



AWS EC2 Virtualization: Introducing Nitro

Chris Schlaeger

Director Kernel and Operating Systems

What is the Nitro Project?

What is a Hypervisor?

The evolution of the Nitro Project

Compatibility to Xen

FAQ

What's next?

What is Nitro?

From the C5 launch:

Q. What is the new hypervisor for Amazon EC2?

The new hypervisor for Amazon EC2, introduced with the launch of C5 instances, is a component that primarily provides CPU and memory isolation for C5 instances. VPC networking and EBS storage resources are implemented by dedicated hardware components that are part of all current generation EC2 instance families. It is built on core Linux Kernel-based Virtual Machine (KVM) technology, but does not include general purpose operating system components.

What is Nitro?

From the C5 launch:

Q. What is the new hypervisor for Amazon EC2?

The new hypervisor for Amazon EC2, introduced with the launch of C5 instances, is a component that primarily provides CPU and memory isolation for C5 instances. VPC networking and EBS storage resources are implemented by dedicated hardware components that are part of all current generation EC2 instance families. It is built on core Linux Kernel-based Virtual Machine (KVM) technology, but does not include general purpose operating system components.

The Nitro Hypervisor is the “new hypervisor,” but more than just a hypervisor

What is a Hypervisor?

A hypervisor consists of:

- Virtual Machine Monitor
- Many device models (10 to 100s)
- Scheduler, memory manager, etc.

This was state of the art in 1974

Not all of the assumptions held true though...

From 1974 to 2006

- Early Intel processors did not trap
- The Xen project found a clever solution
- Paravirtualization modifies the OS to trap
- Hypercalls directly invoke the VMM
- EC2 launched using Xen Paravirtualization

```
<_start>:
  e9 59 e1 17 00    jmpq  ffff82d08037e15e
  0f 1f 00          nopl  (%rax)

<multiboot1_header_start>:
  02 b0 ad 1b 03 00  add  0x31bad(%rax),%dh
  00 00            add  %al,(%rax)
  fb              sti
  4f 52           rex.WRXB push %r10
  e4 0f          HYPERCALL io_in

<multiboot1_header_end>:
  0f 1f 40 00      nopl  0x0(%rax)

<multiboot2_header_start>:
  d6              (bad)
  50              push  %rax
  52              push  %rdx
  e8 00 00 00 00  callq ffff82d080200020
  88 00          mov  %al,(%rax)
```

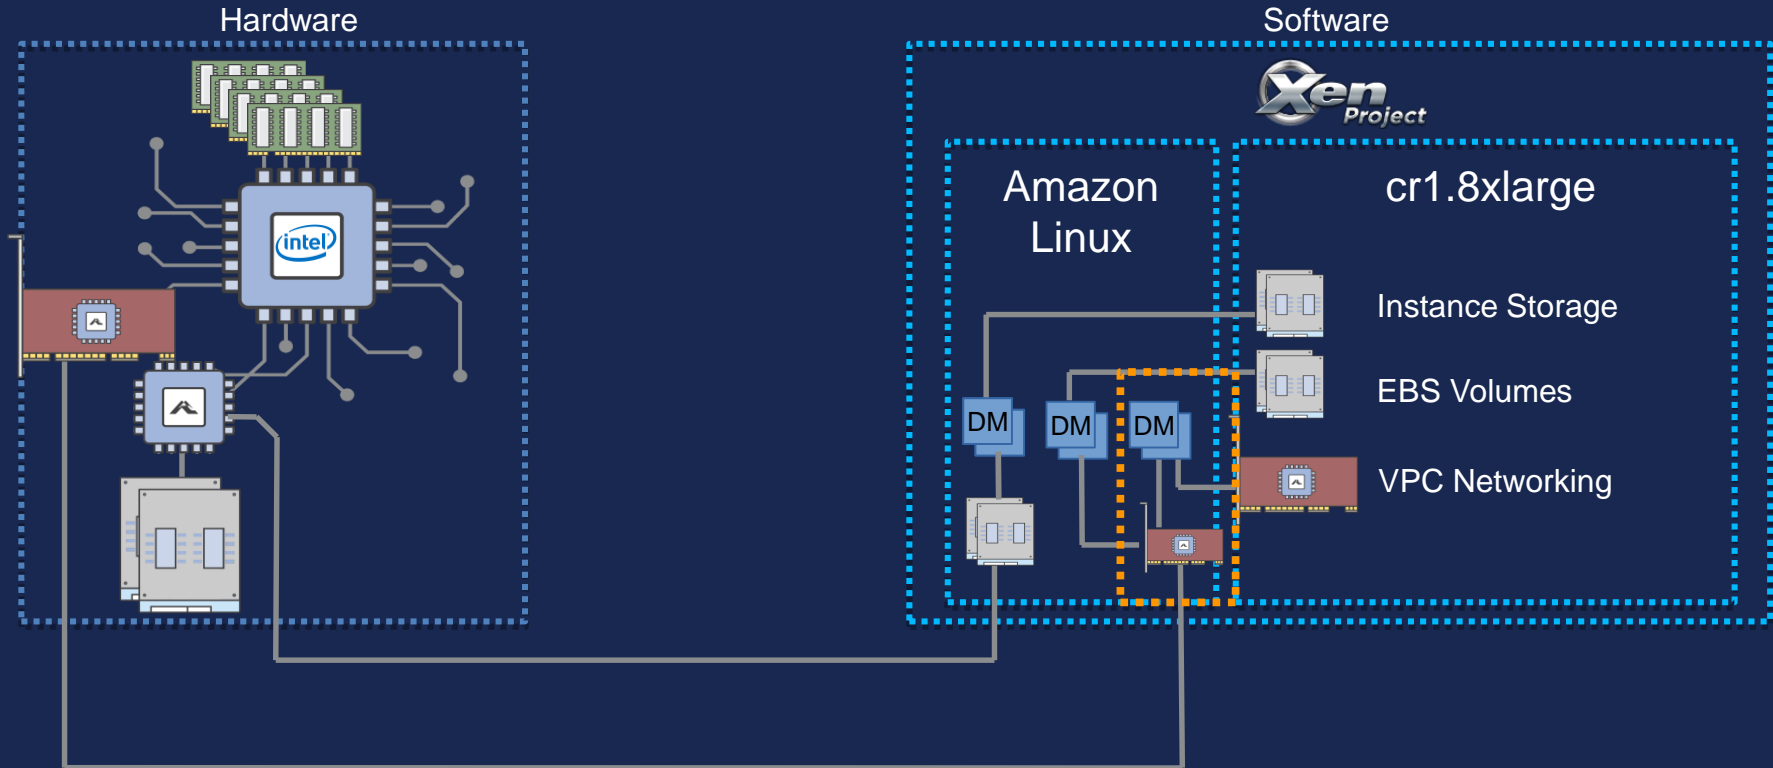
Evolution of the Nitro System

Circa 2012

- Can we do better than the software-only hypervisor architecture?
- Device models compete for CPU and system resources, jitter is hard to avoid.
- Can we decompose the hypervisor and shuffle components around?
- Let's begin our journey with the state of the art instance type from 2012.

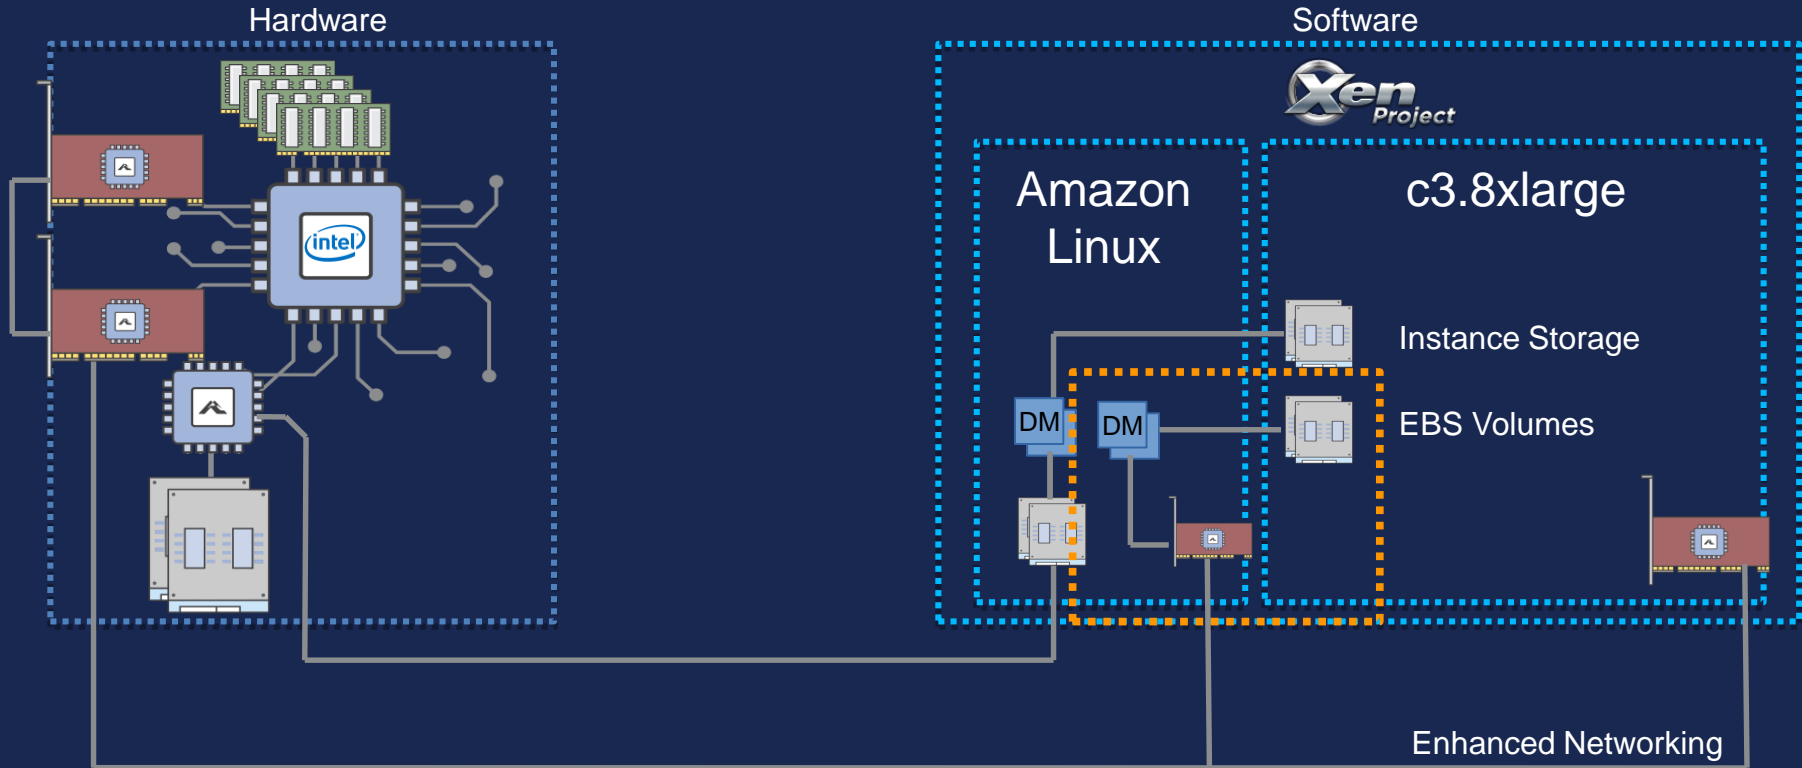
CR1 (no Nitro)

Jan 2013




C3 (early Nitro)

Nov 2013



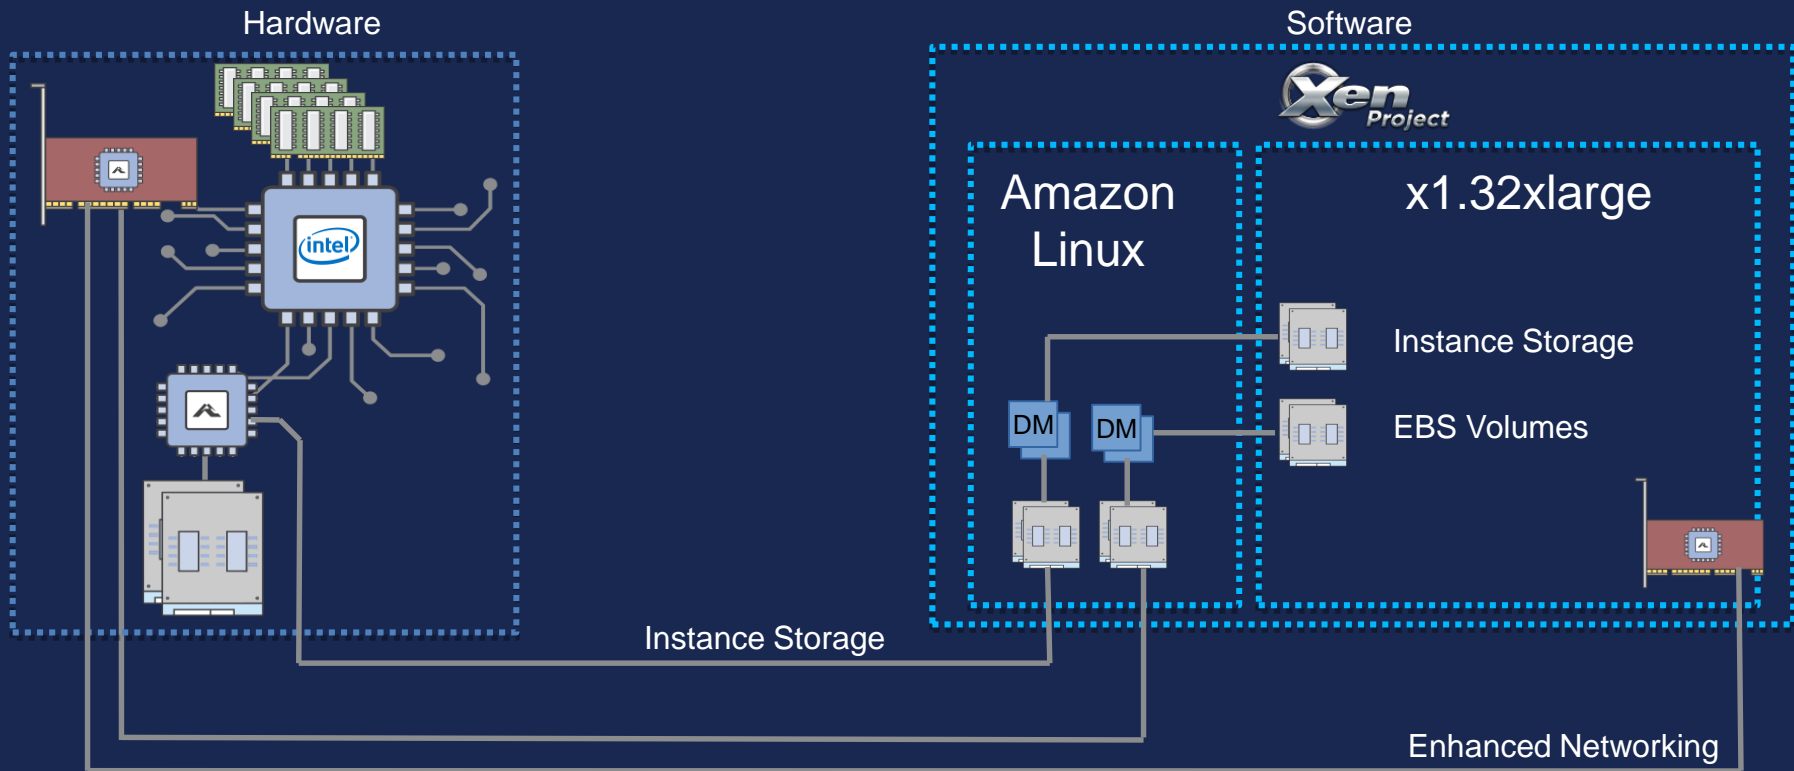
C4

Jan 2015



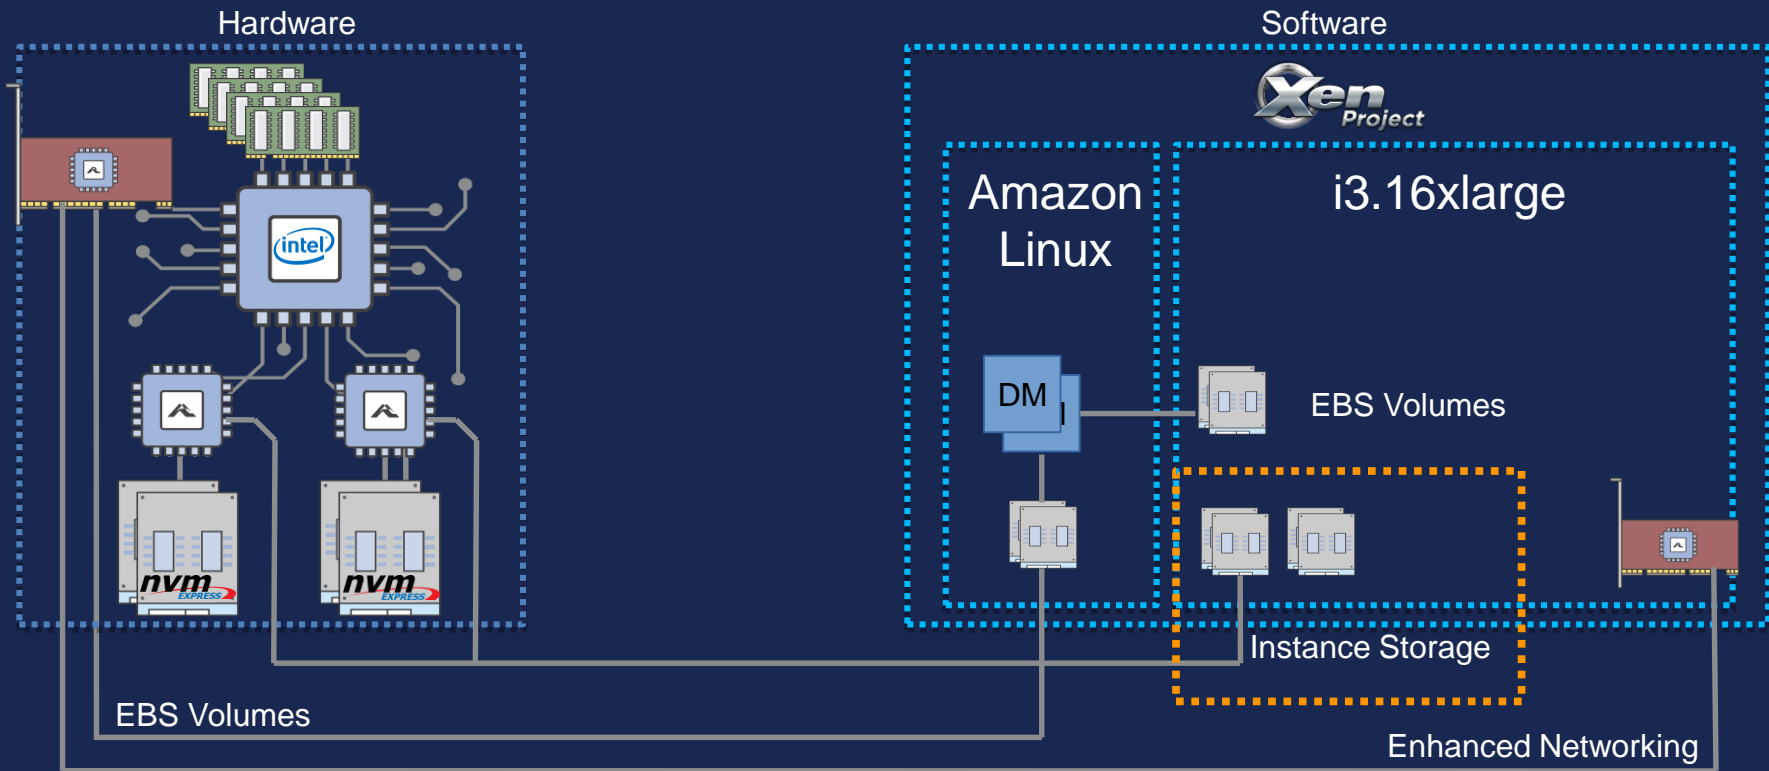
X1

May 2016



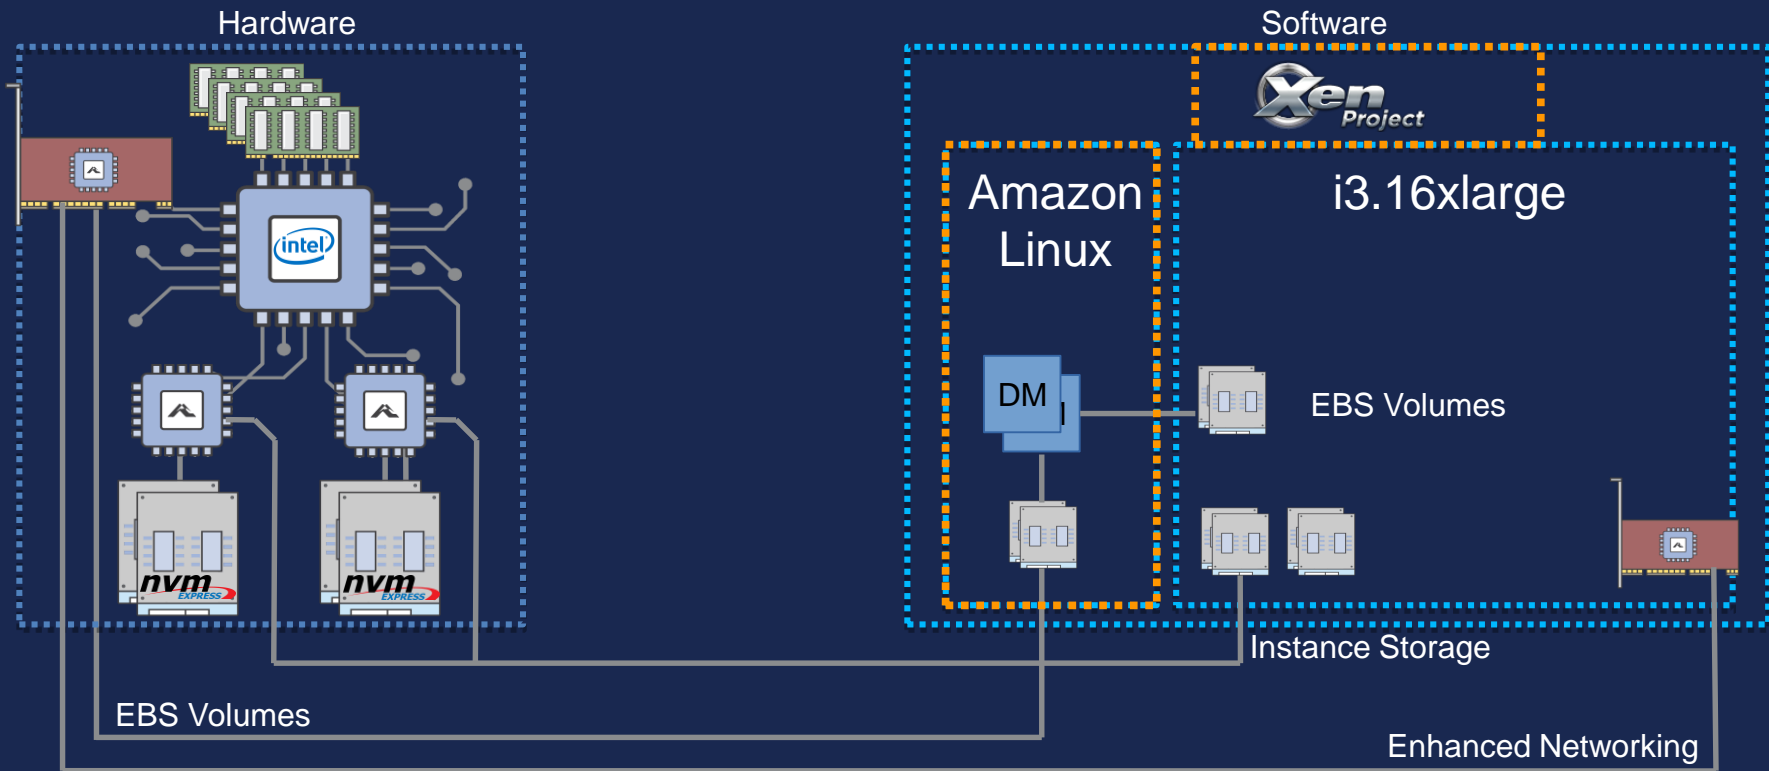
13

Feb 2017




13

Feb 2017




C5

Nov 2017



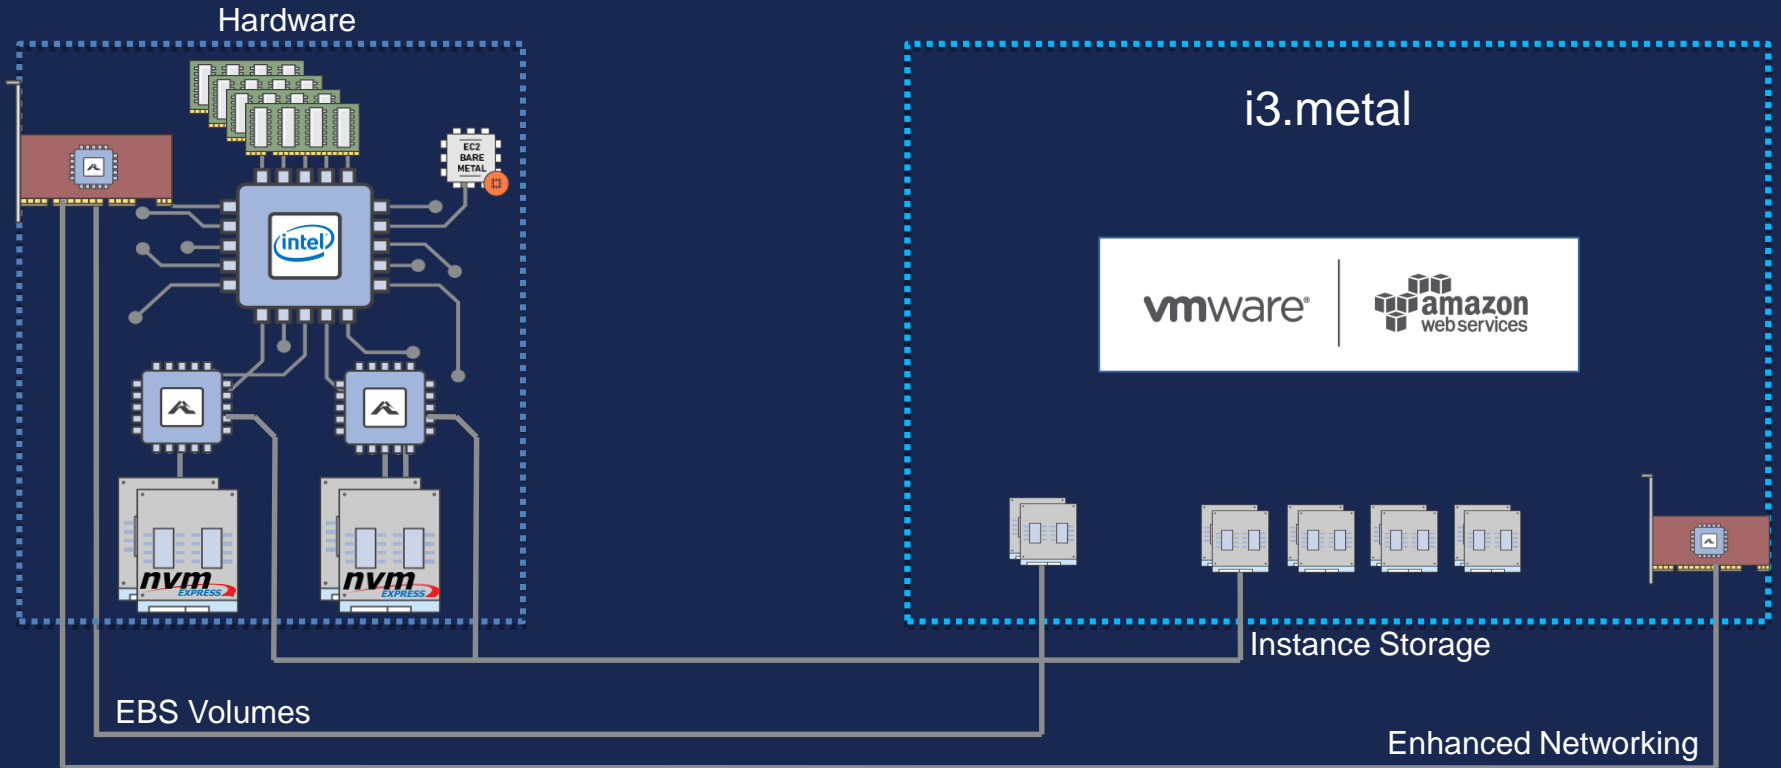
EC2 Bare Metal

Nov 2017



VMware on AWS

Aug 2017



The Nitro System

- Nitro Hypervisor
 - Lightweight hypervisor
- Nitro Card
 - Storage
 - Networking
 - Management
 - Monitoring
 - Security
- Nitro Security Chip
 - Integrated into the motherboard

FAQs

1) Will my existing AMIs work on Nitro-based instances?

Yes. Most ENA capable AMIs have the necessary drivers.

FAQs

2) Will applications need to be modified?

Most of the time, no. Some applications have relied on undocumented behavior to detect they are running within EC2 and they may require adjustment.

FAQs

3) Will all new instance types be based on the Nitro System?

In the fullness of time, we expect most (if not all) new instance types to be Nitro-based. We have no plans to convert existing instance types to Nitro and expect to continue to launch Xen based instance types where appropriate.

What's next?

Thank you!