

Amazon Nitro

Presented by Yiyang Zhang
Slides mainly from Vansh Singh and Kunlin Yang

Administrivia

- Quiz 2 at the end of this lecture
- Project progress report graded


Agenda

- Why Nitro
- Nitro Overview and History
- Nitro Cards
- Nitro Security Chip
- Nitro Hypervisor
- Other Infra Trends


Why Nitro?

Before Nitro - AWS EC2 CR1 (2013):

- Traditional software-based virtualization
- Relatively simple virtual machine monitor
- Access to device models through privileged OS



Structure of XEN



Why Nitro?

Drawbacks of pre-Nitro virtualization system

- Reserve CPU cores for Dom0 acceleration
- Device models compete for CPU and system resources
- Dom0 OS is too big and complicated
- and a potential security vulnerability (more later this quarter)

How can we ensure precious CPU cycles and memory go to user applications?
and the hypervisor is lightweight and fast?

Nitro Overview


What is Nitro?

- AWS's virtualization platform for next-gen EC2 instances (C3-C5)
- A combination of dedicated hardware and lightweight hypervisor

Benefits


- Better performance and price
- Faster innovation
- Enhanced security

EC2 “instance” host architecture



CR1 (no Nitro)


Jan 2013




C3 (early Nitro)

Nov 2013


Hardware



Software




2012 EC2 “instance” host architecture




C3 (early Nitro)

Nov 2013

Hardware



Software




Enhanced Networking

C4


Jan 2015



Hardware



Software




EBS Volumes

Enhanced Networking


C4

Jan 2015

Hardware




Software




X1


May 2016




X1


May 2016






2013 EC2 “instance” host architecture






C5

Nov 2017



2017 Introducing Nitro Architecture



Nitro Hypervisor



KVM-based hypervisor
with custom MM and
small userspace




Only executes on behalf
of instance, quiescent.

With Nitro, the hypervisor can be fast and simple


EC2 Bare Metal

Nov 2017




2018 Nitro enabling bare metal instances

- 2018 Nitro enabling bare metal instances




Nitro in three parts

Nitro Cards



VPC Networking
Amazon Elastic Block Store
(Amazon EBS)
Instance Storage
System Controller

Nitro Security Chip



Integrated into motherboard
Protects hardware resources
Hardware Root of Trust

Nitro Hypervisor



Lightweight Hypervisor
Memory and CPU allocation
Bare Metal-like performance

Nitro Cards



ENA PCIe Controller

VPC Data Plane



NVMe PCIe Controller

EBS Data Plane



NVMe PCIe Controller


Transparent Encryption



Instance
Storage

System Control


Root of Trust



Nitro
Control
|



Nitro Card for VPC



ENA Controller


Drivers available for all major operating systems
Independent of fabric

VPC Data Plane

Encapsulation
Security Groups
Limiters
Routing



Nitro Card for EBS



NVMe Controller

Standard drivers broadly available


EBS Data Plane

Encryption support

NVM to remote storage protocol



Nitro Card for Instance Storage



Instance
Storage

NVMe Controller

Standard drivers broadly available

Instance Storage Data Plane


Transparent Encryption

Limiters

Drive monitoring



Nitro Card Controller



Nitro
Controller

System Control

- Provides passive API endpoint
- Coordinates all other Nitro Cards
- Coordinates with Nitro Hypervisor
- Coordinates with Nitro Security Chip


Hardware Root of Trust

- Provides measurement and attestation

Nitro Security Chip




Custom microcontroller
that traps all I/O to non-
volatile storage




Controllable from the
Nitro Controller to hold
system boot

Provides a simple, hardware-based root of trust

Nitro architecture



Nitro architecture




Better Performance and Price

	AWS Instance	vCPU	RAM (GiB)	CPU:RAM	Price (\$/Hour)
W/ Nitro	c5.large	2	4	1:2	0.0085
W/o Nitro	c4.large	2	3.75	1:1.875	0.1
W/ Nitro	c5.4xlarge	16	32	1:2	0.68
W/o Nitro	c4.4xlarge	16	30	1:1.875	0.796
W/ Nitro	c5.24xlarge	96	192	1:2	4.08
W/o Nitro	c4.8xlarge	36	60	1:1.875	1.591

Source: <https://aws.amazon.com/ec2/pricing/on-demand/>

Nitro performance



In Sum

No Dom0 in Nitro hypervisor—greater simplicity and safety


No SSH or other interactive modes anywhere—no direct human access

All access via 100% AuthN/AuthZ APIs with logging/auditing

—no APIs for memory access

Only the Nitro controller has access to the physical Amazon EC2 network; the mainboard does not

End-to-end Nitro system is developed, deployed, and managed by DevSecOps process



In Sum ...

Big benefits from applying microservice concepts to hardware and full-system design

- Unneeded functionality not present

- What remains is better-defined and easier to reason about: build/test/validate


Stronger single root of trust and greater separation of concerns (and code, and teams) along every dimension

Nitro building blocks will continue to be

applied in Amazon EC2 and beyond


- Firecracker, Outposts, etc.

Lots of security value already, and there's more that we can do!




Other Infrastructure Trends

What AWS customers using today :



Amazon EC2 trn1 instance



trn1/trn1n

Powered by AWS Trainium



Trn1

MATH
ENGINE
FREQUENCY

3 GHz

BF16/FP16

3.4 PFLOPS

TF32

3.4 PFLOPS

FP32

840 TFLOPS

AGGREGATE
ACCELERATOR
MEMORY

512 GB

PEAK MEMORY
BANDWIDTH

13.1 TB/sec


NEURONLINK
BANDWIDTH
BETWEEN CHIPS

768 GB/sec

NETWORK
CONNECTIVITY


800 Gbps EFA
1600 Gbps EFA

Data Center of the Future



The Infrastructure Processing Unit

Providing new data center value




Highly intelligent **infrastructure acceleration**

System-level **security, control, and isolation**

Common software frameworks

HW and SW **programmable**, built to customer needs

OCTEON 10 innovations




- 5nm TSMC process
 - Enables fanless designs
- First inline DPU ML Engine
- Hardware VPP acceleration
- Inline crypto processor
- Arm Neoverse N2 cores
 - Highest SPECint in industry
- PCIe 5.0, DDR5 support
- Integrated with 16x 50GE switch
- 56G SerDes

Programmable Switches

Barefoot Tofino 2-powered 12.8 Tbps 32xQSFP56-DD System





Fixed Pipeline: features and table-sizes are baked in at design time

You declare which headers are recognized



Programmable Pipeline: all stages identical, customer-defined match-action logic