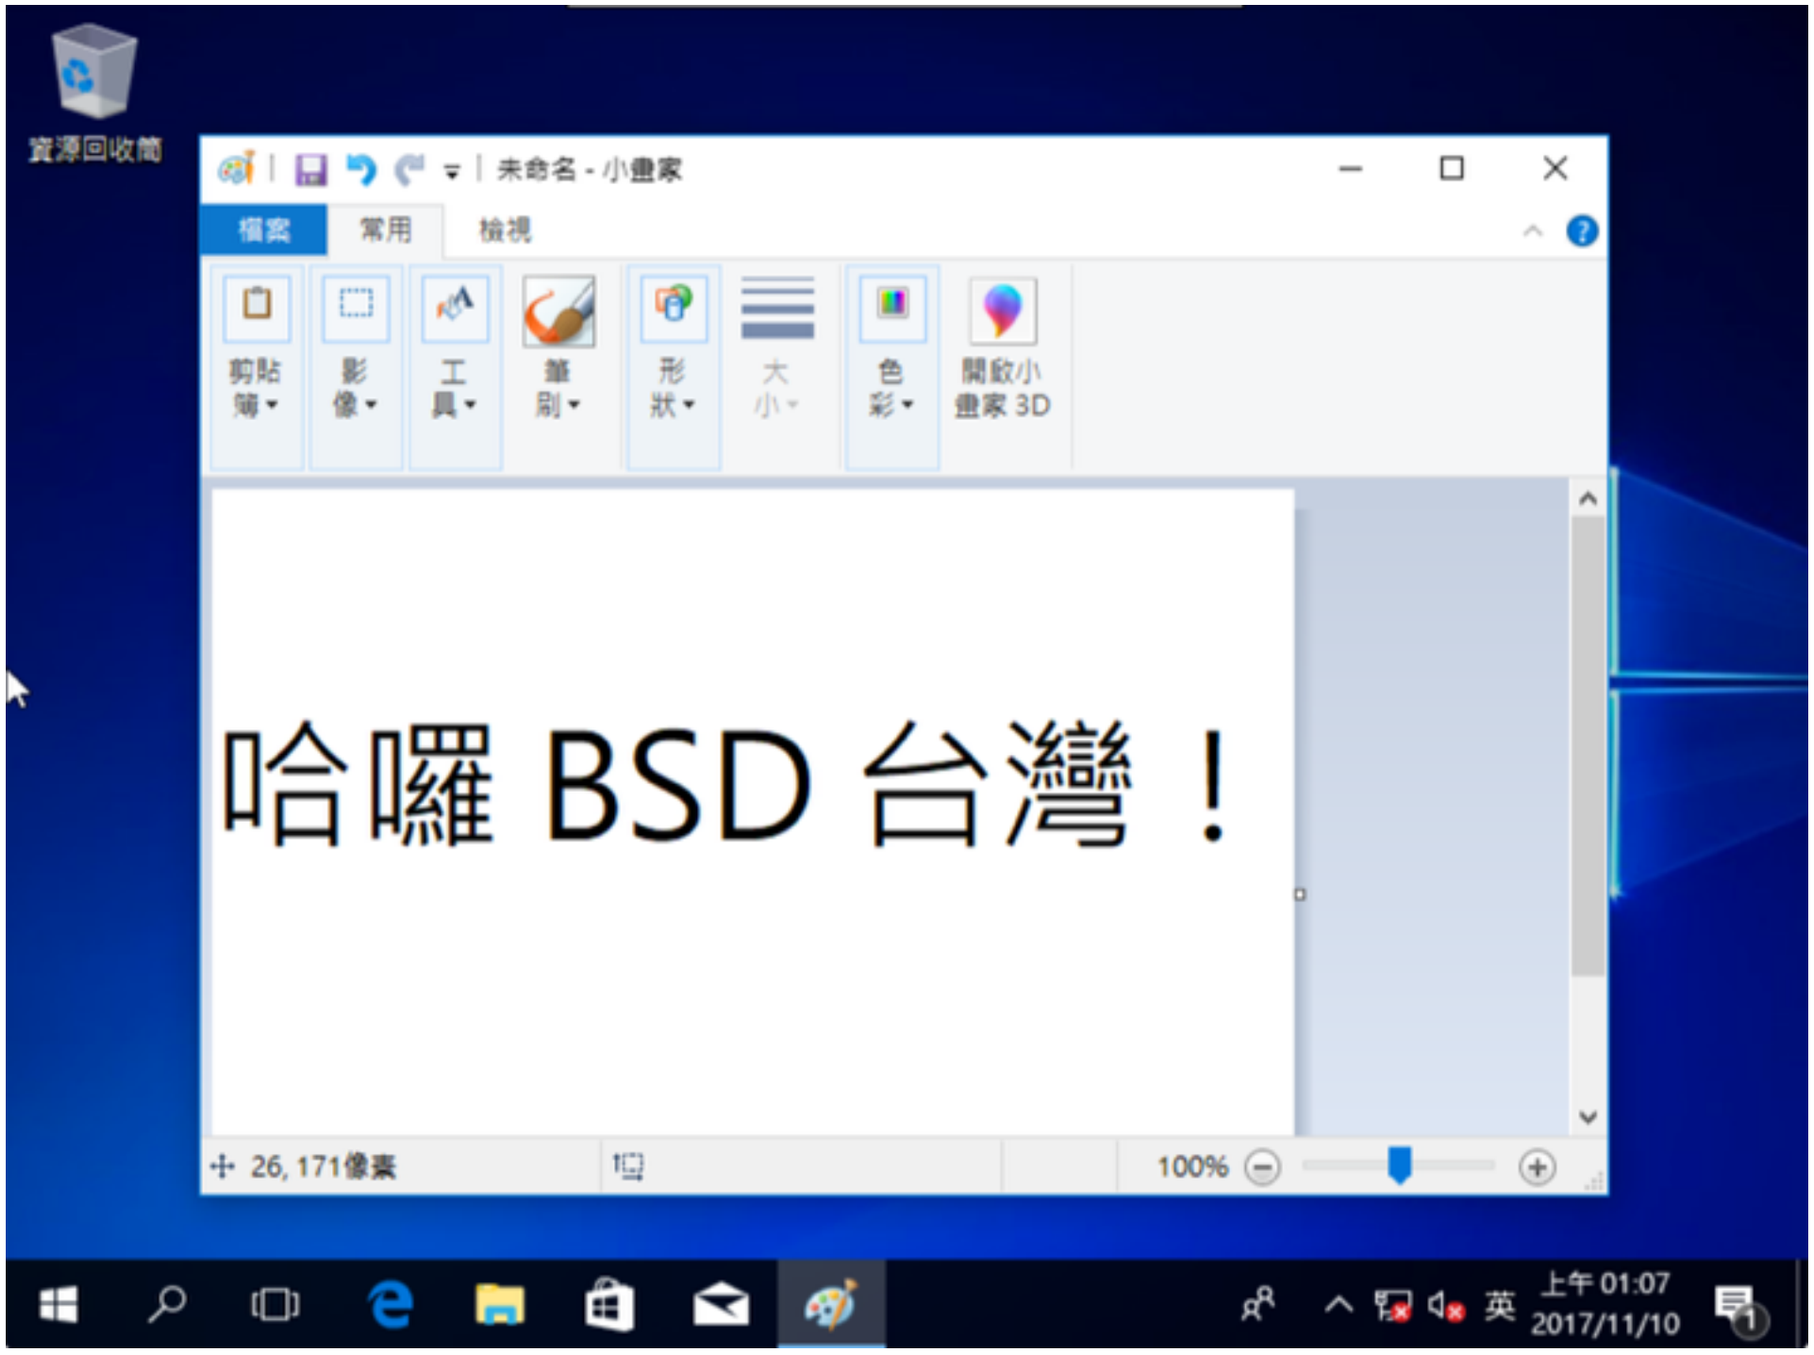
資源回收筒
未命名 - 小畫家
檔案
常用
檢視
剪貼簿
影像
工具
筆刷
形狀
大小
色彩
開啟小畫家 3D
哈囉 BSD 台灣！
26, 171像素
100%
英
上午 01:07
2017/11/10