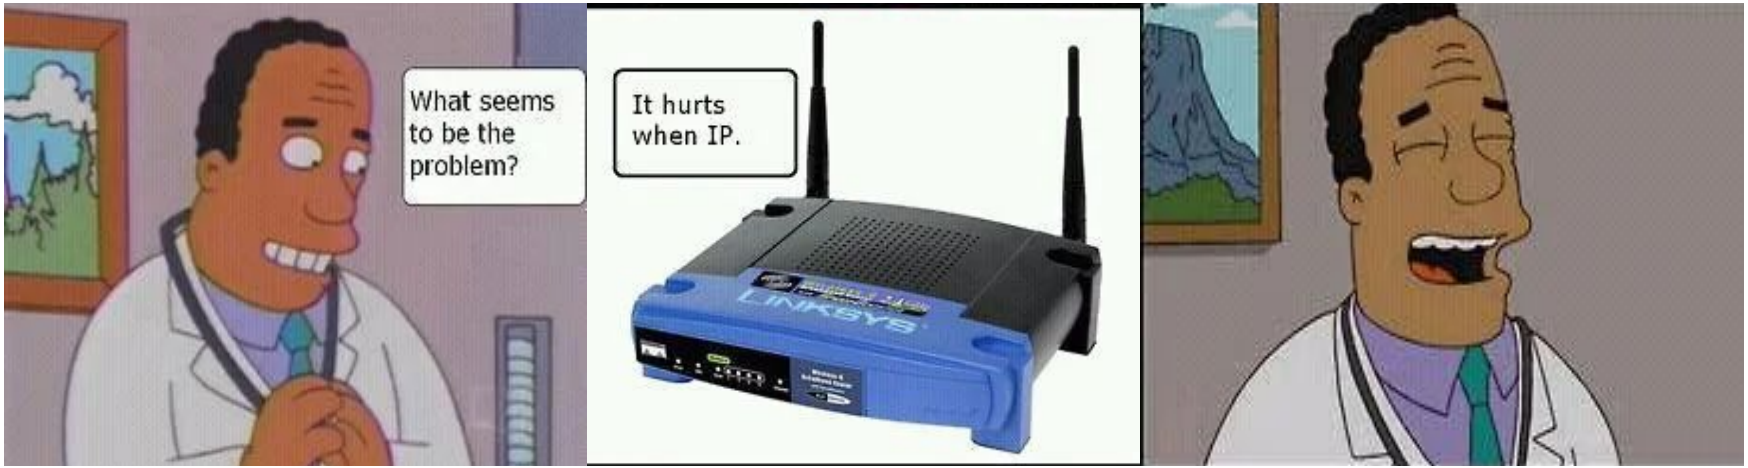

That’s NAT funny

source: https://www.reddit.com/r/pics/comments/fcytn/it_joke_of_the_day/

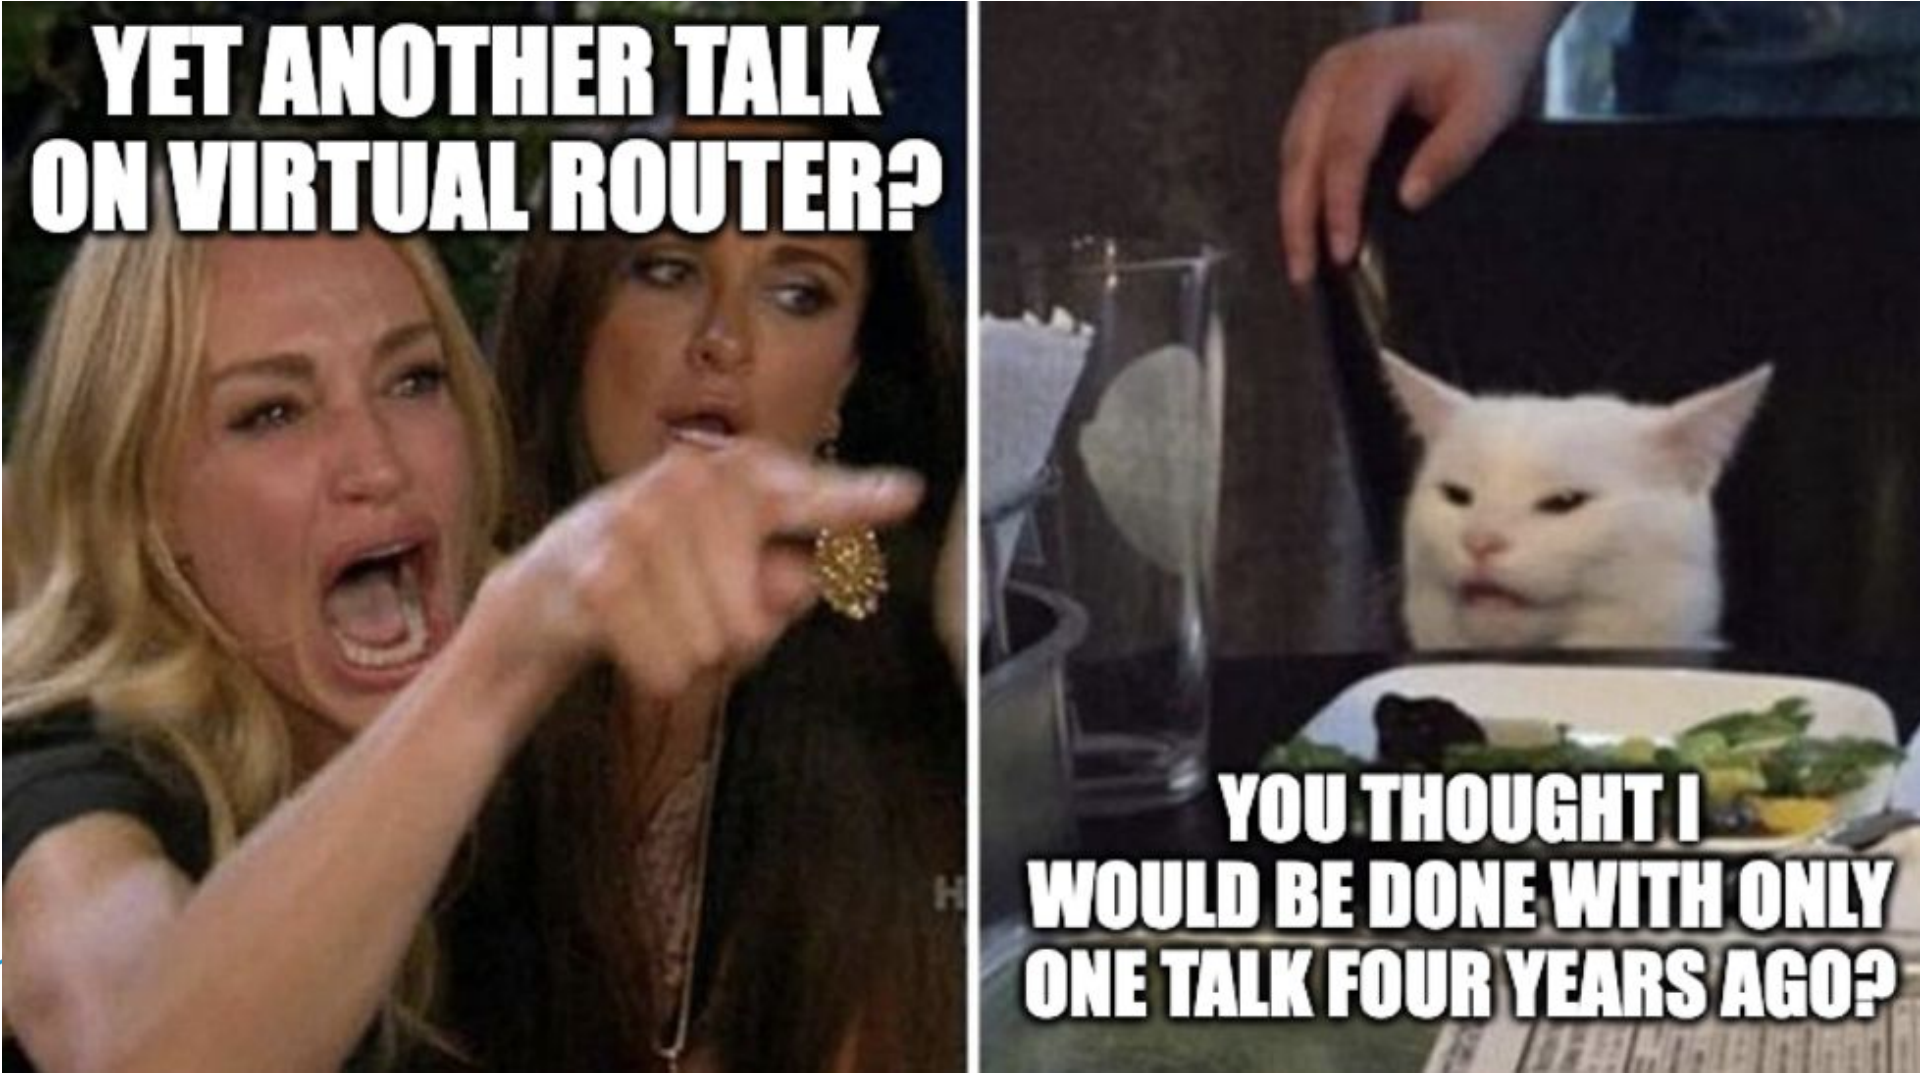
YET ANOTHER TALK ON VIRTUAL ROUTER?
YOU THOUGHT I WOULD BE DONE WITH ONLY ONE TALK FOUR YEARS AGO?

WHY?

Service tickets by logical areas

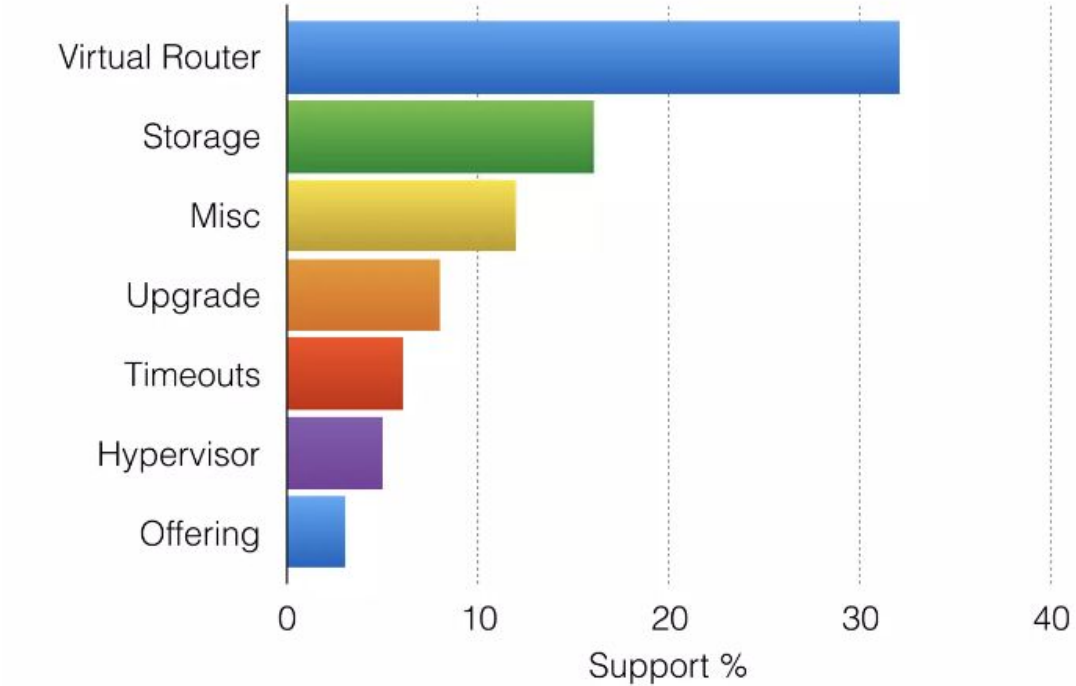

Source: (2017) https://www.slideshare.net/ShapeBlue/cloudstack-top-5-technical-issues-and-troubleshooting

Episode IV: A New Hope

Next-Gen Virtual Router

CCC Virtual, 11 November 2021

The Future of CloudStack Virtual Router

CCCNA Las Vegas, 10 September 2019

CloudStack Virtual Router: Past, Present, Future

CCCNA Montreal, 24 September 2018

Next-Gen Virtual Router And Zero Downtime Upgrades

cloud STACK COLLABORATION CONFERENCE 2022

Rohit Yadav

VP Engineering, ShapeBlue

rohit.yadav@shapeblue.com

cloud STACK COLLABORATION CONFERENCE

Shape Blue

$ whoami

- **Rohit Yadav,** VP Engineering @ ShapeBlue
- From Gurugram, India
- 10+ years CloudStack Committer and PMC
- Mentoring and training, leadership, design and architecture
- Past: authored several flagship features and tools, release manager and maintainer
- Married, 2x🐈 cats, loves mentoring and programming

About ShapeBlue

"ShapeBlue are expert builders of public & private clouds. They are the leading global Apache CloudStack integrator & consultancy"

https://www.shapeblue.com/shapeblue-has-become-an-employee-owned-business/

Agenda

- History & Changes
- SystemVM Template
- Live-Patching
- Zero-Downtime Upgrades
- Demo
- Q&A

How SystemVMs are created/patched then?

- Non-standard files, patching process
- Attached to a systemvm while booting:
  - KVM: kvm agent attaches local systemvm.iso
  - VMware: attached systemvm.iso on sec storage
  - XenServer: attached systemvm.iso from local location

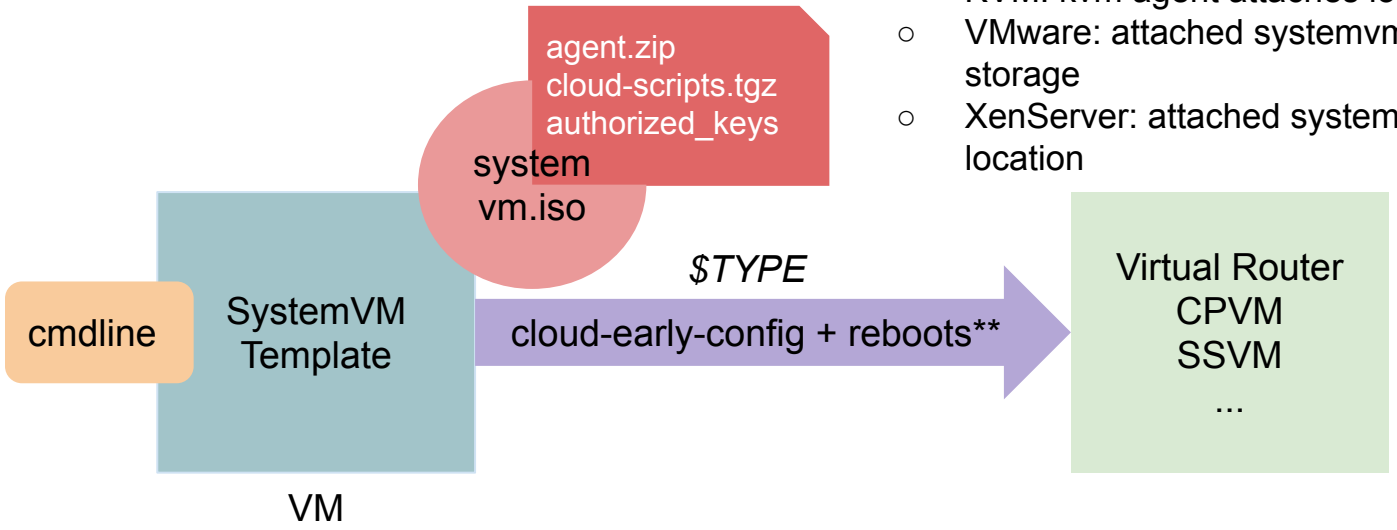

How SystemVMs are created/patched now?

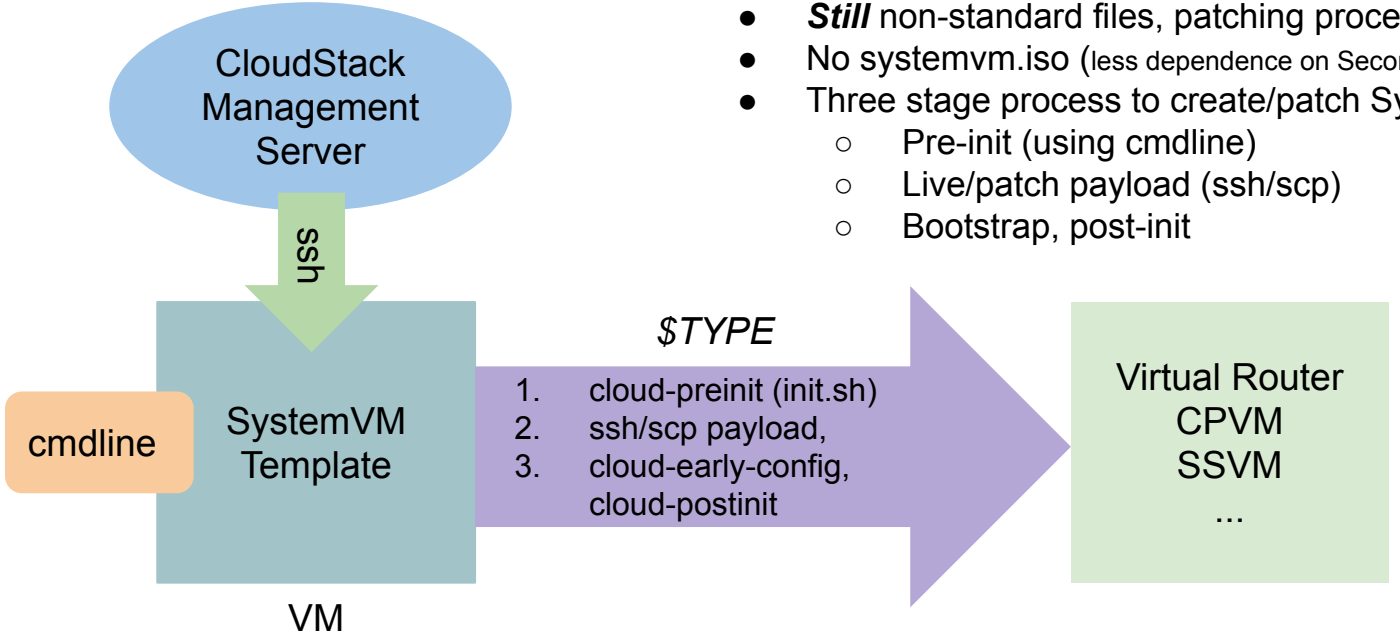

- ***Still*** non-standard files, patching process
- No systemvm.iso (less dependence on Secondary Storage)
- Three stage process to create/patch SystemVMs/VRs:
  - Pre-init (using cmdline)
  - Live/patch payload (ssh/scp)
  - Bootstrap, post-init

SystemVM Template Changes

- In ACS 4.16 and onwards
- Turnkey Packages: SystemVM templates bundled within cloudstack-management package (metadata.ini, maven profile)
- Automated Lifecycle: seed/install/upgrade
- Also used for CKS template
- Upgrade made simpler!

  apt upgrade | dnf update

Further reference: https://www.shapeblue.com/systemvm-template-upgrade-improvements/

SystemVM Software Setup

- “~~systemvm.iso~~” **is history!**
- Software upgrade/patching via ssh/scp
- Equivalent or Faster Live-Patching!
- Outcomes:
  - Faster management server startups - no more injectkeys.sh
  - Secure payload transfer (scp/ssh)
  - No dependency of secondary storage

Further reference: https://www.shapeblue.com/systemvm-template-upgrade-improvements/

SystemVM and VR Lifecycle Changes

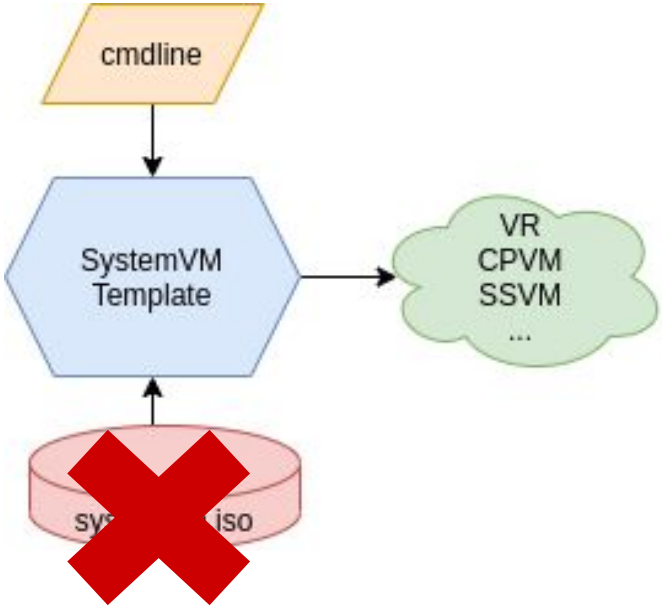

4.16/4.17 onwards

- Build (packer)
- Packaged in rpm/deb (+ hosted 3rd/community servers)
- Automated lifecycle (install, upgrade…)
- Lifecycle:
  - Copy disk, setup network (link-local/private nic)
  - Init/Config: cmdline
  - Live/Patching: ~~systemvm.iso~~ ssh/scp
  - VR Comms: ssh + databags, json/xml
  - SSVM/CPVM Comms: agent cmd/answer
  - Upgrade:
    - Recreate old SystemVM/VR
    - Zero-downtime live-patch!

Zero-Downtime Upgrades!

- Supported with 4.17 and onwards!
- Live-Patching:
  Operators/Admins can consider upgrading VRs using live-patching mechanism without network-downtime for guest networks** and VMs
- Bring Your Own Router:
  Self-service Shared Networks Feature (Alex's Talk)
  Further Reading: https://www.shapeblue.com/self-service-shared-networks/

** consider upgrade conditions and limitations

Traditional ACS Upgrade Process

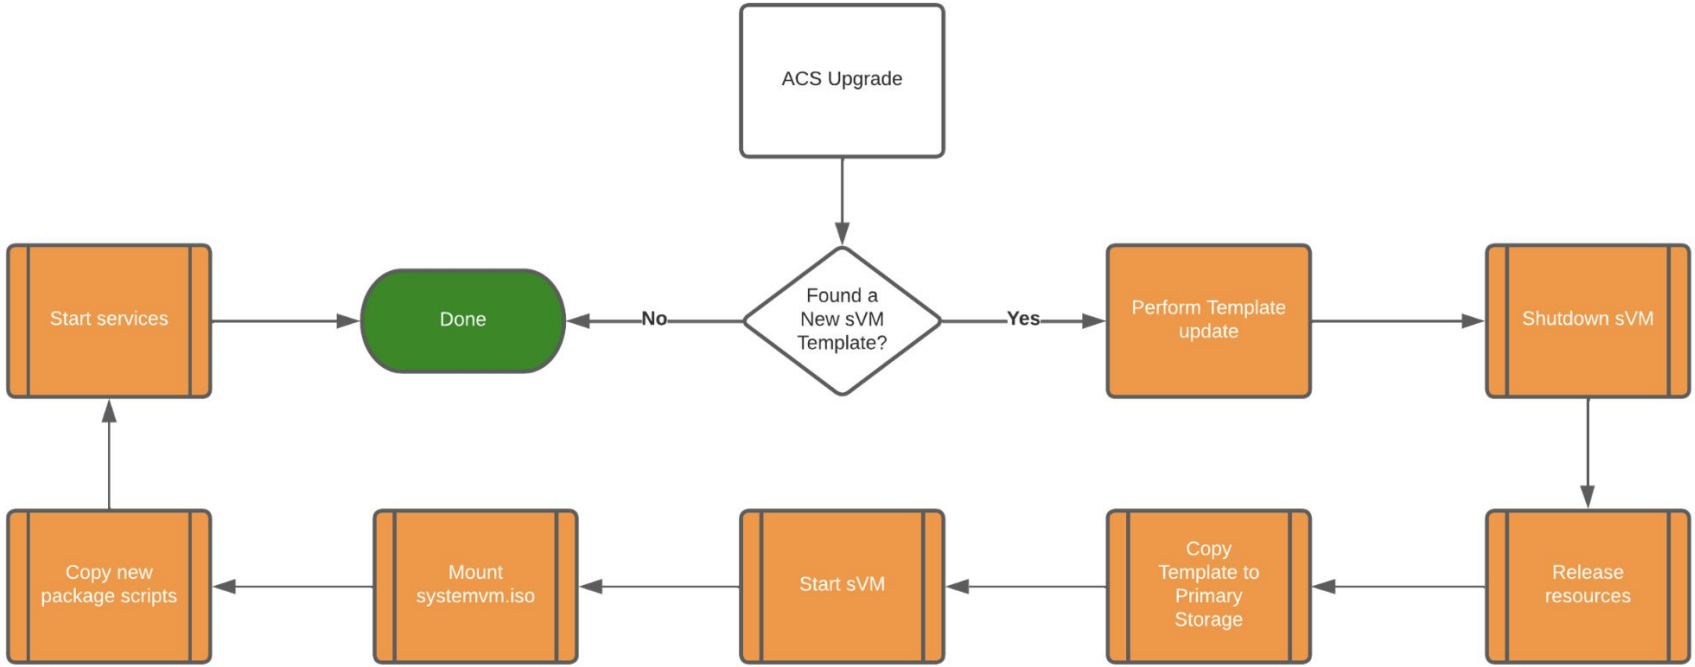

VR: Core vs Non-Core Services?

What causes network downtime?

- Firewall, ACLs
- NAT, SNAT/DNAT
- Forwarding Rules
- Guest Networking
- Static Routes
- VPN**

Typically Linux Kernel Feature

What causes service downtime?

- DNS, DHCP (dnsmasq)
- Password Server
- Metadata (apache2)
- LB (haproxy)
- Redundancy (Keepalived, conntrackd)
- Misc (health, monitoring, network stats)

Typically User-Space Apps

Live-Patching

- New Admin API and API changes:
  - patchSystemVM (forced=true|false)
  - restartNetwork (livepatch=true)
  - restartVpc (livepatch=true)
- Secure copy (scp) agent.zip, cloud-scripts.tgz to SystemVM/VR
- Runs *patch-sysvms.sh*:
  - Backup existing software (scripts, config and certificates)
  - Patch using latest payload software (scripts)
  - Restart services as per /var/cache/cloud/enabled_svcs
  - Post-patch checks, store checksum of payload (cloud-scripts-signature)
  - On services restart fail -> Restore backup software version

Upgrade: Live Patching

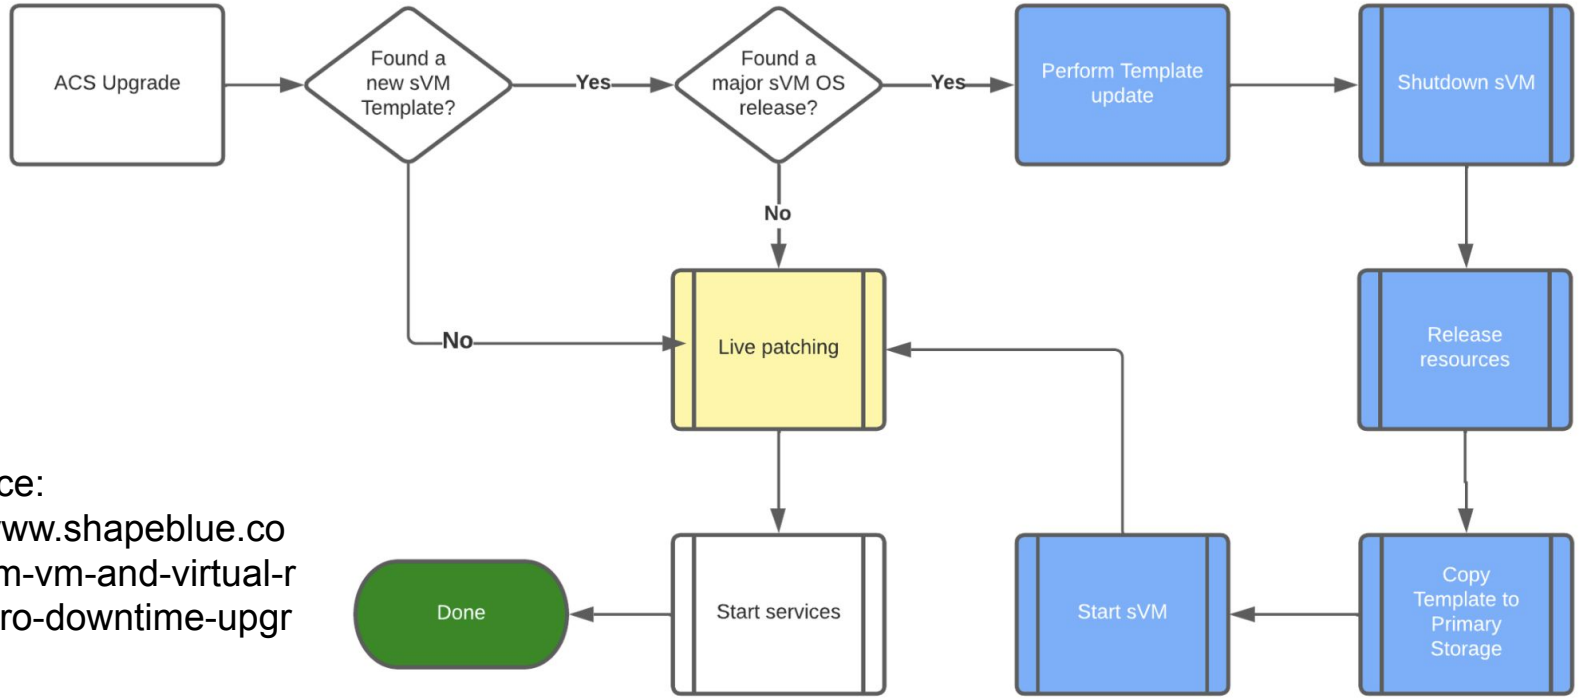

Reference:
https://www.shapeblue.com/system-vm-and-virtual-router-zero-downtime-upgrade/

Live-Patching Limitations

System VM	Services
SSVM	cloud, apache2, portmap
CPVM	cloud
VRs	haproxy, apache2, dnsmasq

Services that will be restarted

- Some (non-core) services will be restarted
- SystemVM template upgrade is unavoidable:
  - base template changes or base template OS EOL
  - core package change (JRE)
  - security issue (kernel)
  - template built-in/patching scripts

Live-Patch/Upgrade Considerations

ACS Version	Upgrade Version	Live Patching Support	Reason / Comment
<=4.13	4.17+	No	Update in the openJDK version
4.14	4.17+	Yes	May notice some issue with remove access VPN due to older version of Strongswan used in 4.14 template
>=4.15	4.17+	Yes	N/A

Further reference: https://www.shapeblue.com/system-vm-and-virtual-router-zero-downtime-upgrade/

Demo Time!


Do we get to mess with him now?
No. We wait until the worst possible moment.
CI/CD Demo
Rashtab 2018

Future Roadmap

- Standard init config (user-data/cloud-init?)
- Remove custom init/patching script
- Smaller SystemVM template:
  - Bundle single hypervisor image
  - Use qemu-img for on-the-fly conversion
- Container based systemvm-app upgrades (podman+systemd)

Future Work: Episode V?

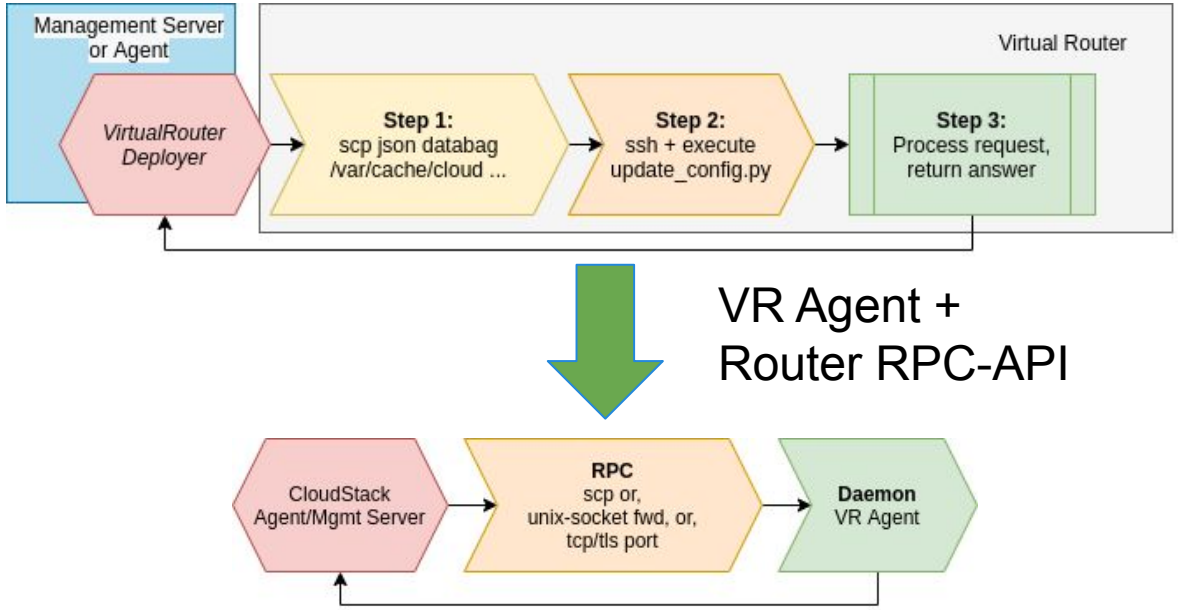

Q&A