Live Patching, Virtual Machine Introspection & Vulnerability Management

Lars Kurth
Community Manager, Xen Project
Chairman, Xen Project Advisory Board
Director, Open Source, Citrix


larskurth

Cheng Zhang
Software Engineer, Citrix
Currently working for XenServer Livepatch integration and new packaging framework


aiwei2013214


Presentation on xenbits.xenproject.org/people/larsk/

Xen is the Engine, XenServer is the Car

Simon Crosby, XenSource Inc.

Xen Project: www.xenproject.org
XenServer: www.xenserver.org

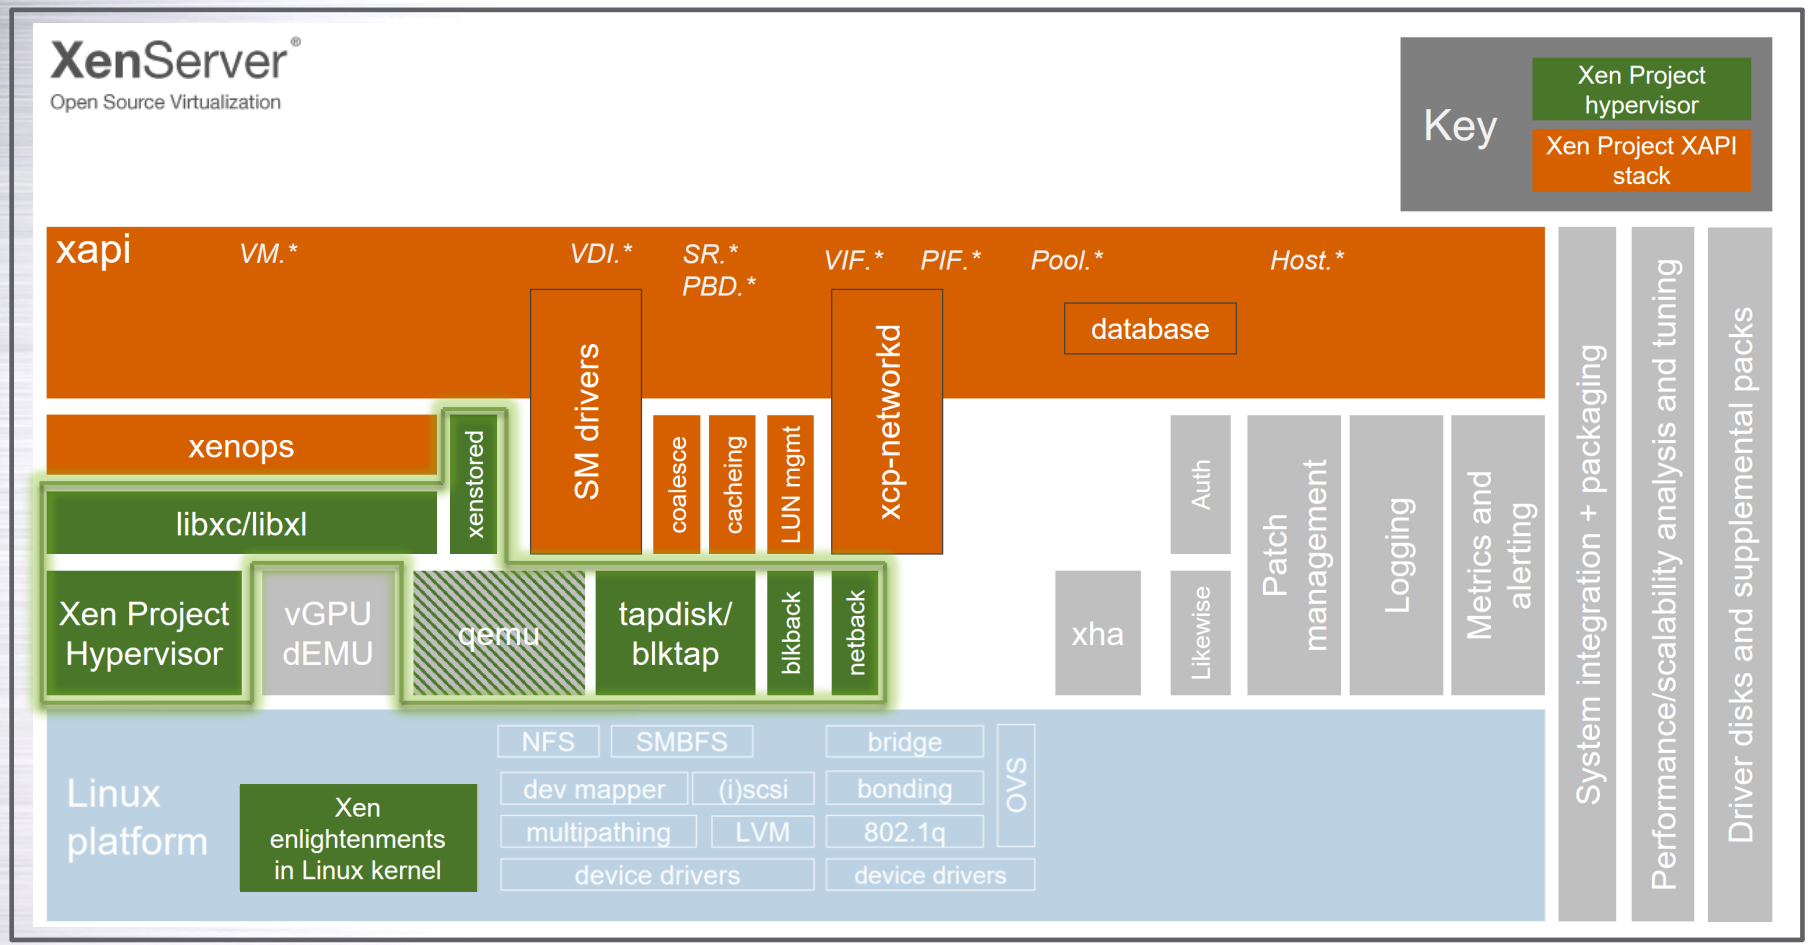
XenServer
Open Source Virtualization
Key
Xen Project hypervisor
Xen Project XAPI stack
xapi
VM.*
VDI.*
SR.*
PBD.*
VIF.*
PIF.*
Pool.*
Host.*
database
SM drivers
xcp-networkd
xenops
xenstored
libxc/libxl
coalesce
cacheing
LUN mgmt
Xen Project Hypervisor
vGPU dEMU
qemu
tapdisk/ blktap
blkback
netback
xha
Auth
Likewise
Patch management
Logging
Metrics and alerting
System integration + packaging
Performance/scalability analysis and tuning
Driver disks and supplemental packs
Linux platform
Xen enlightenments in Linux kernel
NFS
SMBFS
bridge
dev mapper
(i)scsi
bonding
multipathing
LVM
802.1q
device drivers
device drivers
OVS


SOFTLAYER
an IBM Company
ORACLE
CLOUD
Amazon Lightsail
kt
CITRIX
SUSE
ORACLE
amazon
web services
Tencent 腾讯
inspur 浪潮
rackspace
the open cloud company
DATAPIPE
HUAWEI
Alibaba Cloud
aliyun.com
XZD
XENZYNQ
DISTRIBUTION
NXP
Bitdefender
ais
epam
Xenbedded
GlobalLogic
Bromium
DORNERWORKS
BAE SYSTEMS
OpenXT
QUBES OS
A REASONABLY SECURE OPERATING SYSTEM
CRUCIBLE
XILINX
galois
BAE SYSTEMS
Xen
Project

Story 1:
Virtual Machine Introspection

A new way to protect against malware

Enablers: from xenaccess/xenprobes to LibVMI

Interesting research topic

Originally used for forensics (too intrusive for server virt)

VMI: enabling commercial applications

Hardware assisted VMI solves the intrusion problem

Collaboration between: Zentific, Citrix, Bitdefender, Intel and others

Traditional Cloud Security

Installed in-guest agents, e.g. anti-virus software, VM disk & memory scanner, network monitor, etc.

Can be disabled by rootkits and advanced persistent threats (APT)

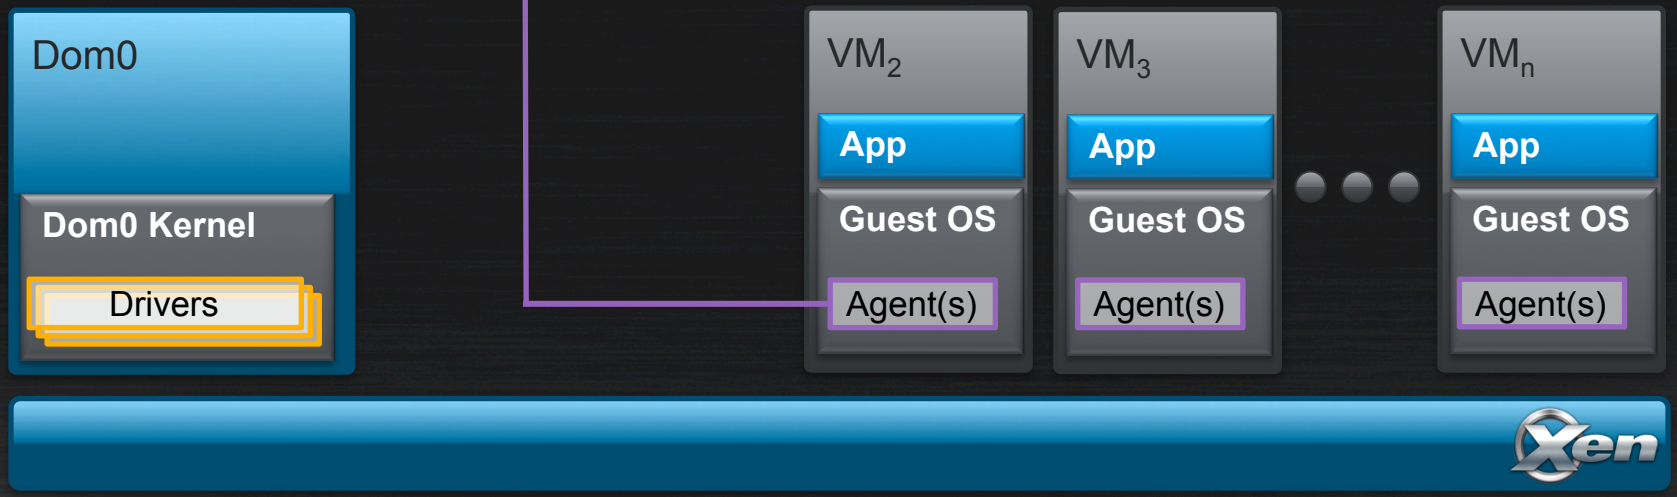

A new model for Cloud Security?

Uses HW extensions to monitor memory (e.g. Intel EPT) ➔ Low Intrusion

Register rules with Xen to trap on and inspect suspicious activities (e.g. execution of memory on the dynamic heap)

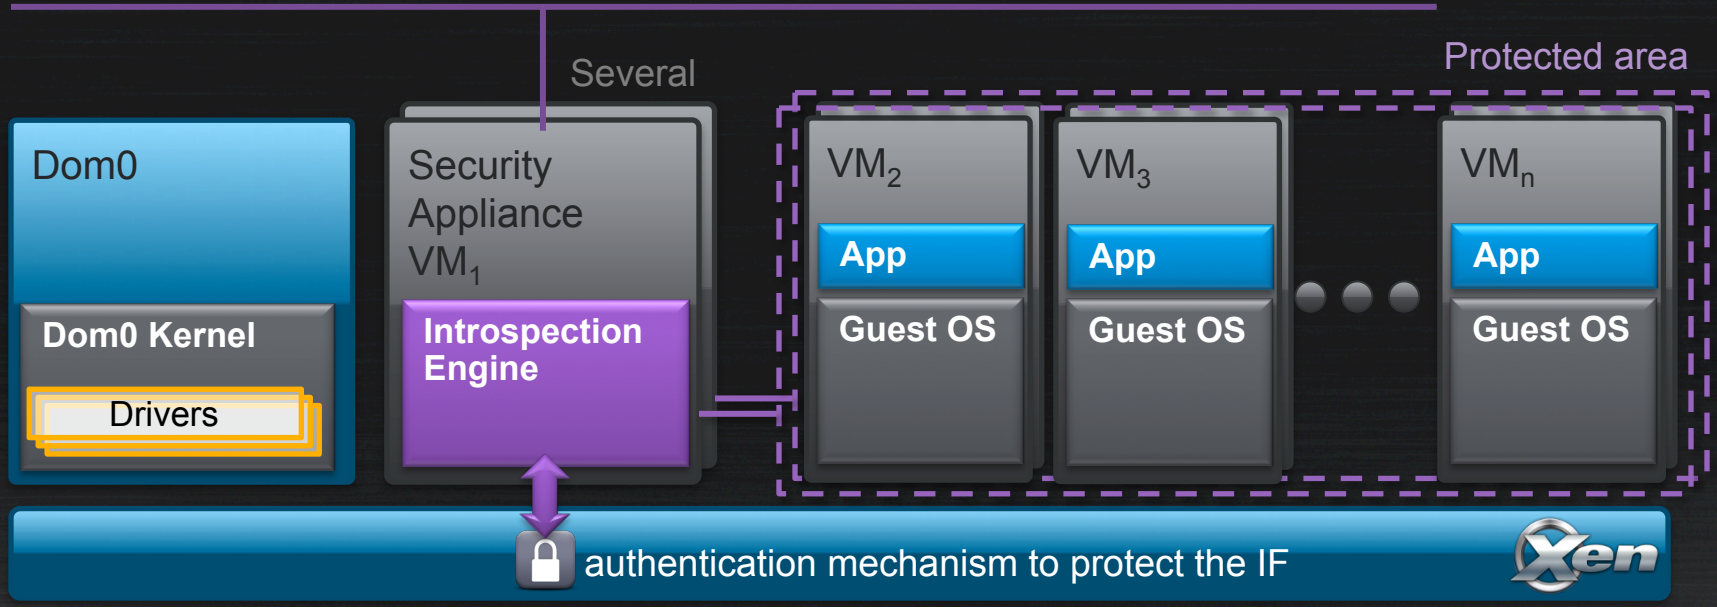

Protection against attack techniques

All malware need an attack technique to gain a foothold
Attack techniques exploit specific software bugs/vulnerability

The number of available attack techniques is small
Buffer Overflows, Heap Sprays, Code Injection, API Hooking, …

Because VMI protects against attack techniques
It can protect against entirely new malware

Verified to block these advanced attacks in real-time
APT28, Energetic Bear, DarkHotel, Epic Turla, Regin, ZeuS, Dyreza, EternalBlue[1]
… solely by relying on VMI

WannaCry/EternalBlue blocked in real installations

[1] businessinsights.bitdefender.com/hypervisor-introspection-defeated-enternalblue-a-priori

Protection against rootkits & APTs

Rootkits & Advanced Persistent Threats
Exploit 0-days in Operating Systems/System Software
Can disable agent based security solutions (mask their own existence)

VMI solutions operate from outside the VM
Thus, it cannot be disabled using traditional attack vectors

BUT:
VMI is not a replacement, for traditional security solutions
It is an extra tool that can be used to increase protection

If you want to know more …

Documentation

wiki.xenproject.org/wiki/Virtual_Machine_Introspection

Products

Bitdefender HVI
XenServer
www.bitdefender.com

Protection & Remedial
Monitoring & Admin

Citrix Ready

AIS Introvirt
XenServer
www.ainfosec.com

Zentific Zazen (June 17)
Xen & XenServer & …
www.zentific.com

Protection & Remedial
Monitoring & Admin
Forensics & Data gathering
Malware analysis

Pratap Sankar @ Flickr

Story 2:
Live Patching in Xen Project and XenServer

A tale of close collaboration within the Xen Project Community

Why did we develop Live Patching?

Cloud reboot affected AWS, Rackspace, IBM SoftLayer and others

Deploying security patches may require reboots; Inconveniences users

How did we fix this?

2015: Design with input from AWS, Alibaba, Citrix, Oracle and SUSE

2016: Xen 4.7 came with Live Patching for x86

2016: Xen 4.8 added extra x86 use-cases and ARM support

2017: XenServer 7.1 releases Live Patching in first commercial product

What is Live Patching?

Replacing compiled functions with new code, encoded in an ELF file called payload, while the hypervisor is running without impacting running guests.

```
const char *xen_extra_version(void)
{
return XEN_EXTRAVERSION;
}
```

```
push %rbp
mov %rsp,%rbp
lea 0x16698b(%rip),%rax
leaveq
retq
```

```
const char *xen_extra_version(void)
{
return “Hello World”;
}
```

```
push %rbp
mov %rsp,%rbp
lea 0x29333b(%rip),%rax
leaveq
Retq
```

Design: xenbits.xenproject.org/docs/unstable/misc/livepatch.html

Building Live Patches in Xen

Through livepatch-build-tools (based on kpatch-build)

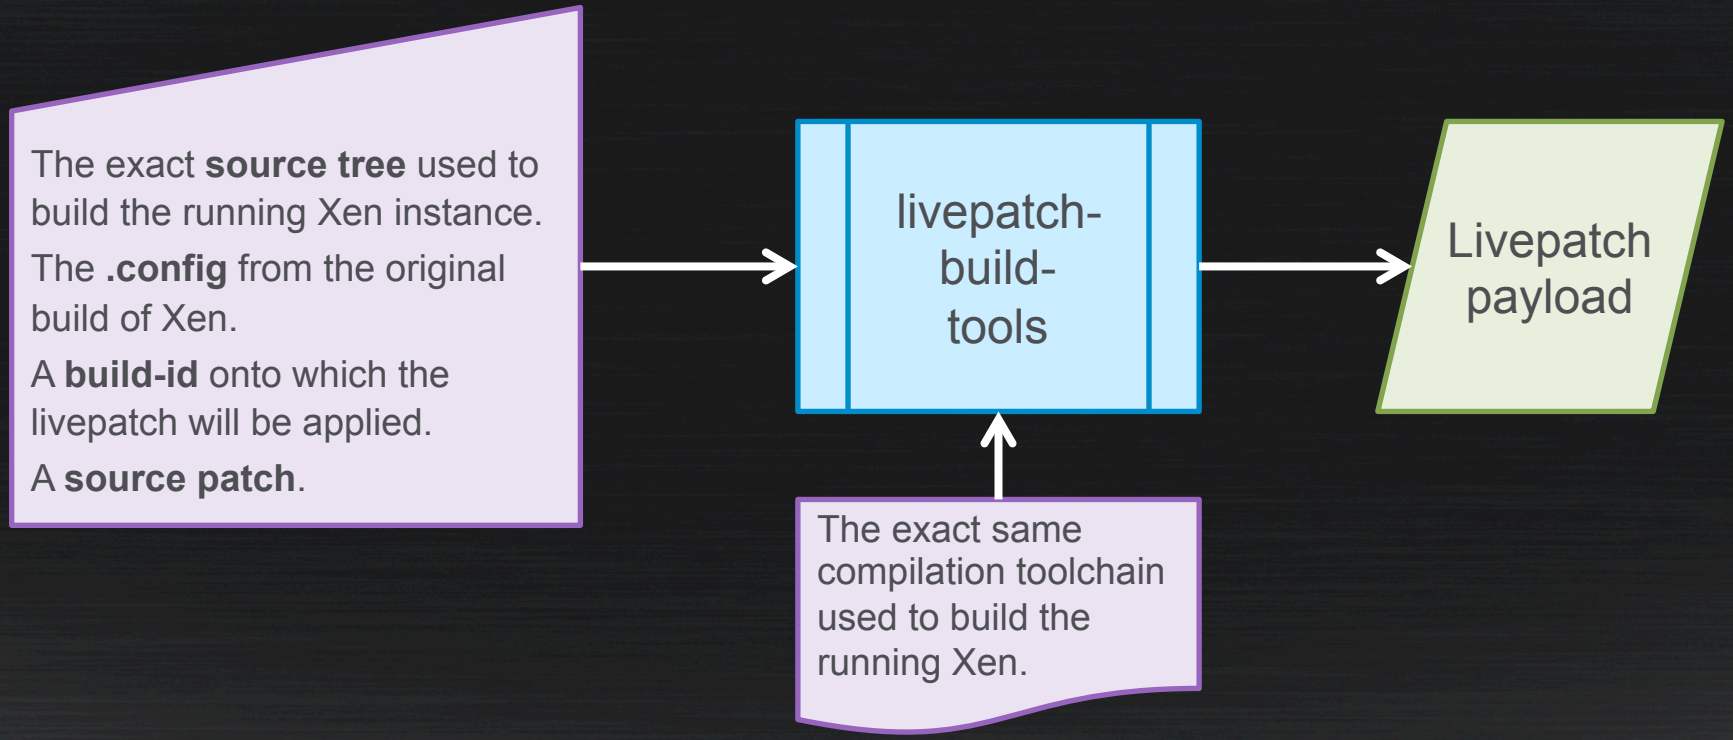

Applying Live Patches in Xen

Through xen-livepatch

Supports **stacking** of different payloads; payloads depend on **build-id**

Functionality:

list: lists loaded and applied live patches

upload: load & verify a live patch

unload: unload a live patch

apply: apply a live patch

revert: un-apply a live patch

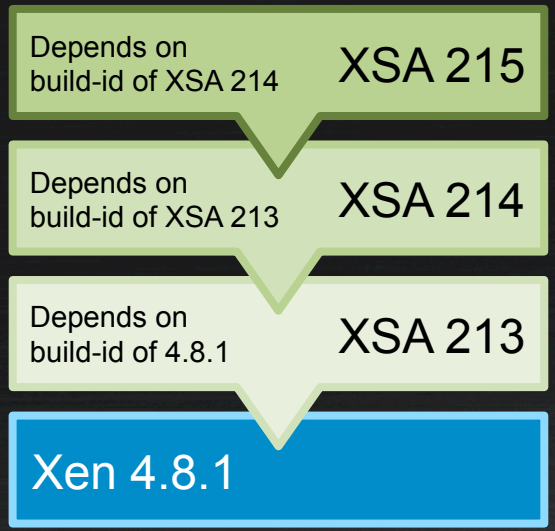

Live Patching in Xen & XenServer

Target	Technology	Function + Data	Data Structures	Inline patching
Dom0 & Guest Linux Kernel	Kernel Live Patching	✔	✗	✗
	kGraft (SUSE)	✔	✗	✗
	kPatch (RedHat)	✔	✔ via hooks	✗
	kSplice (Oracle)	✔	✔	✔
Hypervisor	Xen LivePatch	✔ Xen 4.7	✔ Xen 4.8 via hooks	✗ Future

XenServer LivePatch

- For Dom0 (CentOS)
- For Xen

Integrates different solutions into a single user experience

Live Patches in XenServer

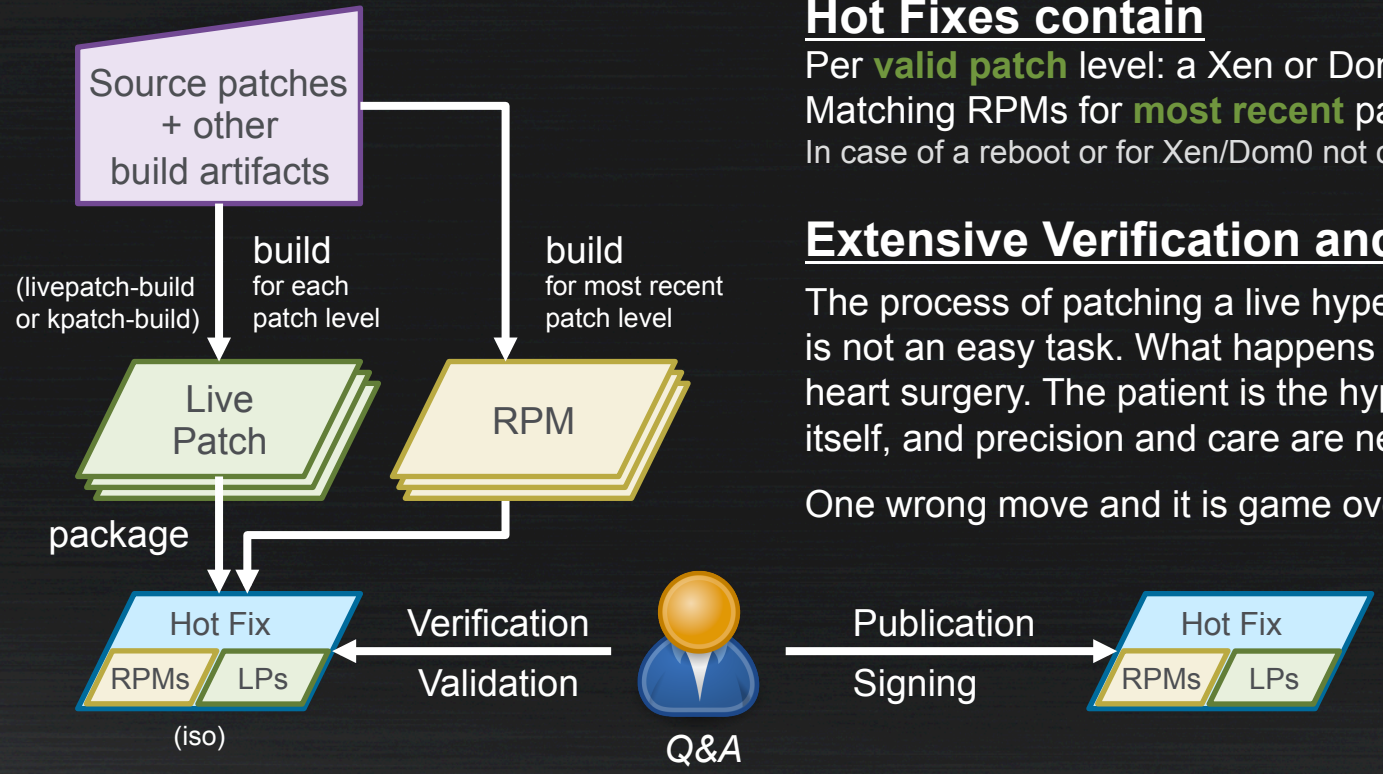

Hot Fixes contain

Per **valid patch** level: a Xen or Dom0 Live Patch
Matching RPMs for **most recent** patch level
In case of a reboot or for Xen/Dom0 not capable of Live Patching

Extensive Verification and Validation:

The process of patching a live hypervisor or kernel is not an easy task. What happens is a little bit like open heart surgery. The patient is the hypervisor and/or Dom0 itself, and precision and care are needed to get things right.

One wrong move and it is game over.

Live Patches in XenServer

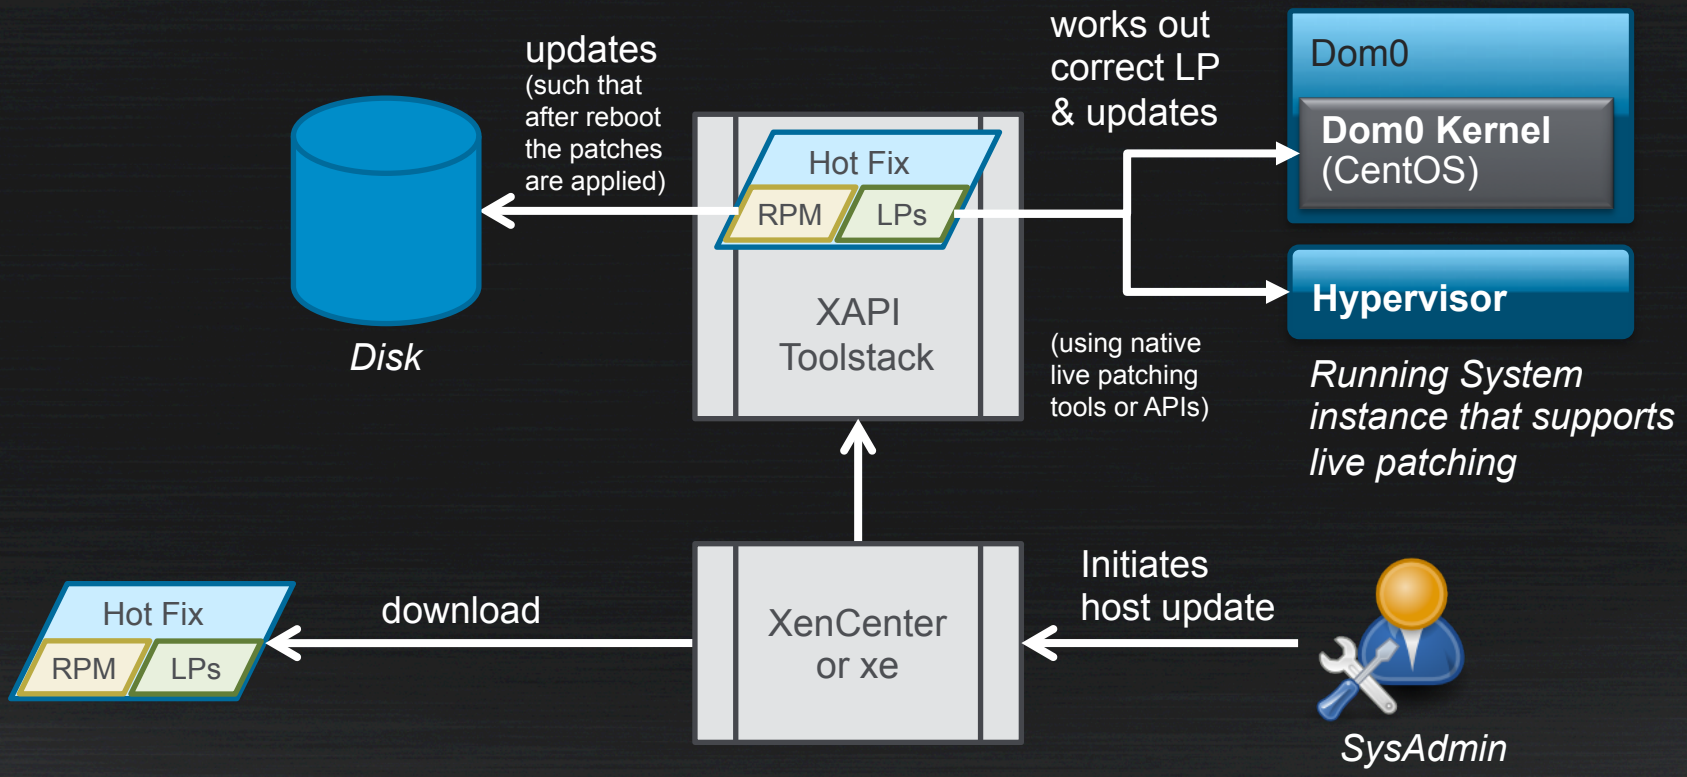

Demo: Live Patches in XenServer

xenbits.xenproject.org/people/larsk/LCC17 - Build LivePatch.mp4
xenbits.xenproject.org/people/larsk/LCC17 - Apply LivePatch.mov

Pratap Sankar @ Flickr

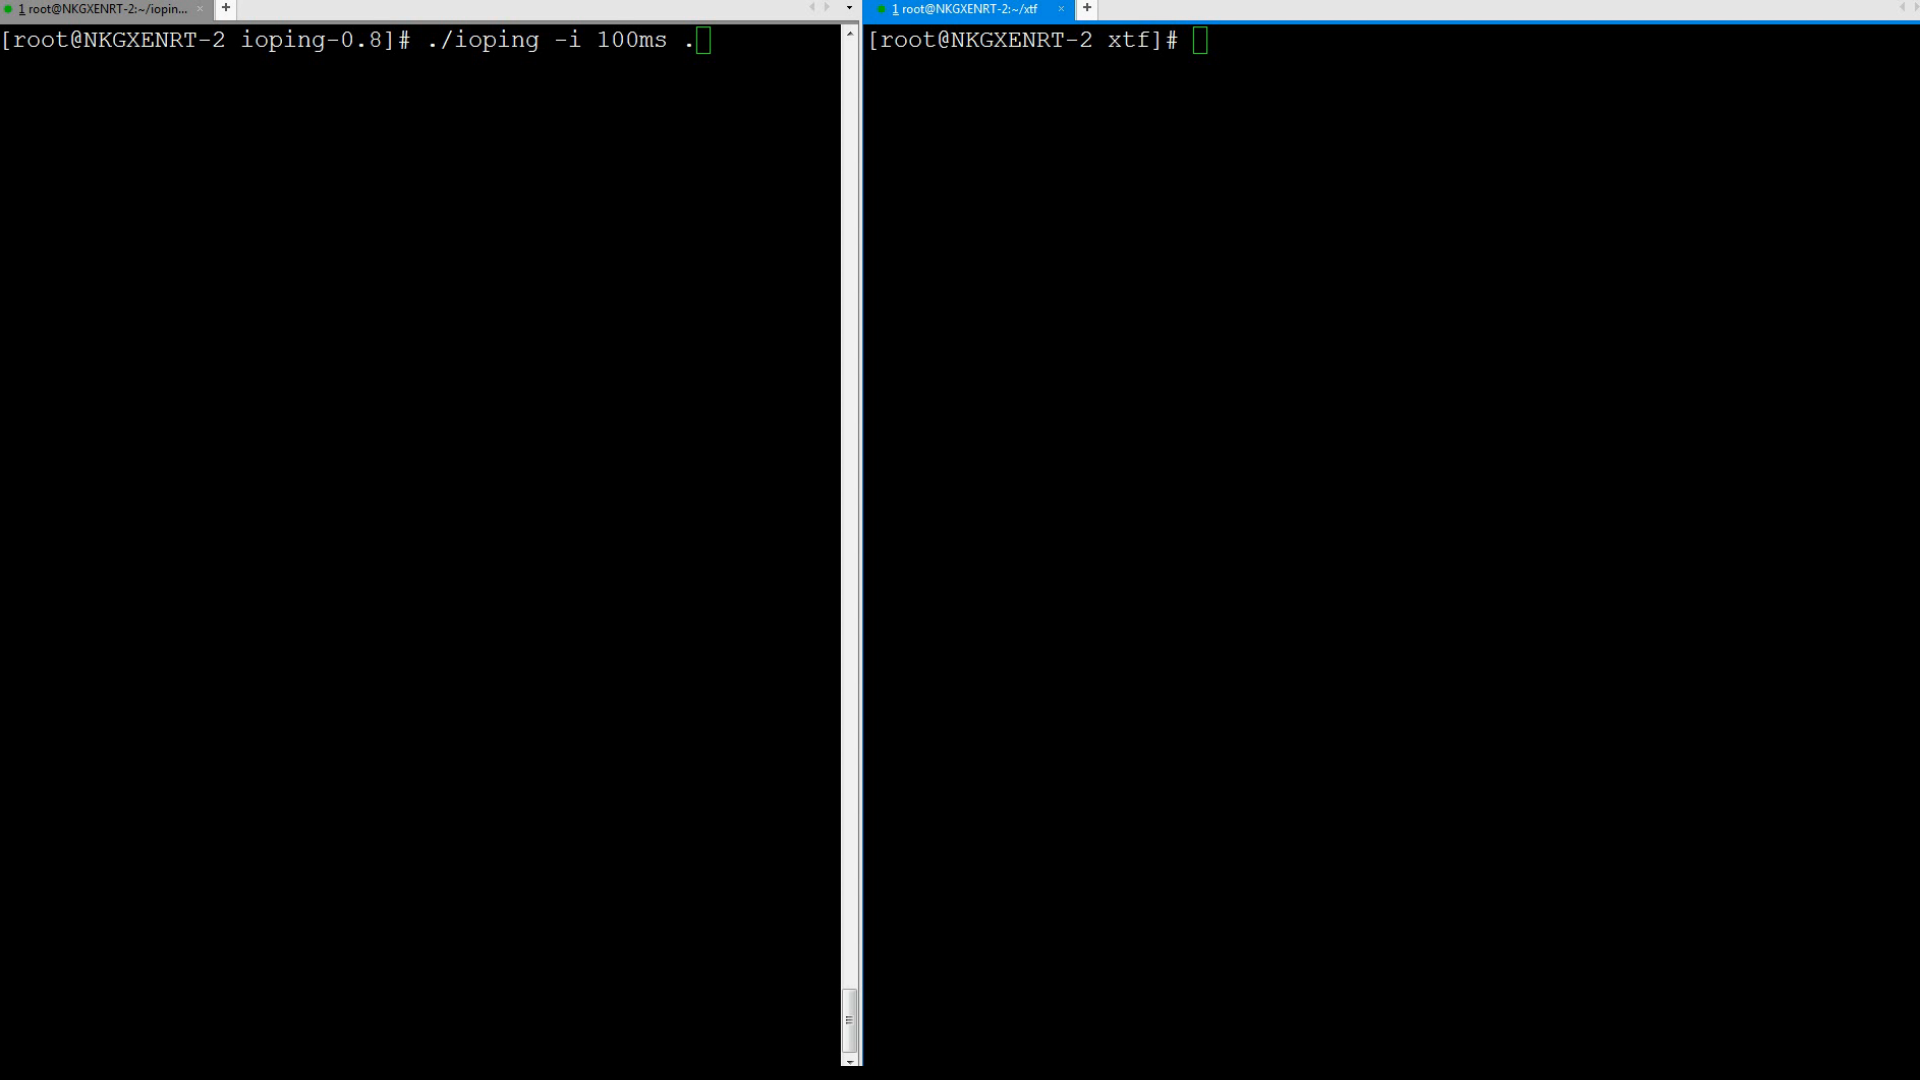
1 root@NKGXENRT-2:~/iopin...
[root@NKGXENRT-2 ioping-0.8]# ./ioping -i 100ms .
1 root@NKGXENRT-2:~/xtf
[root@NKGXENRT-2 xtf]#

Using Live Patching with XenServer makes Live Patching easy!

If you want to know more …

Xen Project LivePatch Specification & Status

xenbits.xenproject.org/docs/unstable/misc/livepatch.html
wiki.xenproject.org/wiki/LivePatch

Xen Project LivePatch Presentations & Videos

xenbits.xenproject.org/people/larsk/FOSDEM17-LivePatch.pdf (Short)
people/larsk/XPDS16-LivePatch.pdf (Long)

Xen Project LivePatch Videos

fosdem.org/2017/schedule/event/iaas_livepatxen/

XenServer

xenserver.org

Pratap Sankar @ Flickr

Story 3: Vulnerability Management in Xen Project

Industry-leading vulnerability management

Software bugs happen
Some will be security vulnerabilities

CVE's discovered

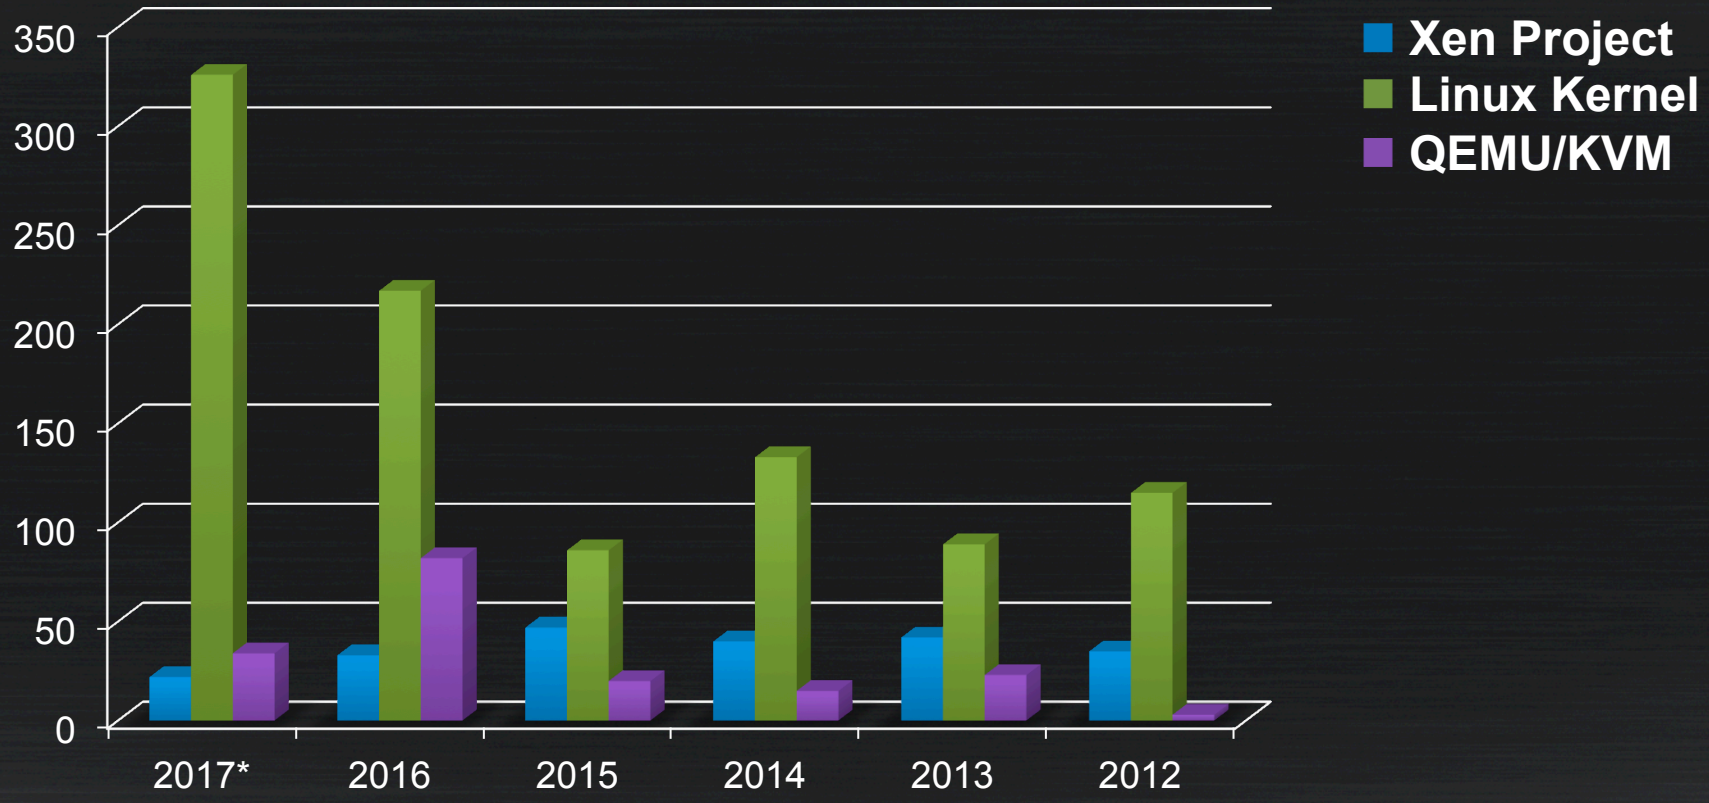

*) Data up to May 31st, 2017

Vulnerability data from cvedetails.com

Xen Project: Responsible Disclosure

xenproject.org/security-policy.html

Fixing Security Bugs:
Dedicated security team =
security experts from within
the Xen Project Community

R P

Security Team:
Triage
Creation of fix/patches
Can a Livepatch can be created?
No? If possible, re-write fix/patches
Validation of fix/patches
Assignment of CVE
Issue description and risk analysis

R: Vulnerability reported to security@xenproject.org
P: Vulnerability pre-disclosed on xen-security-issues@lists.xenproject.org

Xen Project: Responsible Disclosure

xenproject.org/security-policy.html

Fix their systems/software:
Eligible Xen Project Users
are informed **under embargo**
of the vulnerability

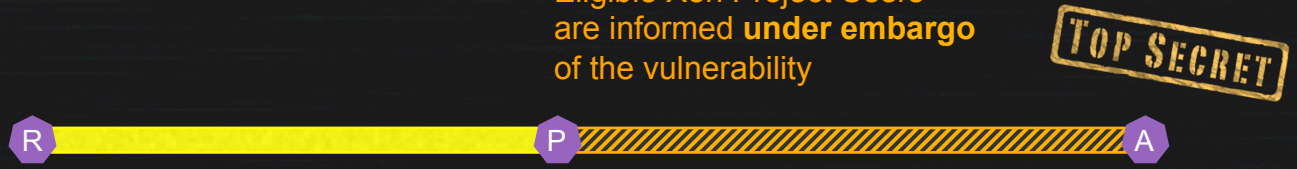

Eligible Users = Pre-disclosure list members:
Product Companies, Open Source & Commercial Distros (e.g. Huawei, Debian)
Service/Cloud Providers (e.g. Alibaba)
Large Private Downstream (e.g. Google)

Allowed to share information via
xen-security-issues- discuss@lists.xenproject.org

R: Vulnerability reported to security@xenproject.org
P: Vulnerability pre-disclosed on xen-security-issues@lists.xenproject.org
A: Vulnerability announced on xen-announce@lists.xenproject.org & xenbits.xen.org/xsa

Xen Project: Responsible Disclosure

xenproject.org/security-policy.html

General Publication:
Information about
vulnerability is made **public**

Everyone else:
Patches their systems either through security updates from distros/products or builds them from source.

Users of service/cloud providers will **not be** impacted

R: Vulnerability reported to security@xenproject.org
P: Vulnerability pre-disclosed on xen-security-issues@lists.xenproject.org
A: Vulnerability announced on xen-announce@lists.xenproject.org & xenbits.xen.org/xsa

Other Disclosure Models

Responsible Disclosure: fix critical systems/software **before** publication

Full Disclosure, post-fix: public disclosure **with** a fix

Full Disclosure, immediate (no-fix): public disclosure **without** a fix

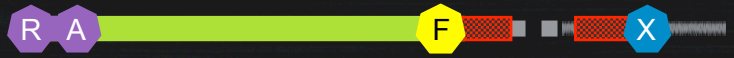

R: Vulnerability reported to security@...
P: Vulnerability pre-disclosed to eligible users
A: Vulnerability announced publicly
F: Fix available

Vulnerability Process Comparison

FOSS Project	Bug Severity [1]	Process Type	Days [2] (Responsible only)	Who? [3] (Responsible only)
Linux Kernel via OSS security distros [4] OSS security [5]	≥ Medium − Critical ≤ Low	Responsible Disclosure Full Disclosure, no-fix	14-19	D [6]
QEMU/KVM via OSS security distros [4] OSS security [5]	≥ Medium − Critical ≤ Low	Responsible Disclosure [7] Full Disclosure, no-fix	14-19	D [6]
OpenStack OSSA **OpenStack OSSN**	≥ Medium − Critical ≤ Low	Responsible Disclosure Full Disclosure, post-fix	3-5	D, S, P
Xen Hypervisor Includes Linux & QEMU vulnerabilities in supported Xen configurations	Low − Critical	Responsible Disclosure	14	D, S, P

1) Is the CVE severity used to handle vulnerabilities differently?
2) Days embargoed (information is secret)
3) D = Distros/Products, S = Public Service, P = Private Downstream

4) http://oss-security.openwall.org/wiki/mailing-lists/distros
5) http://www.openwall.com/lists/oss-security
6) No Chinese companies or distros on pre-disclosure list
7) Only handles x86 KVM bugs (no ARM or other bugs)

Summary …

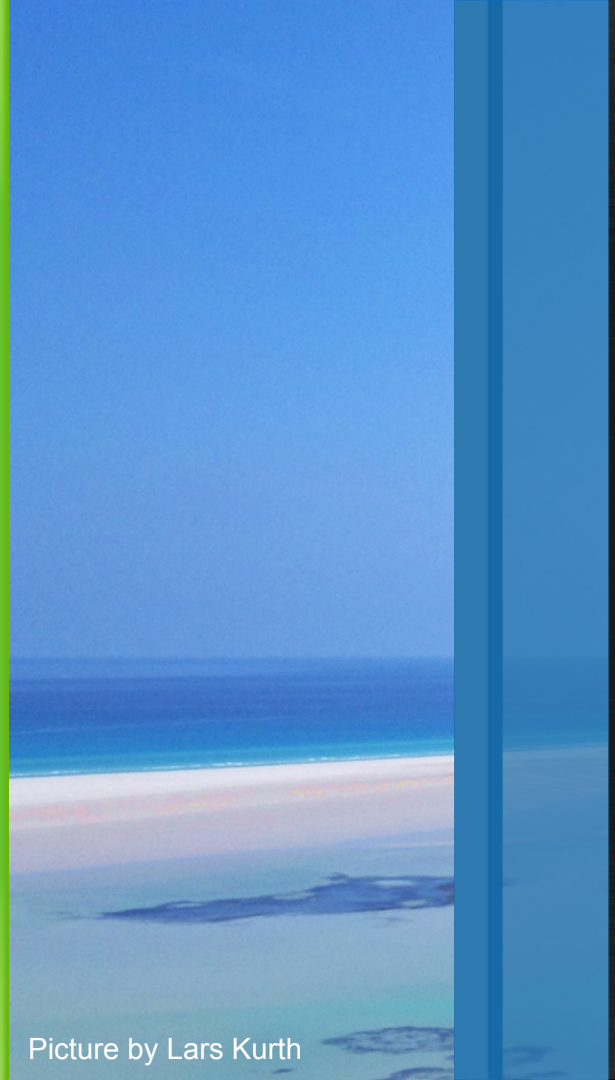

Picture by Lars Kurth

Only Hypervisor with VMI

Protection from new classes of malware
Several security companies working with XenServer

Live Patching

Disruption free application of vulnerabilities
Used by several cloud providers
Used best in commercial products, e.g. XenServer

Industry Leading Vulnerability Process

Includes QEMU and Kernel XSAs
Designed with input from Cloud Providers
Stable number of CVEs

BUT: the Xen Project cannot today distribute XSAs as Live Patches (the project delivers source code only)

Sys Admins

Extra protection = extra piece of mind

Questions

xenbits.xenproject.org/people/larsk

You can also contact Patrick Zhang (patrick.zhang@citrix.com)
after the presentation

Picture by Lars Kurth