AWS re:Inforce

JUNE 13 - 14, 2023 | ANAHEIM, CA

DAP401

Security design of the AWS Nitro System

J.D. Bean

Principal Security Architect, EC2
AWS

The Ship of Theseus

AWS Nitro System

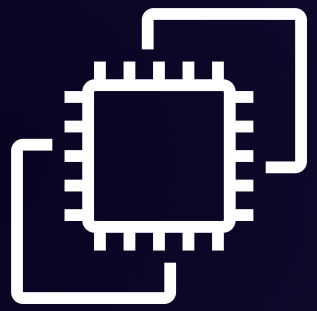

AWS Nitro System

Launched in November 2017

In development since 2012

Purpose-built hardware/software

Hypervisor built for AWS

Five generations of custom chips

All instances launched since 2018 use the AWS Nitro System

AWS Nitro System

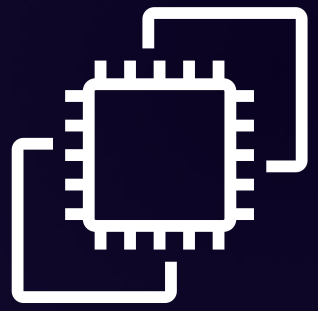

AWS Nitro System

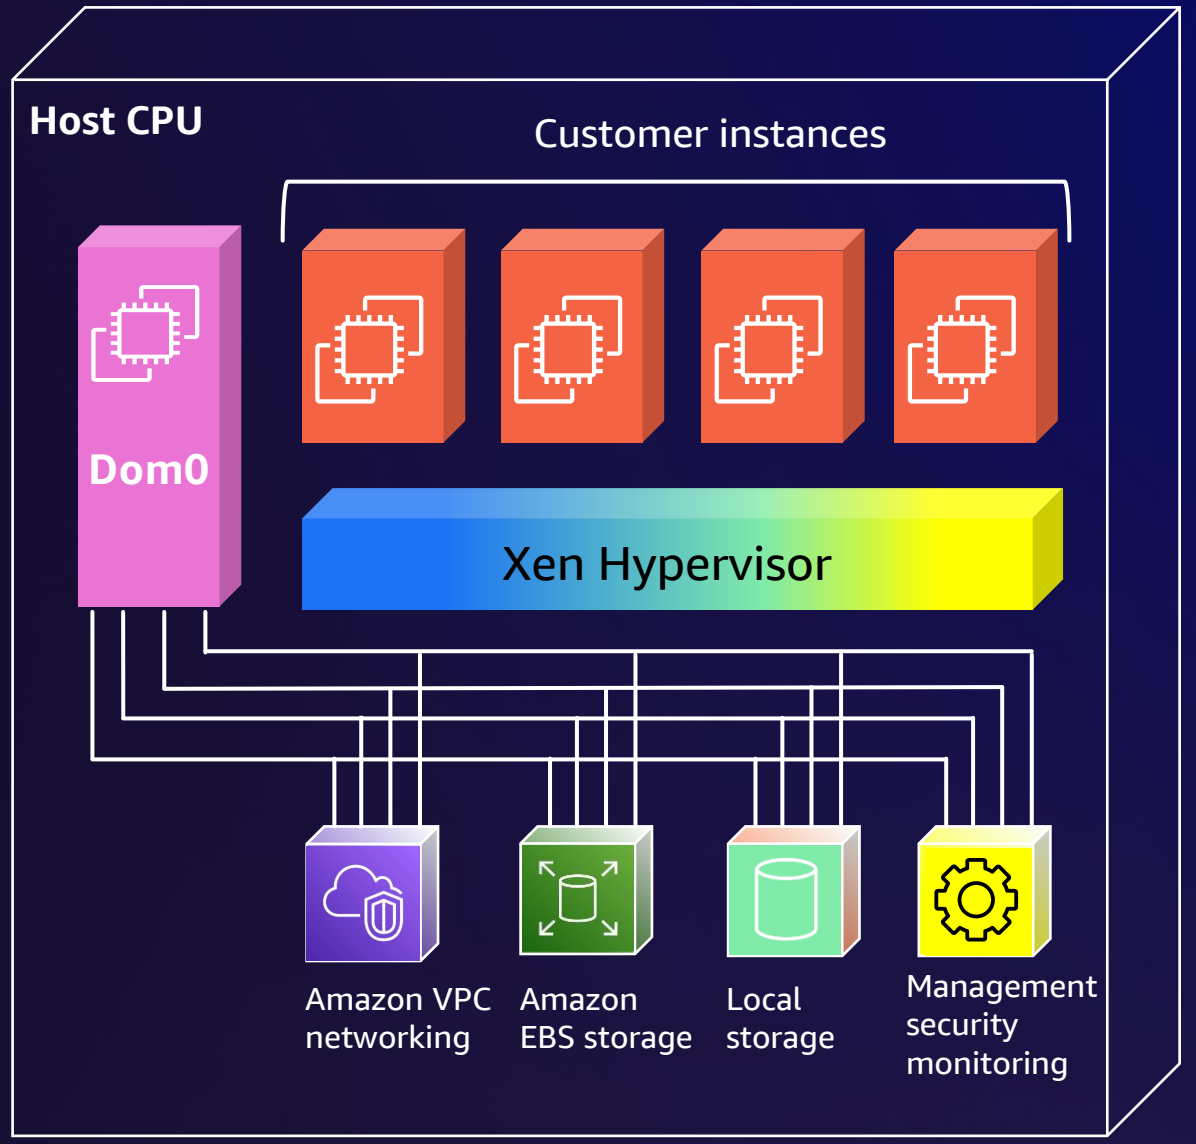

AWS Nitro System

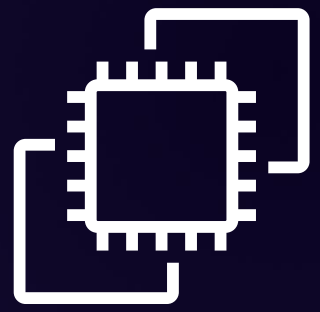

AWS Nitro System

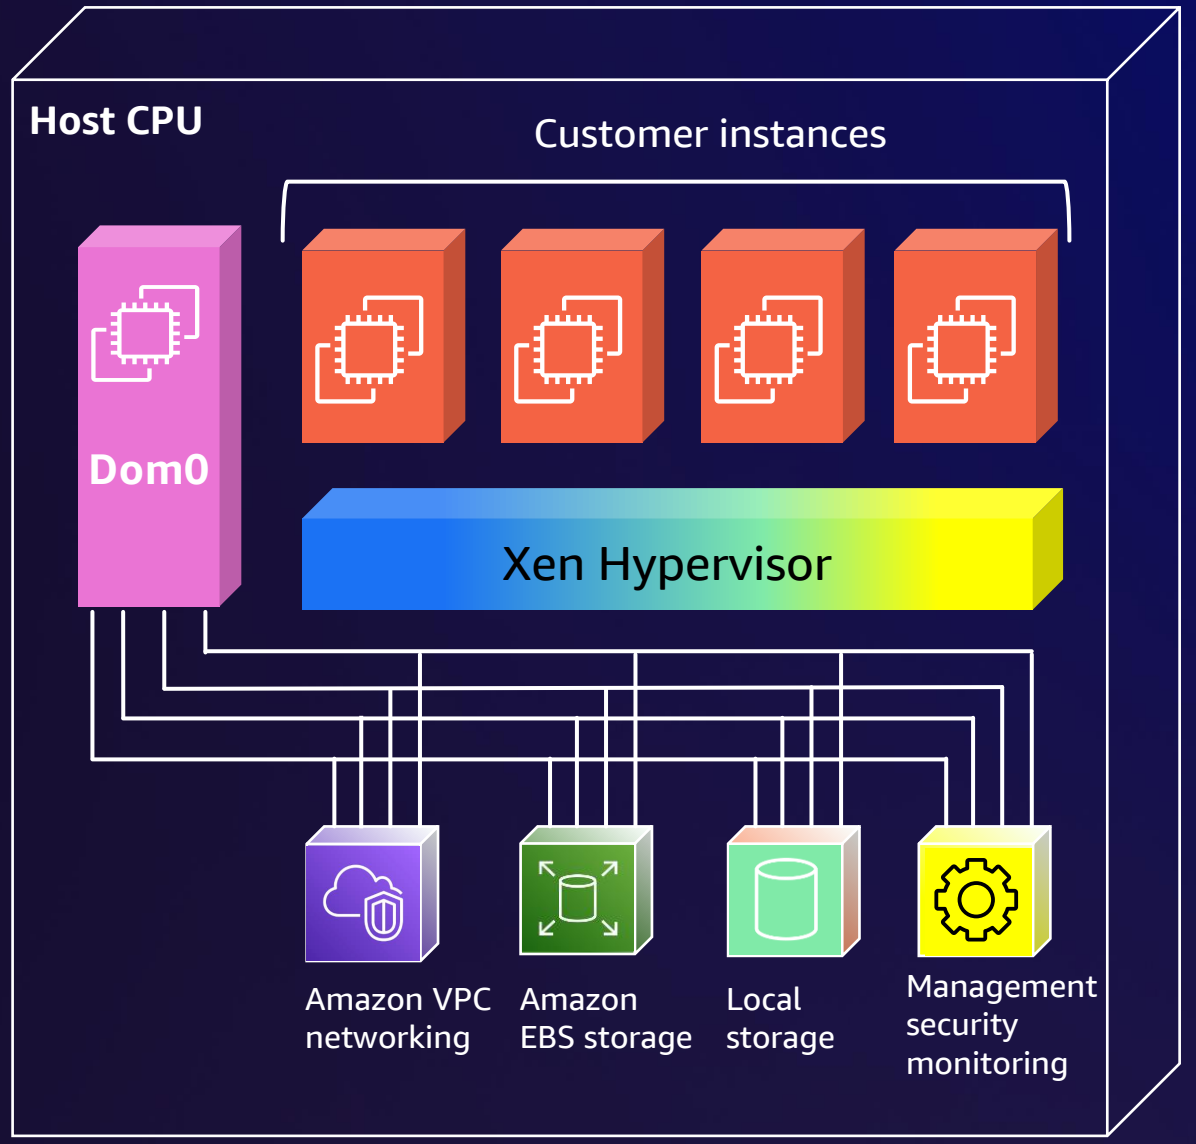

AWS Nitro System

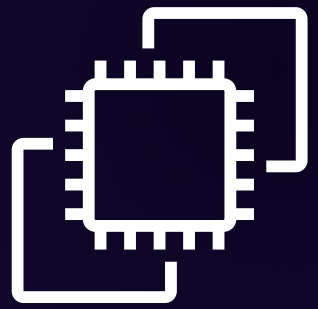

AWS Nitro System

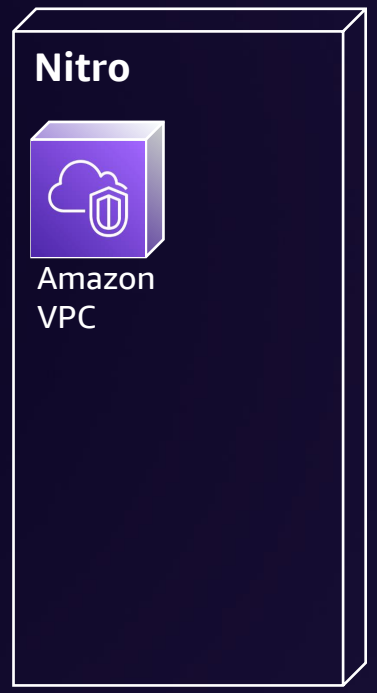

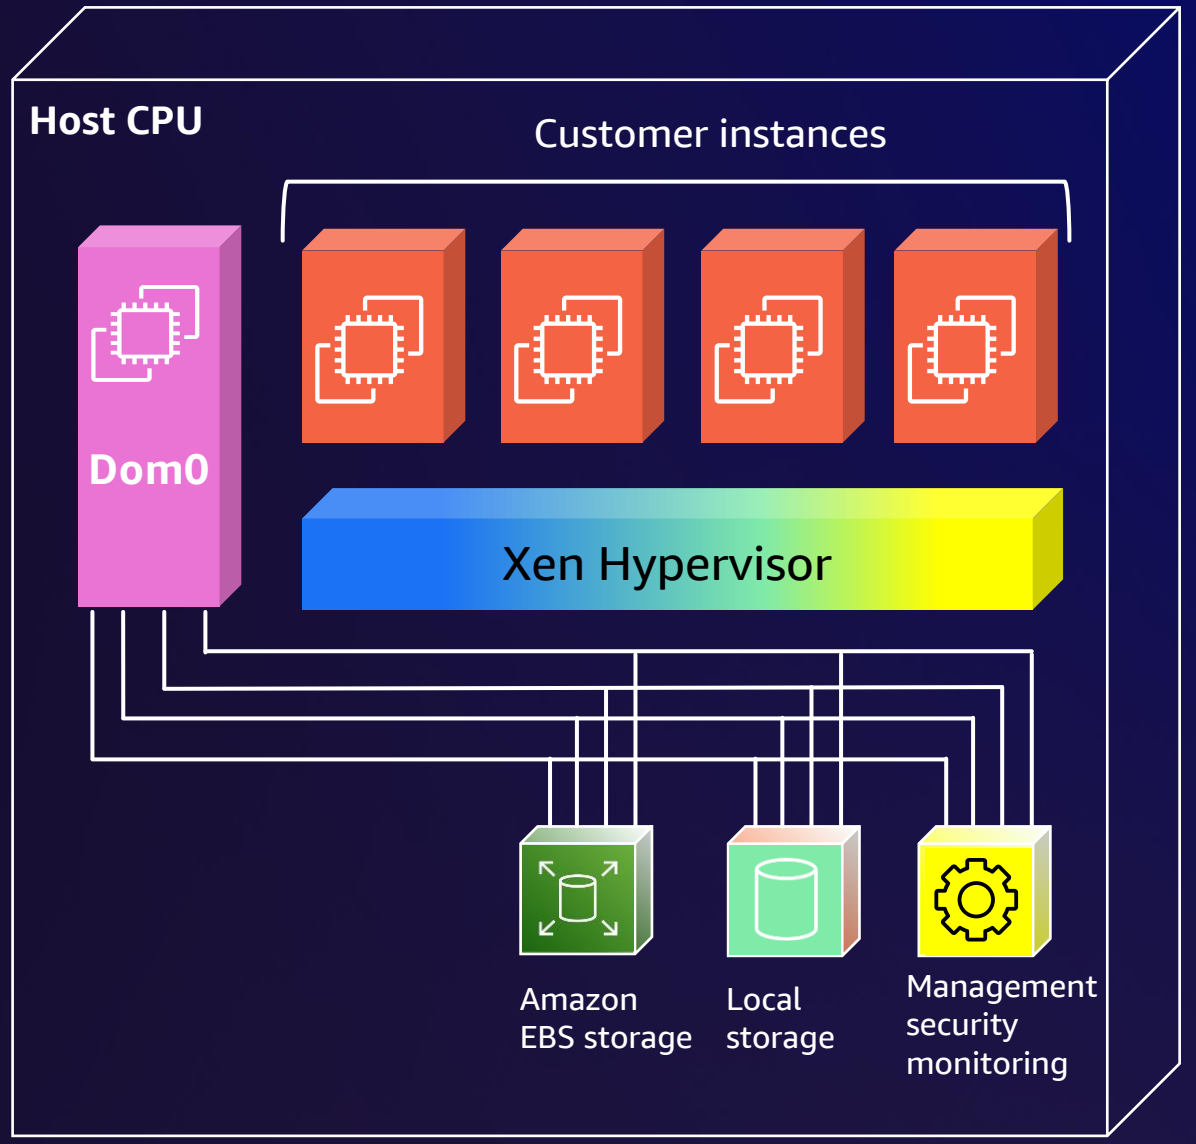

AWS Nitro System

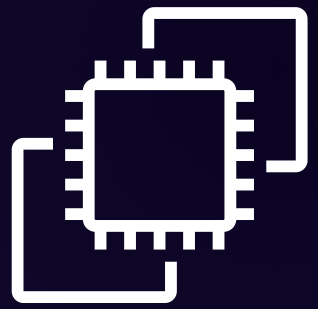

AWS Nitro System

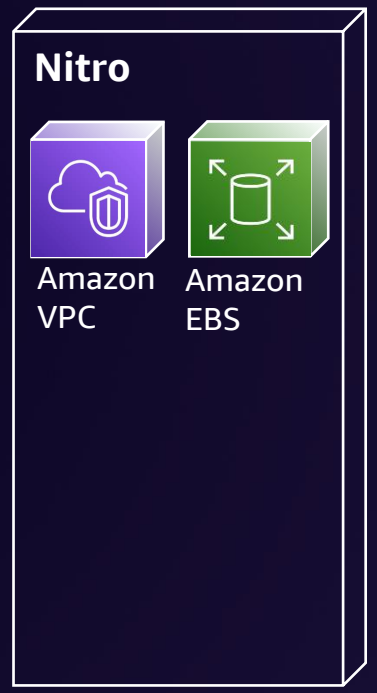

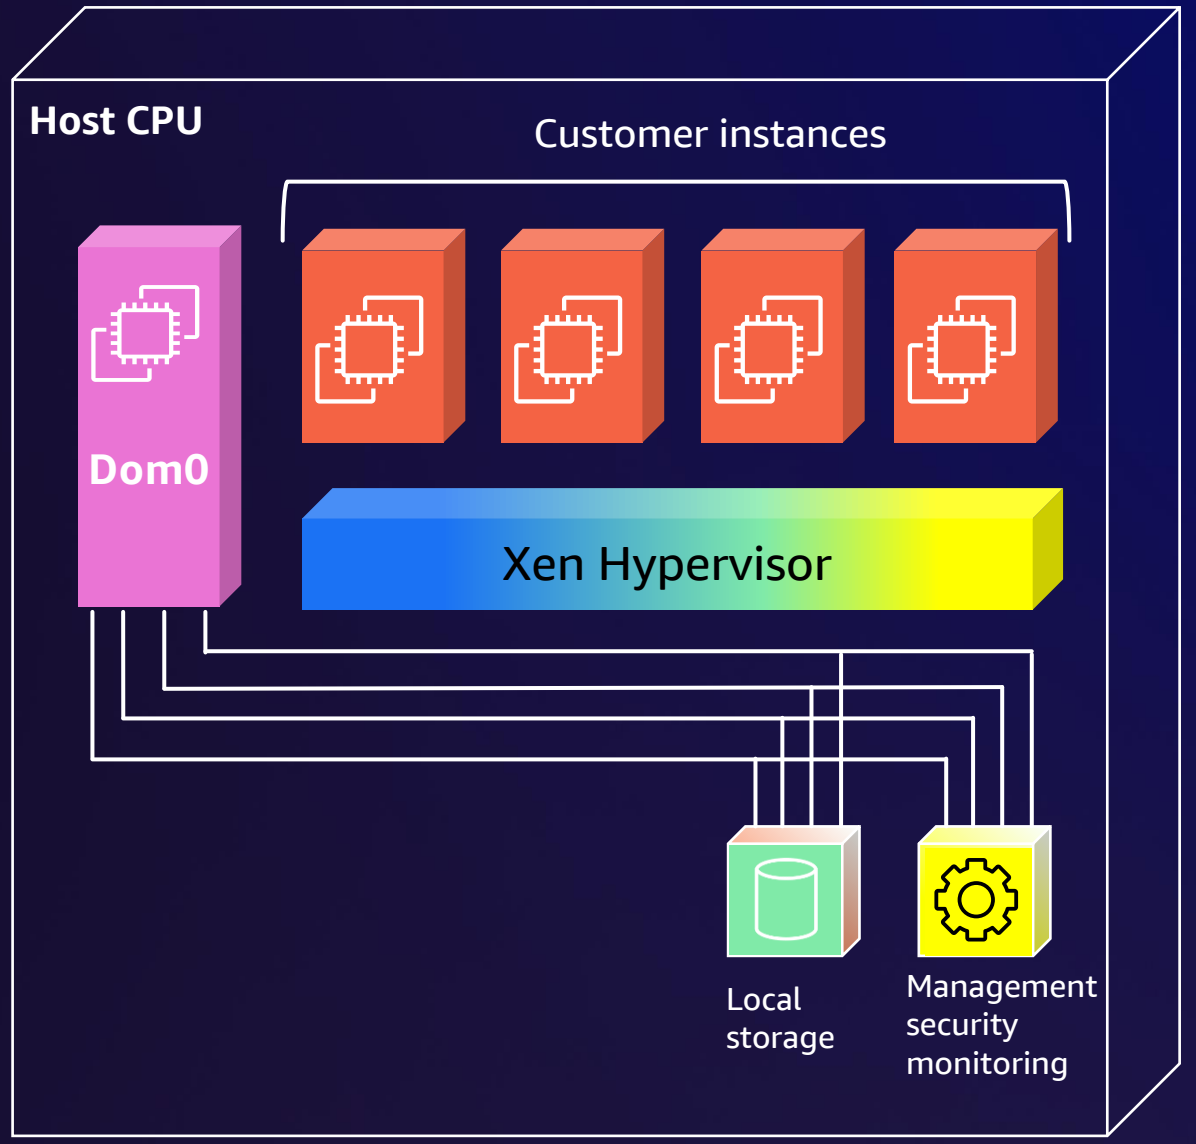

AWS Nitro System

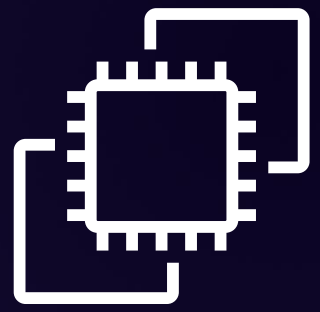

AWS Nitro System

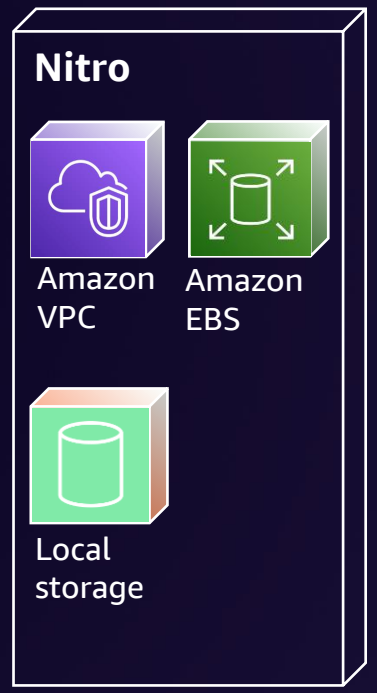

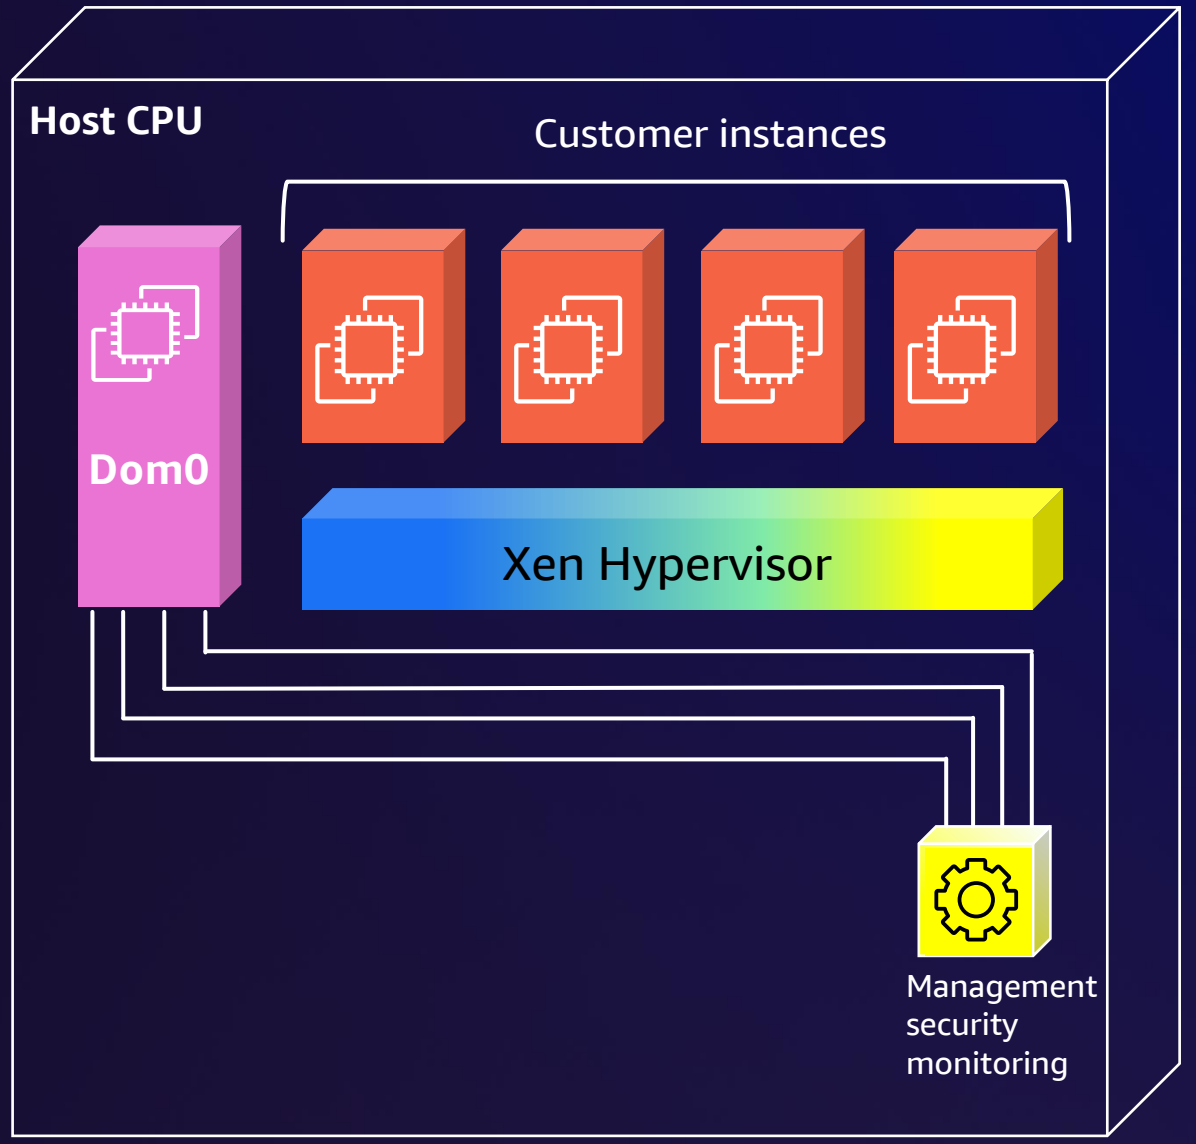

AWS Nitro System

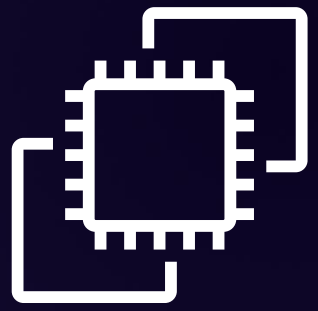

AWS Nitro System

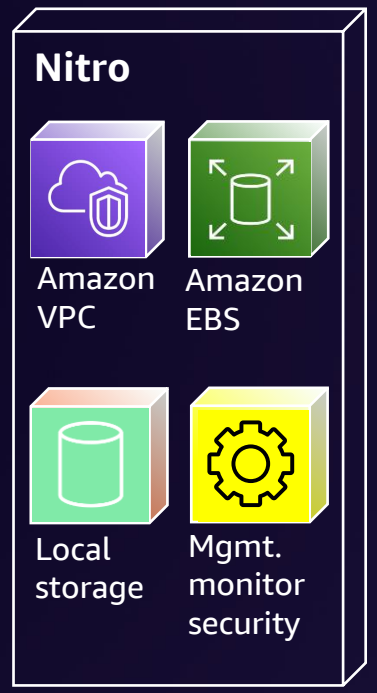

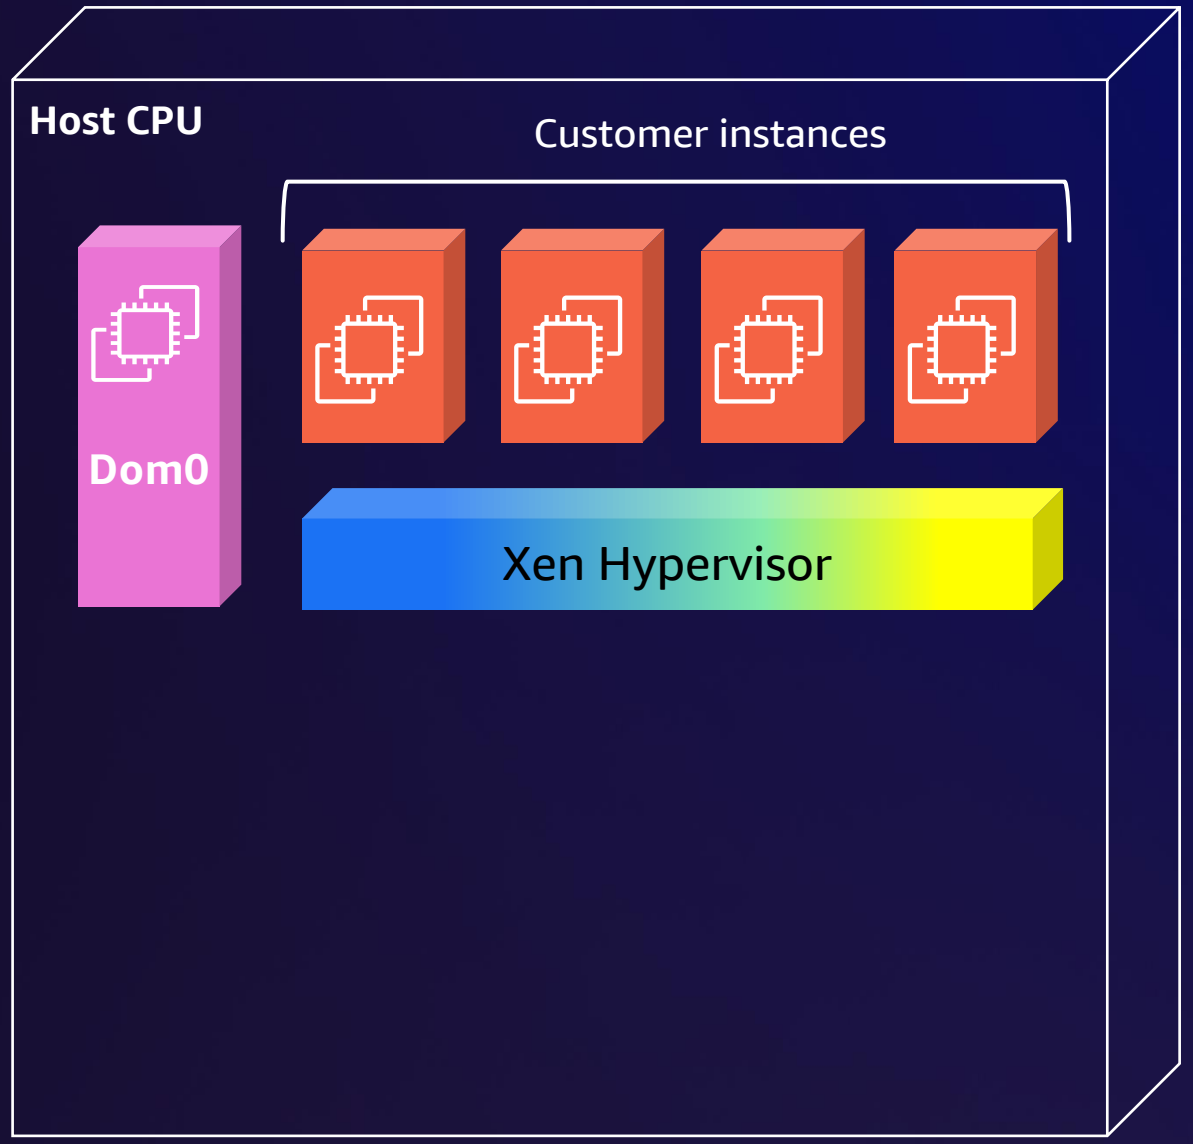

AWS Nitro System

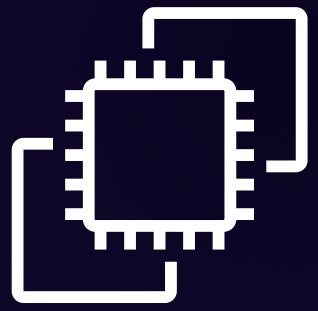

AWS Nitro System

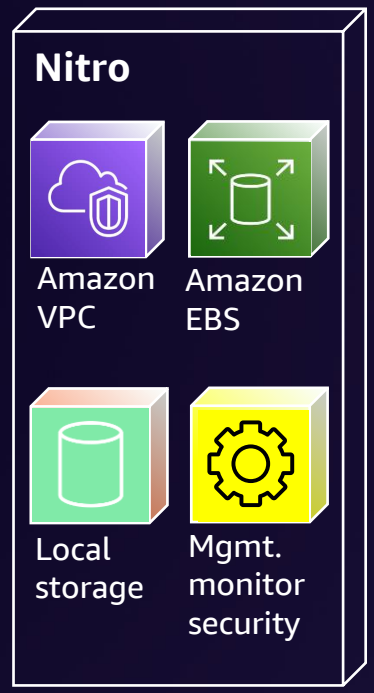

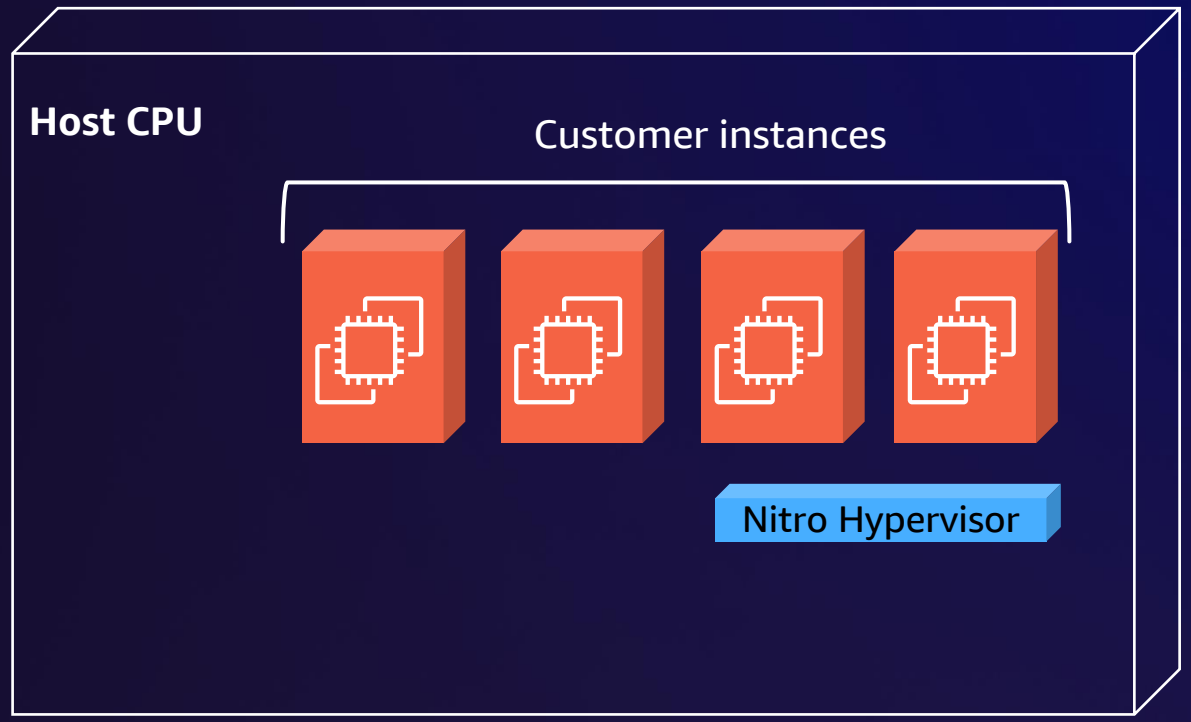

Let’s dig a little deeper into the components

AWS Nitro System

Nitro Cards

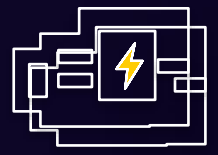

- Local NVMe storage
- Elastic Block Storage
- Instance storage
- System controller
- Hardware root of trust

Nitro Security Chip

- Integrated into motherboard
- Protects hardware resources

Nitro Hypervisor

- Lightweight hypervisor
- Memory and CPU allocation
- Bare metal-like performance

The Nitro Cards

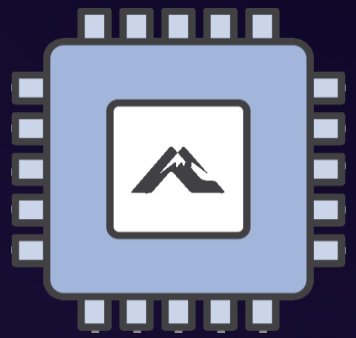

Continuous silicon innovation

NITRO 1

NITRO 2

NITRO 3

NITRO 4

Nitro V5

- 2x more transistors
- 50% faster DRAM speed
- 2x more PCIe bandwidth

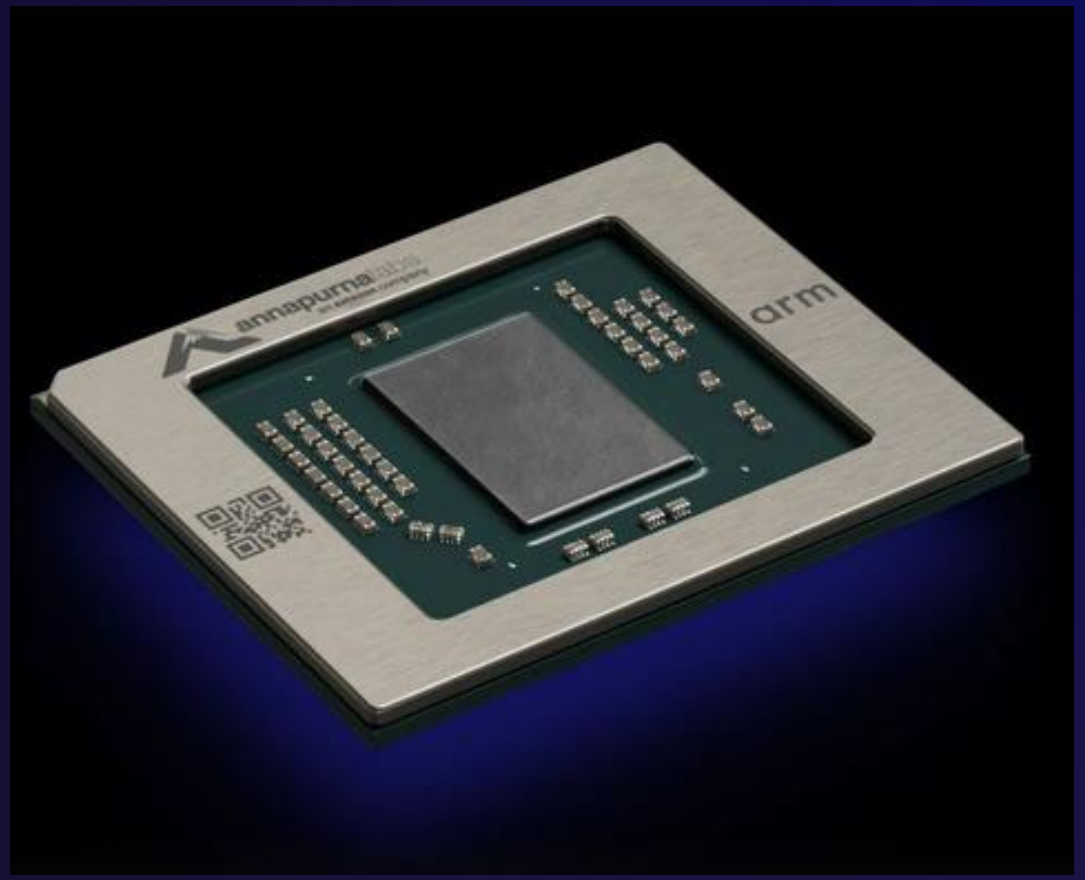

Nitro Cards

ENA PCIe controller

VPC data plane

NVMe PCIe controller

EBS data plane

NVMe PCIe controller

Instance Store

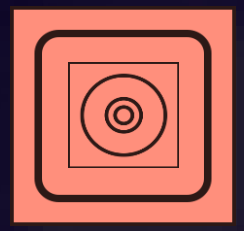

System control

Root of trust

Nitro Controller

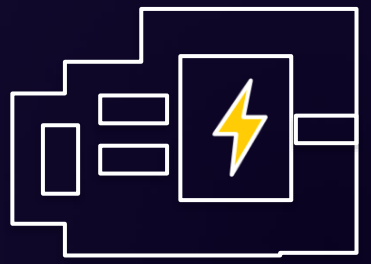

Nitro Controller

System control

- Provides passive API endpoint
- Coordinates other Nitro System components

Nitro Controller APIs

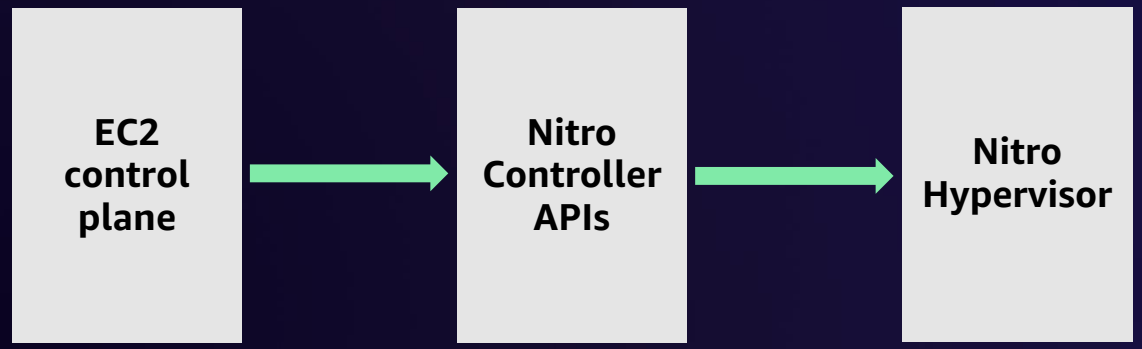

The Nitro Controller presents system management APIS to the dedicated EC2 control plane network.

Every API action is authenticated, authorized, and logged. Each control plane component is only authorized for the set of operations needed for it to complete its business purpose.

Formally verified for memory safety in the face of any configuration file or network request.

Nitro Cards for I/O

Amazon Virtual Private Cloud (Amazon VPC)

Amazon Elastic Block Store (Amazon EBS)

Amazon EC2 instance storage

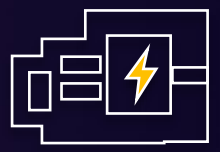

- VPC encapsulation
- Encryption
- Security groups
- Limiters
- Routing

- Encryption
- Communication w/ EBS servers over EBS Protocol

- Encryption
- Limiters
- Drive monitoring

- ENA network card
- EFA device

- NVME block device

- NVME block device

Nitro architecture

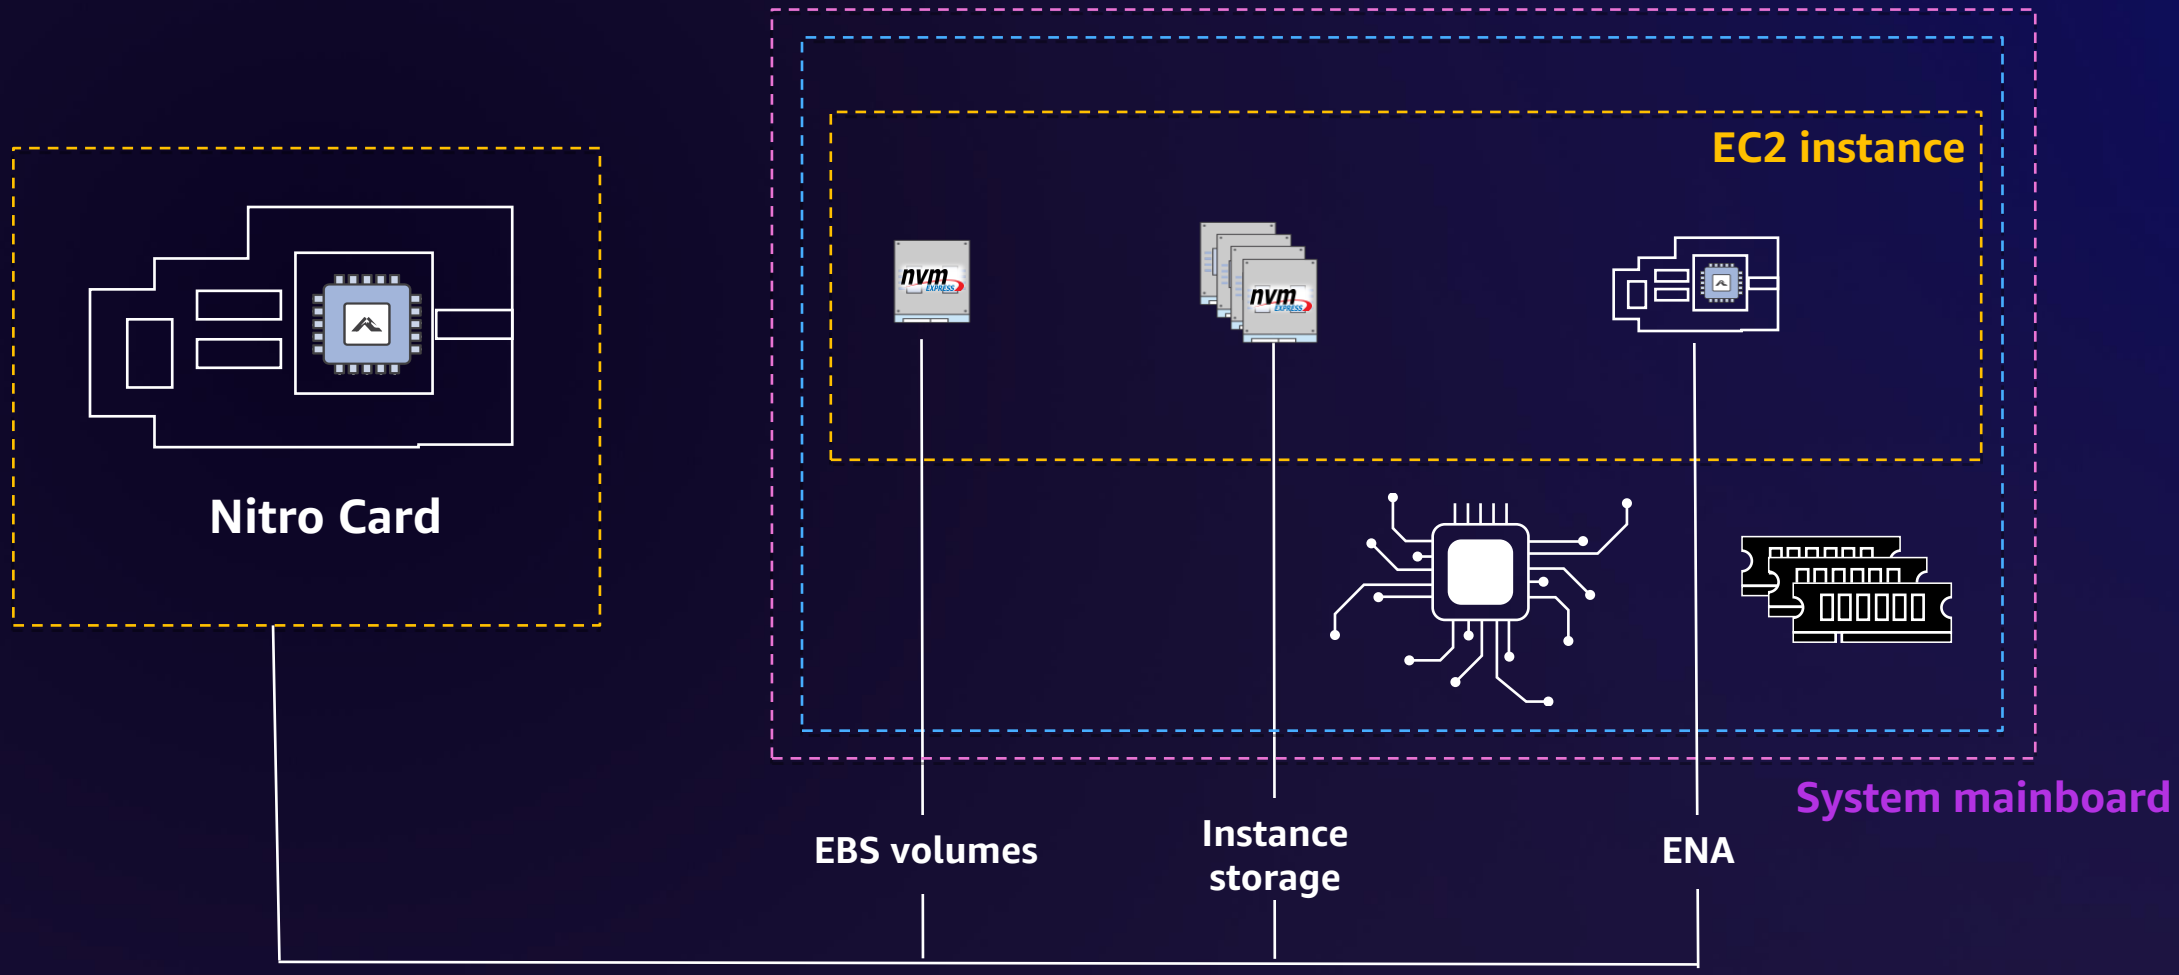

Transparent hardware encryption

The encryption keys used for EBS, local instance storage, and for VPC networking are only ever present on the system in plaintext within the protected memory of a Nitro Card

Encryption key management

Amazon EBS

- Volumes have independent lifetimes (plus snapshots); therefore, key management via AWS KMS (FIPS 140-2 Level 3 validated HSMs)

Instance storage

- Locally generated, used, deleted (instance lifecycle)

VPC

- Seed materials regionally generated
- Seeds distributed, not actual secrets; key derived on Nitro device; **rotated frequently**

Nitro Hypervisor

KVM-based hypervisor with custom memory management and minimized user space

Only executes on behalf of instance, quiescent

With Nitro, the hypervisor is minimal and performant

Nitro Hypervisor: Minimalization

The Nitro Hypervisor has been deliberately minimization

Intentionally excludes

- Networking stack
- General purpose file system implementations
- Peripheral driver support
- SSH server
- Shell
- Etc.

Nitro Hypervisor: Primary functions

1. Receive virtual machine management commands (start, stop, and so on) sent from the Nitro Controller
2. Partition memory and CPU resources by utilizing hardware virtualization features of the server processor
3. Assign SR-IOV virtual functions provided by Nitro hardware interfaces (NVMe block storage for EBS and instance storage, Elastic Network Adapter [ENA] for network, and so on) through PCIe to the appropriate VM

Nitro virtualized architecture

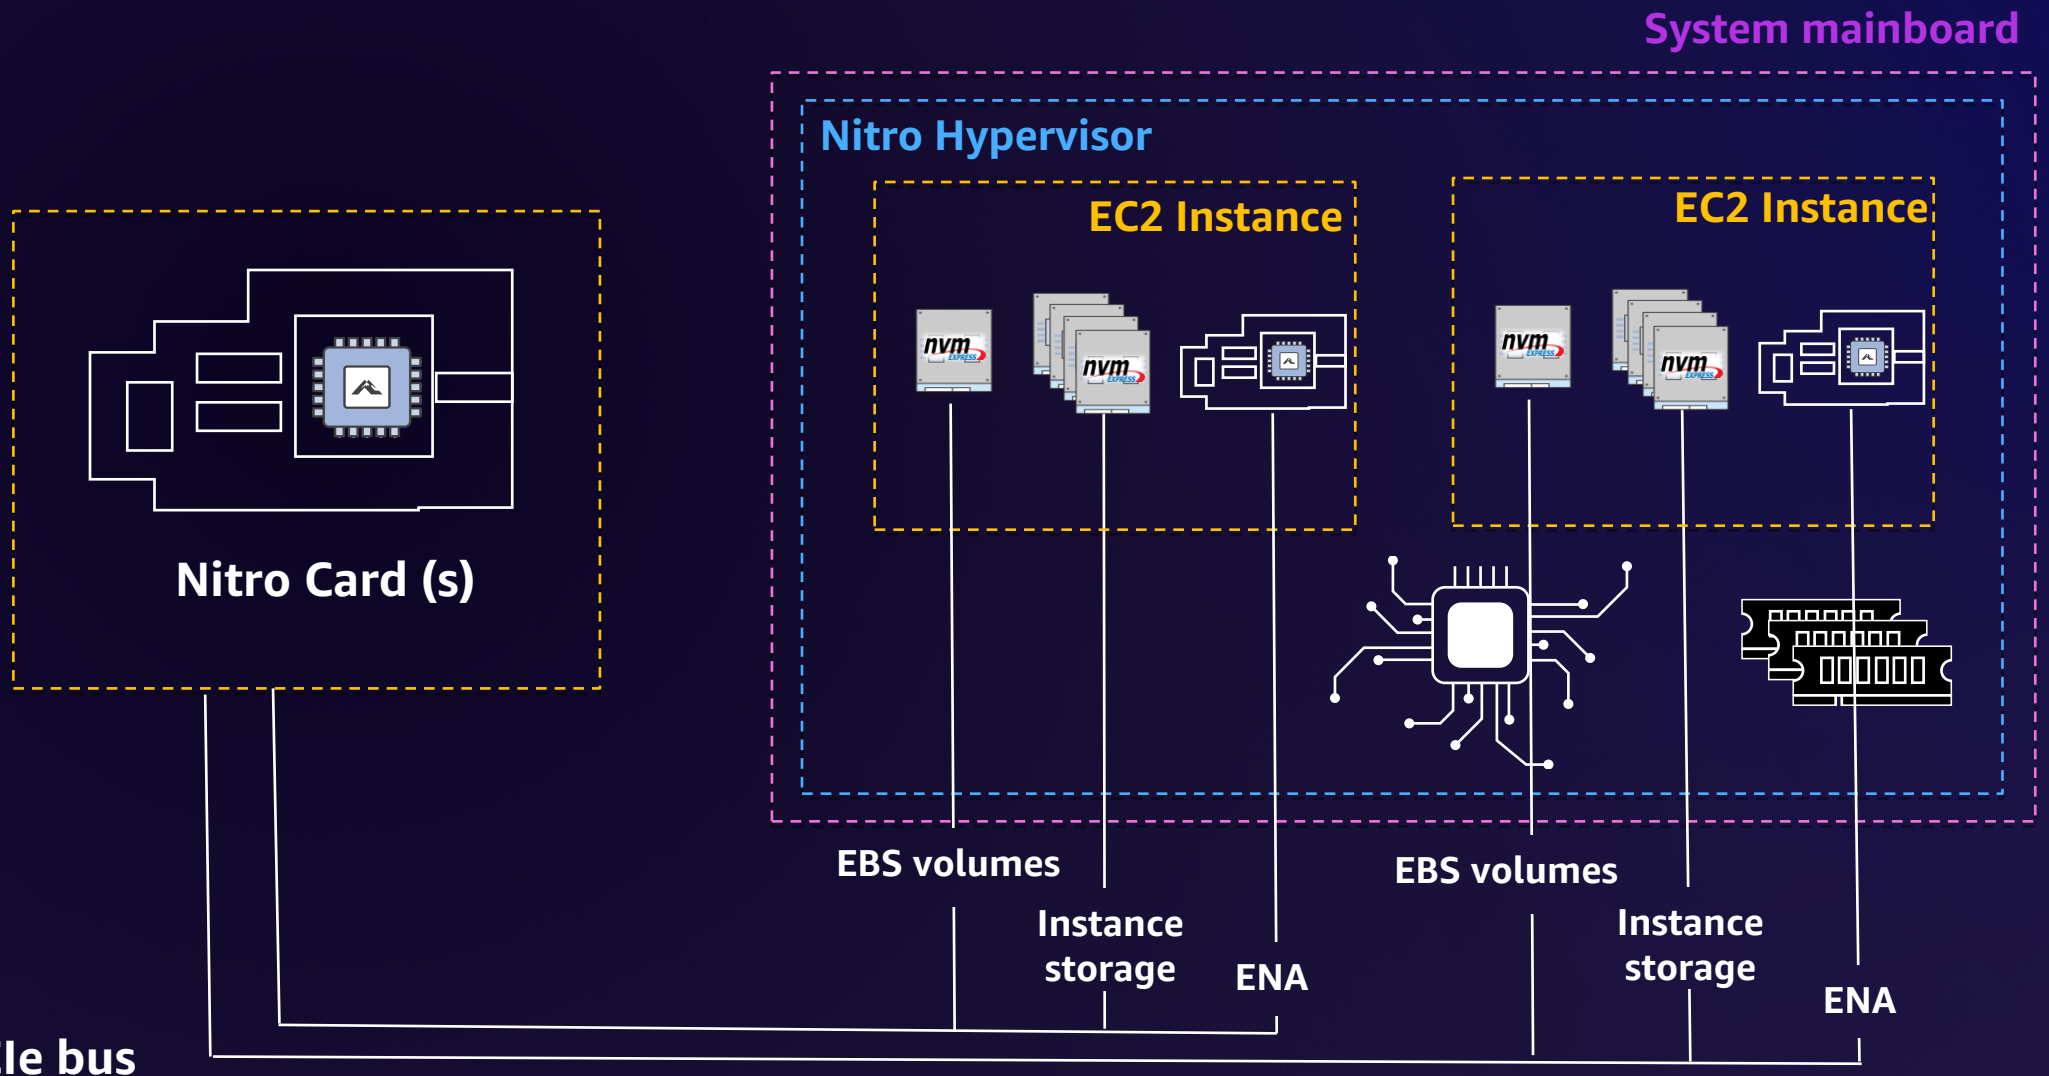

Nitro bare-metal architecture

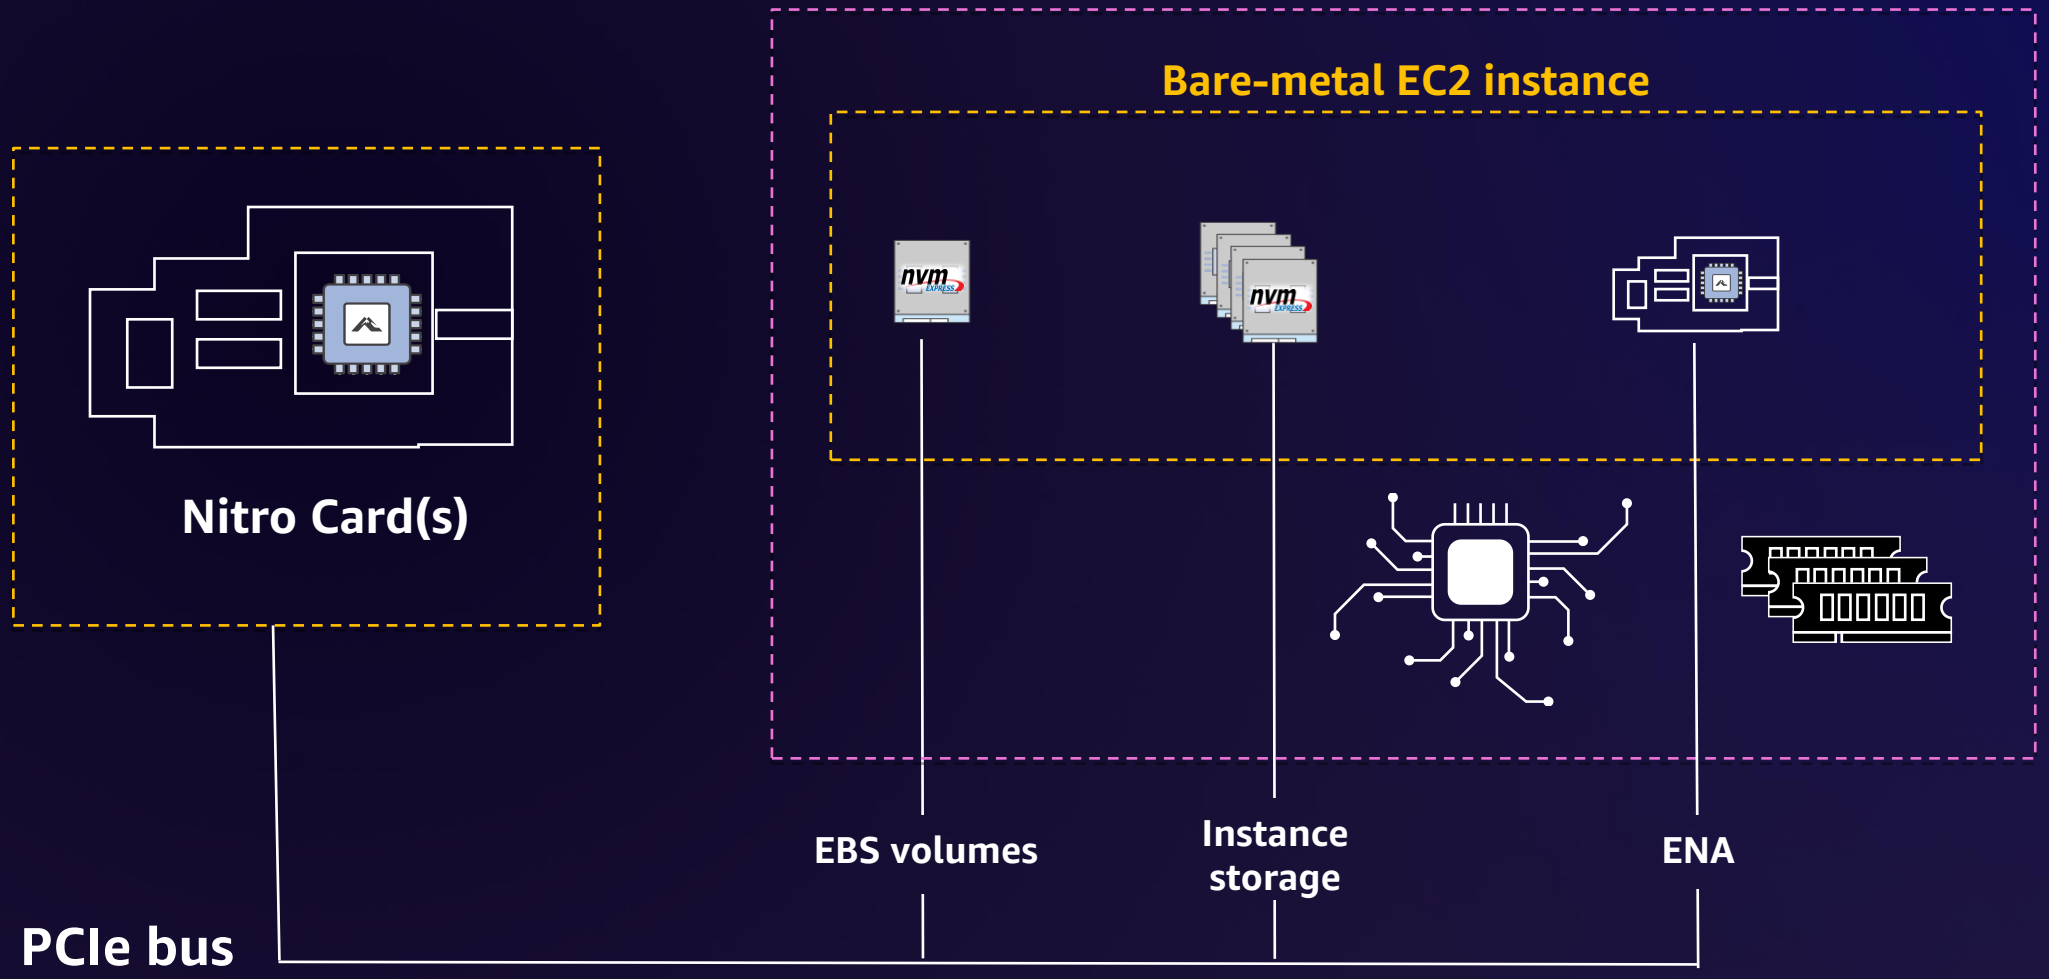

Nitro Hypervisor performance overhead

The Nitro Hypervisor provides performance which is nearly indistinguishable from that of an equivalent bare-metal instance

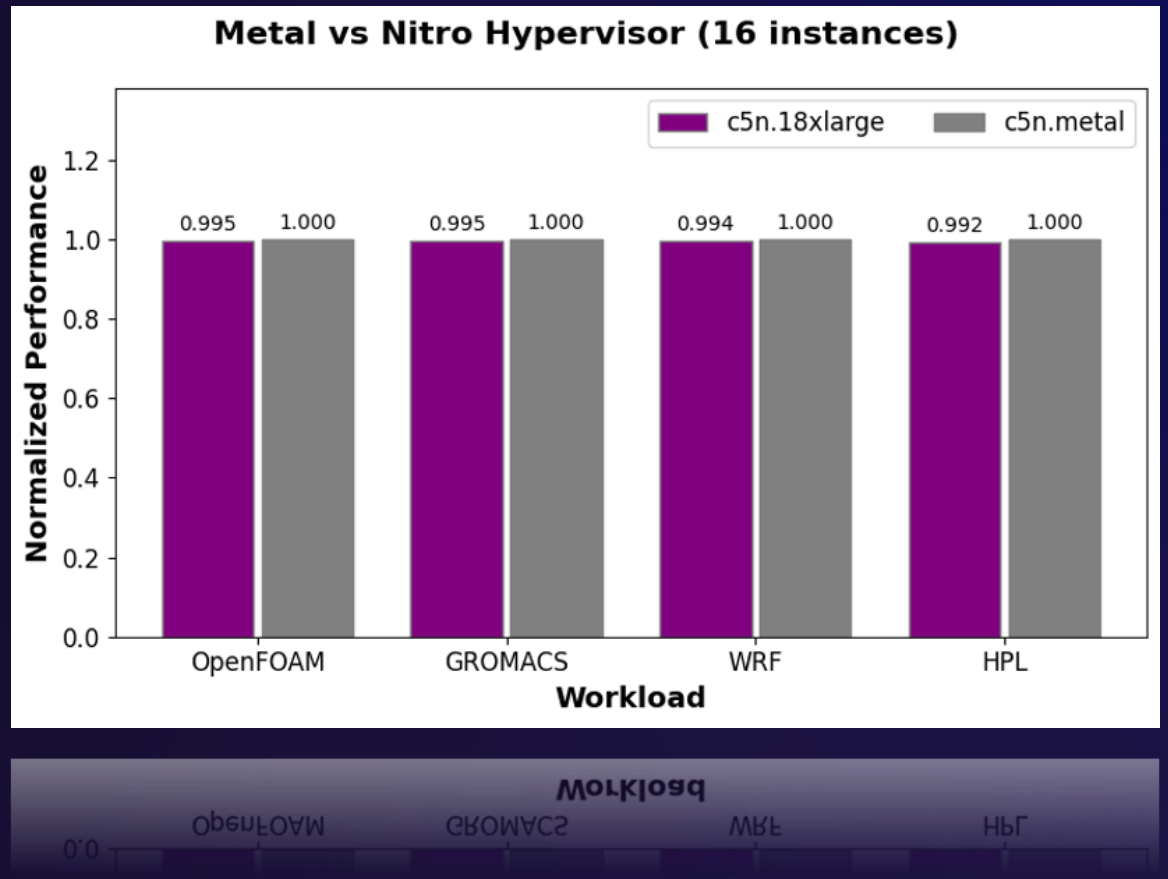

Nitro Security Chip

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

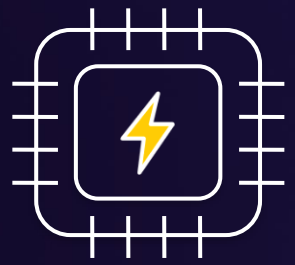

Used by Nitro Controller to monitor hardware, validate and update system firmware/software

Enables a simple, hardware-based root of trust

Nitro Hardware root of trust

Radical simplification enabled by Nitro Cards

All write access to non-volatile storage is blocked in hardware

Simple to understand security due to lack of legacy

Firmware-like hypervisor is injected from controller storage on boot

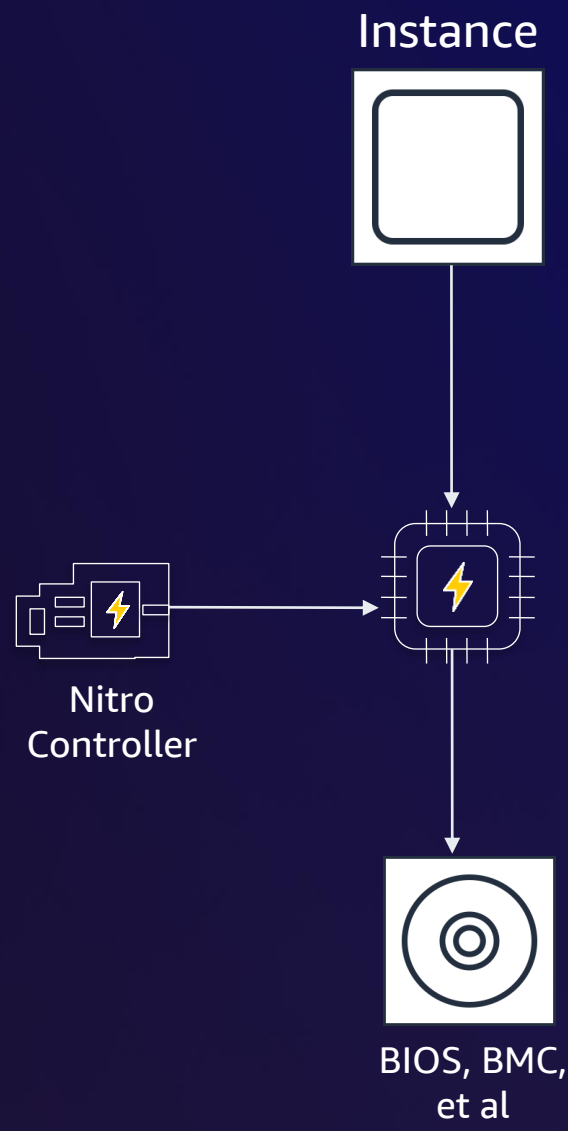

Nitro Hardware root of trust

Radical simplification enabled by Nitro Cards

All write access to non-volatile storage is blocked in hardware

Simple to understand security due to lack of legacy

Firmware-like hypervisor is injected from controller storage on boot

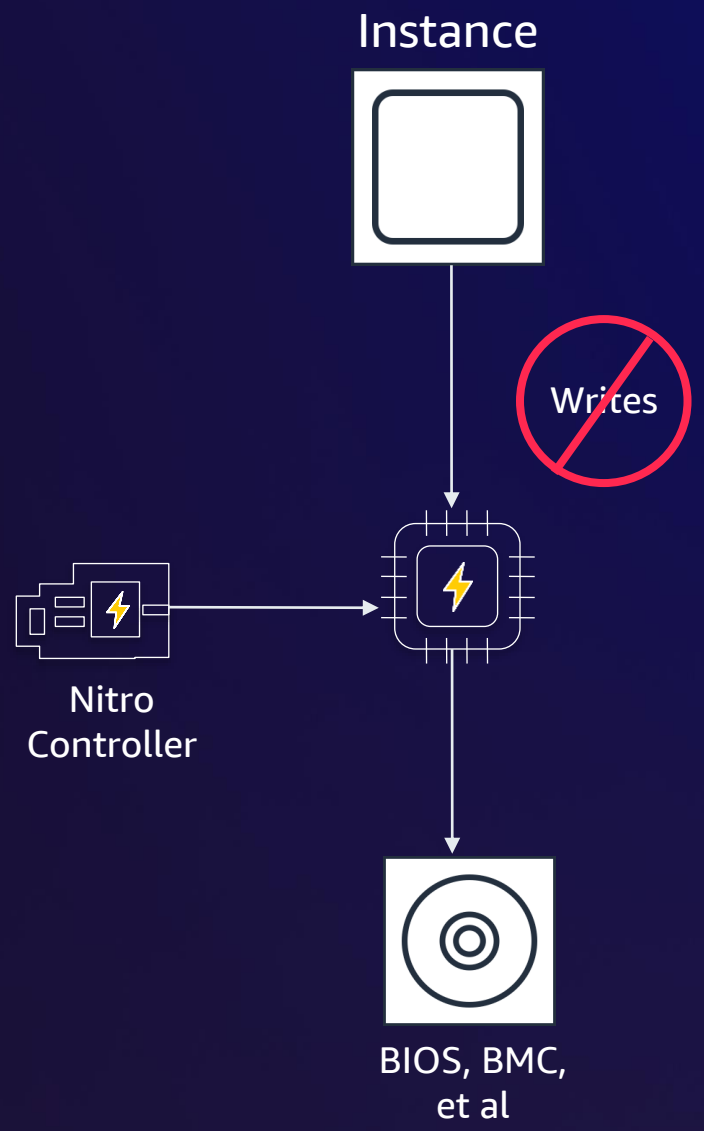

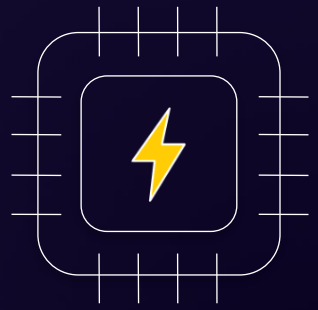

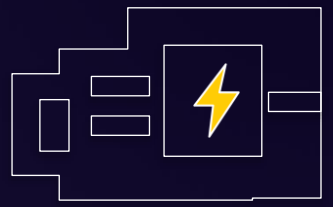

Model Checking Boot Code from AWS Data Centers

Byron Cook[1,2], Kareem Khazem[1,2], Daniel Kroening[3], Serdar Tasiran[1], Michael Tautschnig[1,4(✉)], and Mark R. Tuttle[1]

[1] Amazon Web Services, Seattle, USA
tautschn@amazon.com
[2] University College London, London, UK
[3] University of Oxford, Oxford, UK
[4] Queen Mary University of London, London, UK

Abstract. This paper describes our experience with symbolic model checking in an industrial setting. We have proved that the initial boot code running in data centers at Amazon Web Services is memory safe, an essential step in establishing the security of any data center. Standard static analysis tools cannot be easily used on boot code without modification owing to issues not commonly found in higher-level code, including memory-mapped device interfaces, byte-level memory access, and linker

The Nitro system: Putting it together

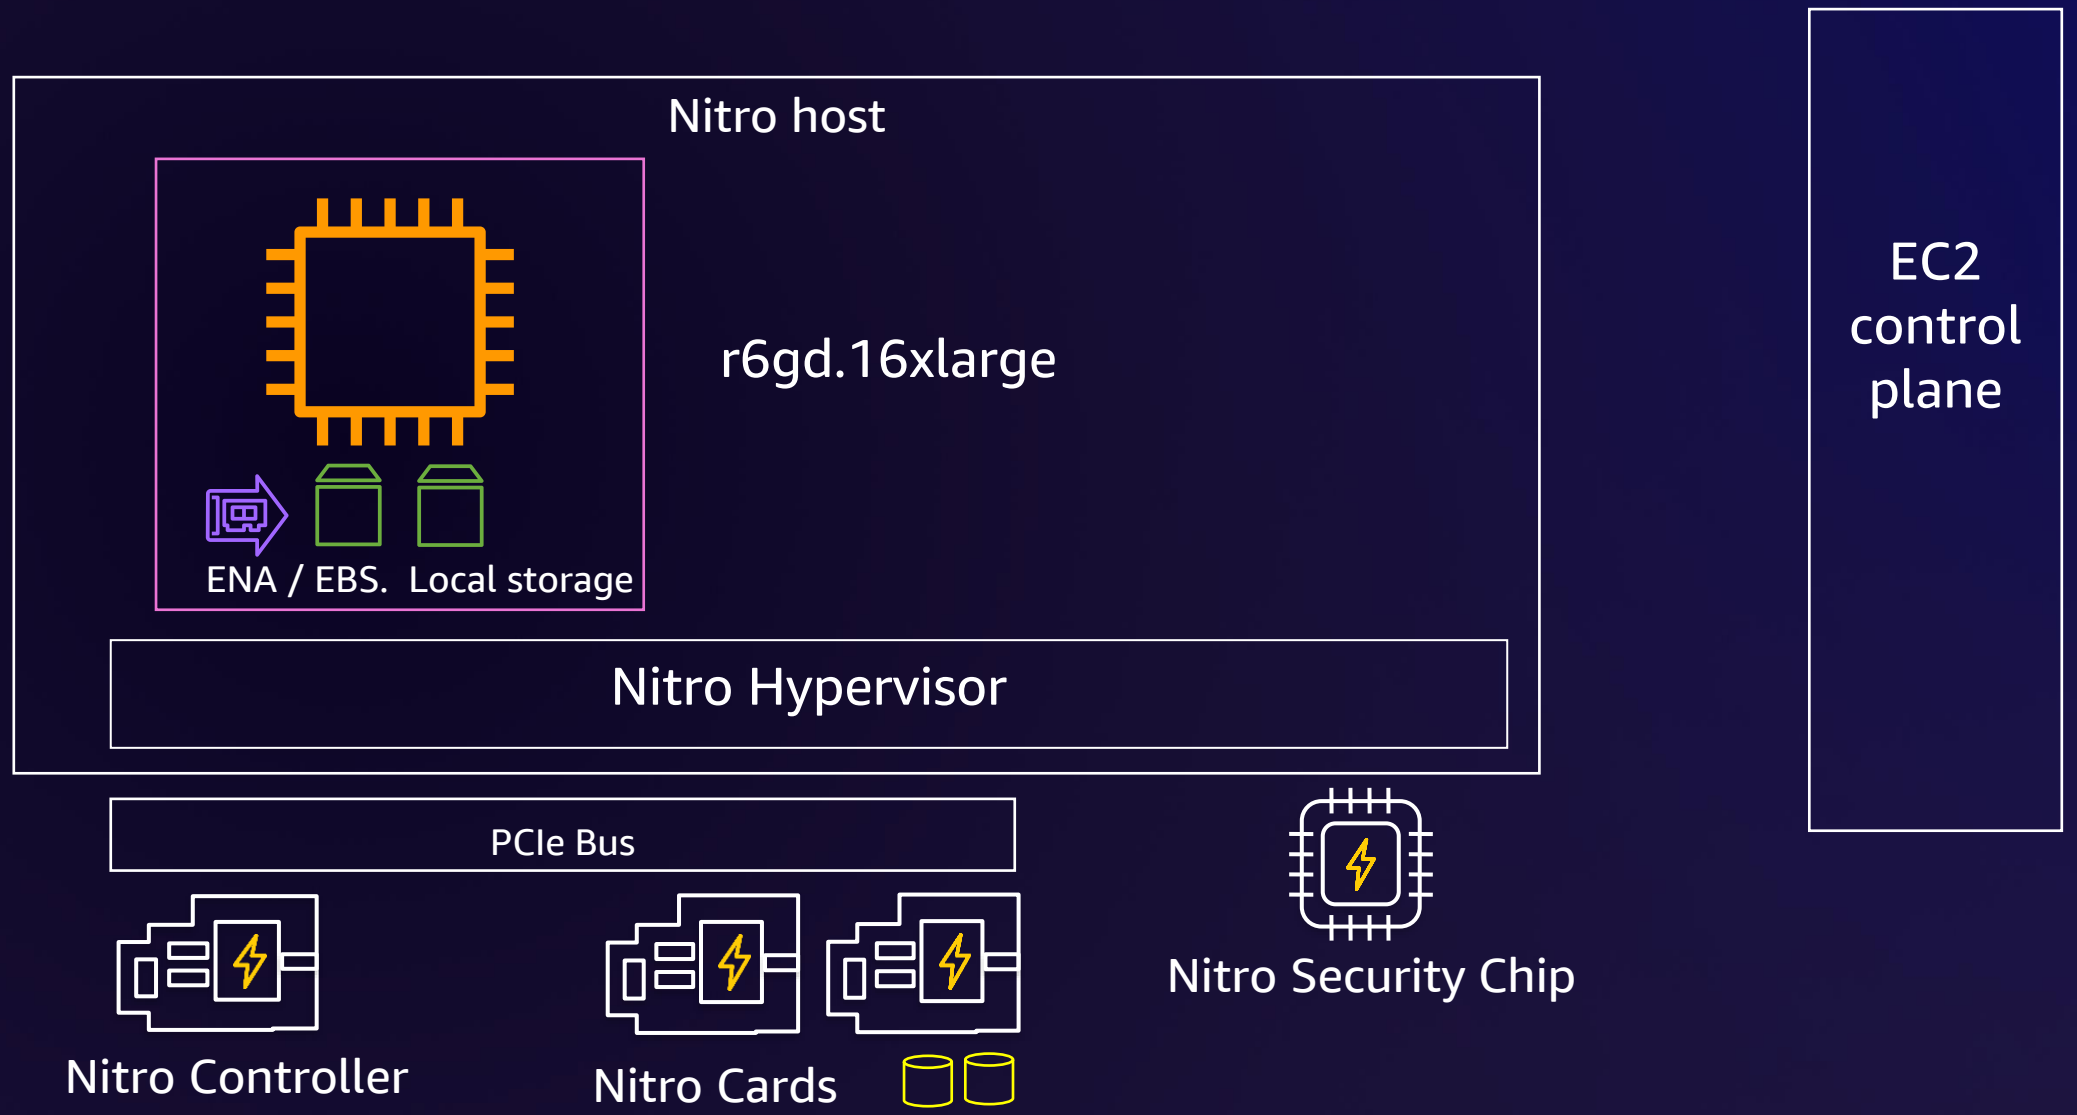

The Nitro system: Putting it together

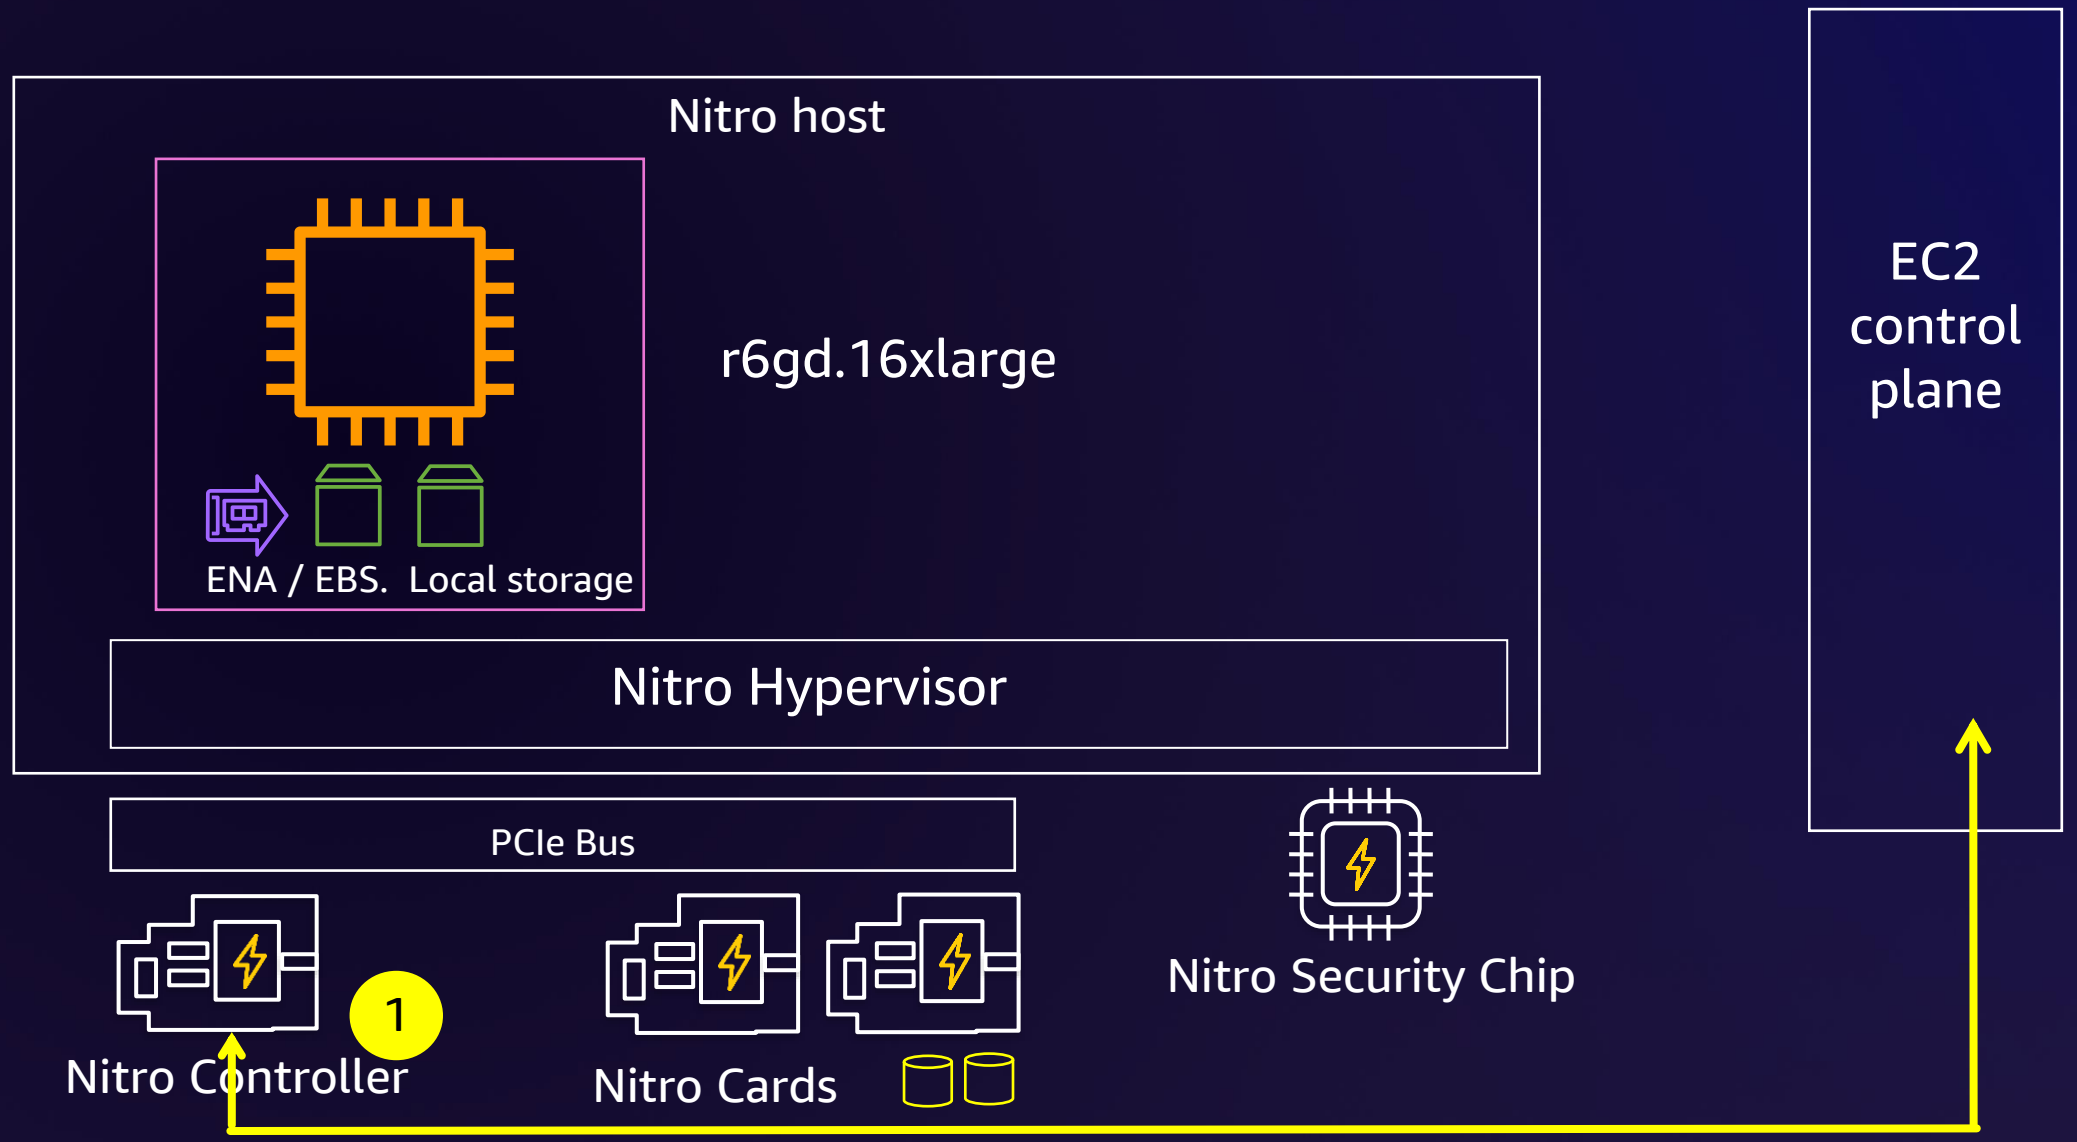

The Nitro system: Putting it together

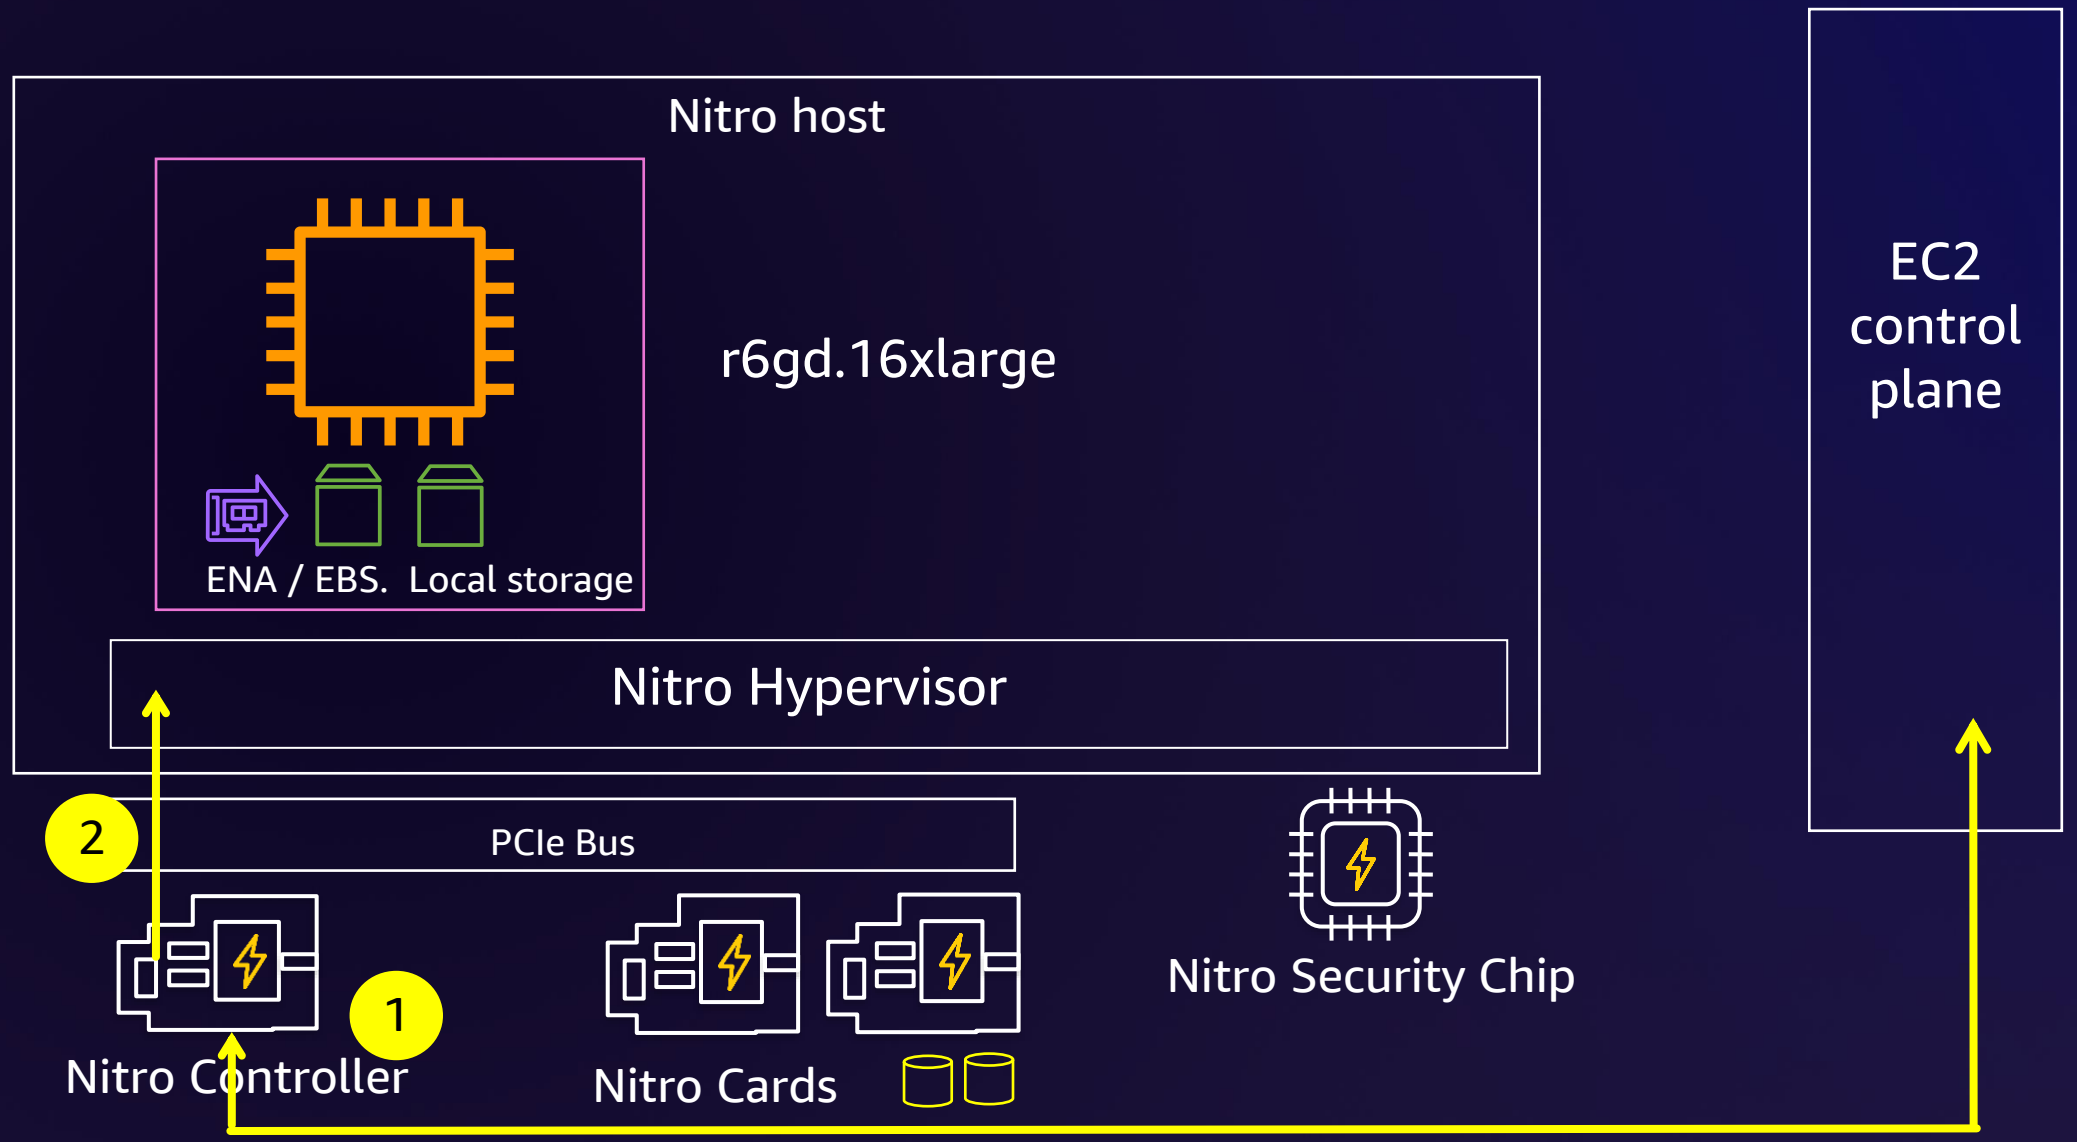

The Nitro system: Putting it together

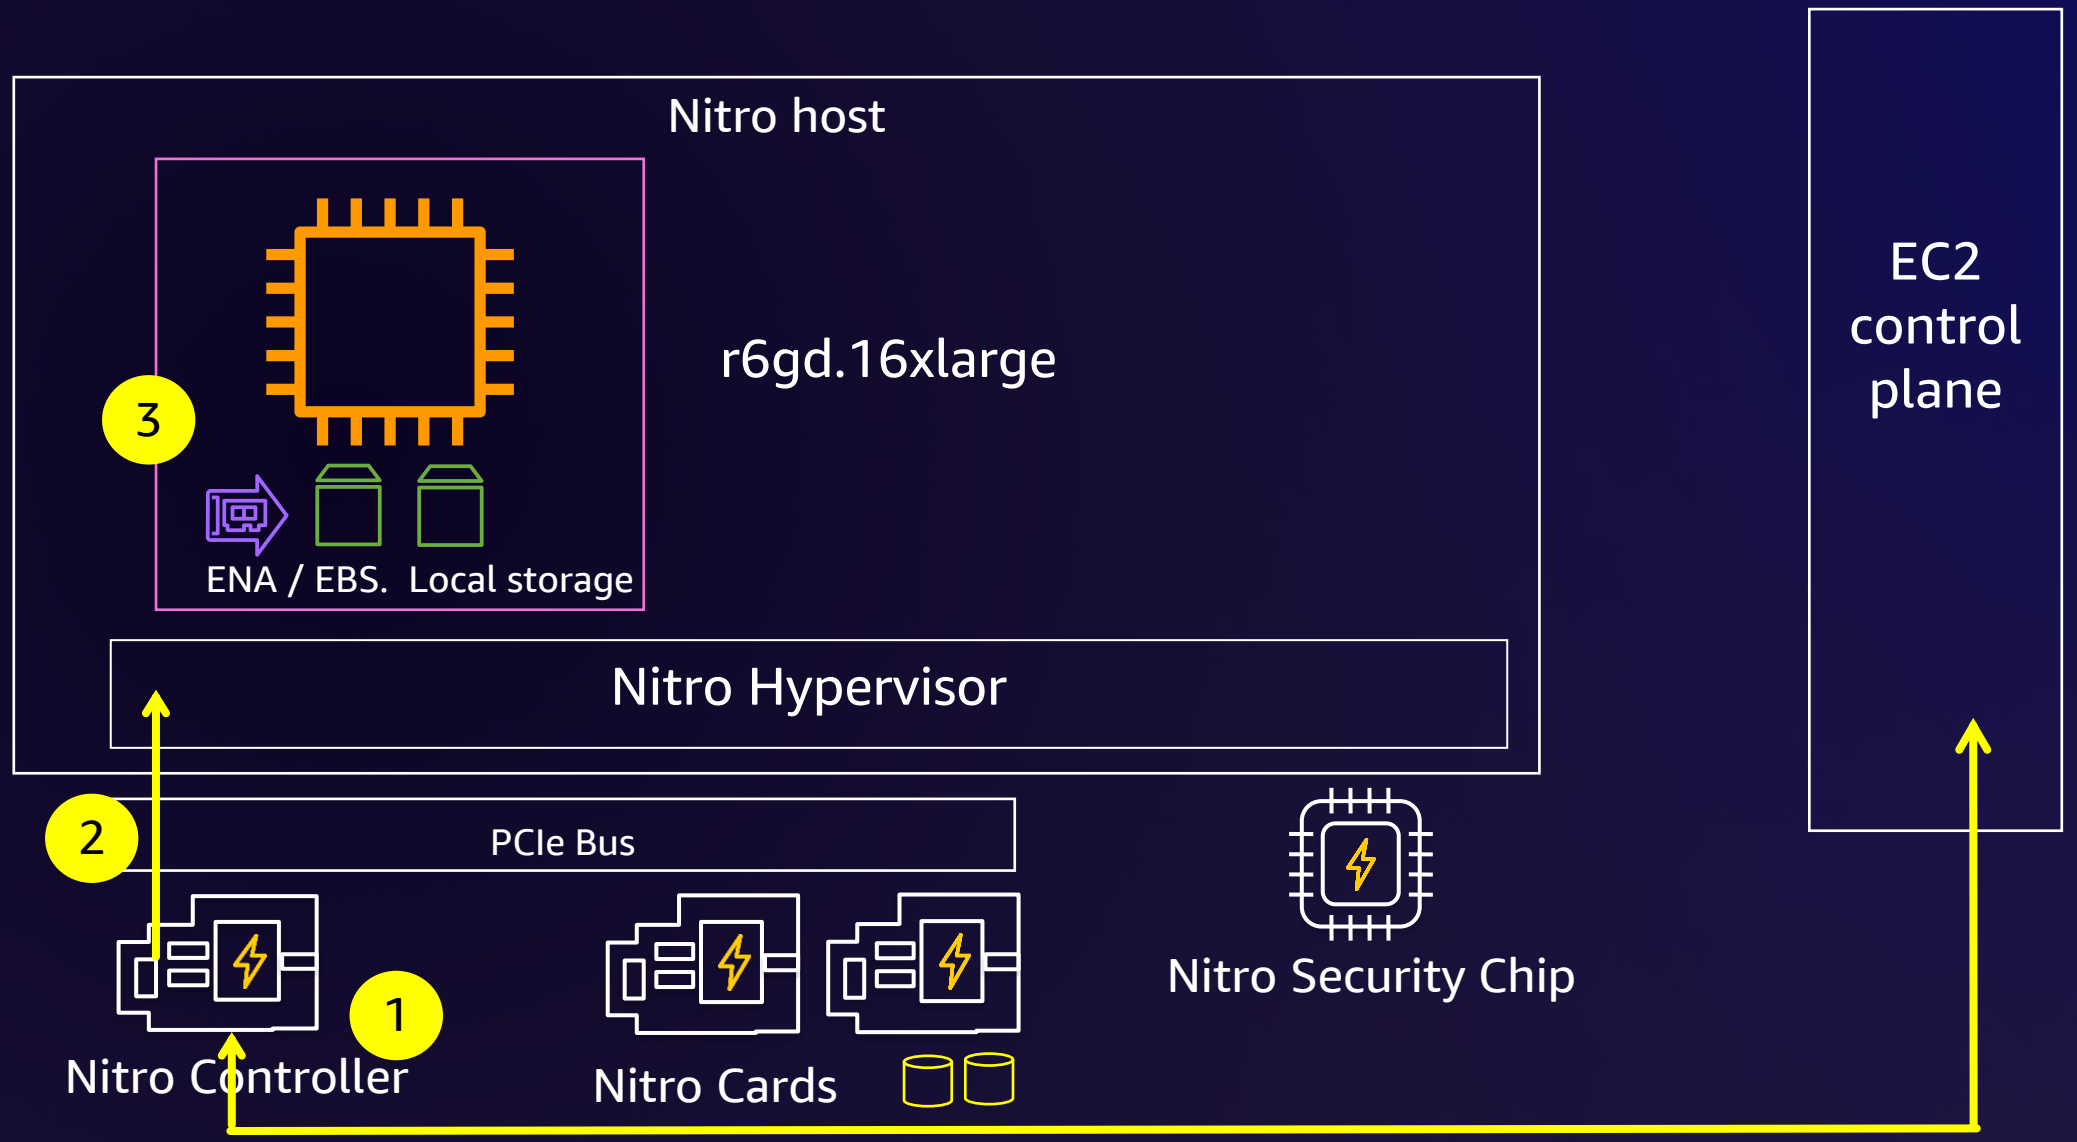

The Nitro system: Putting it together

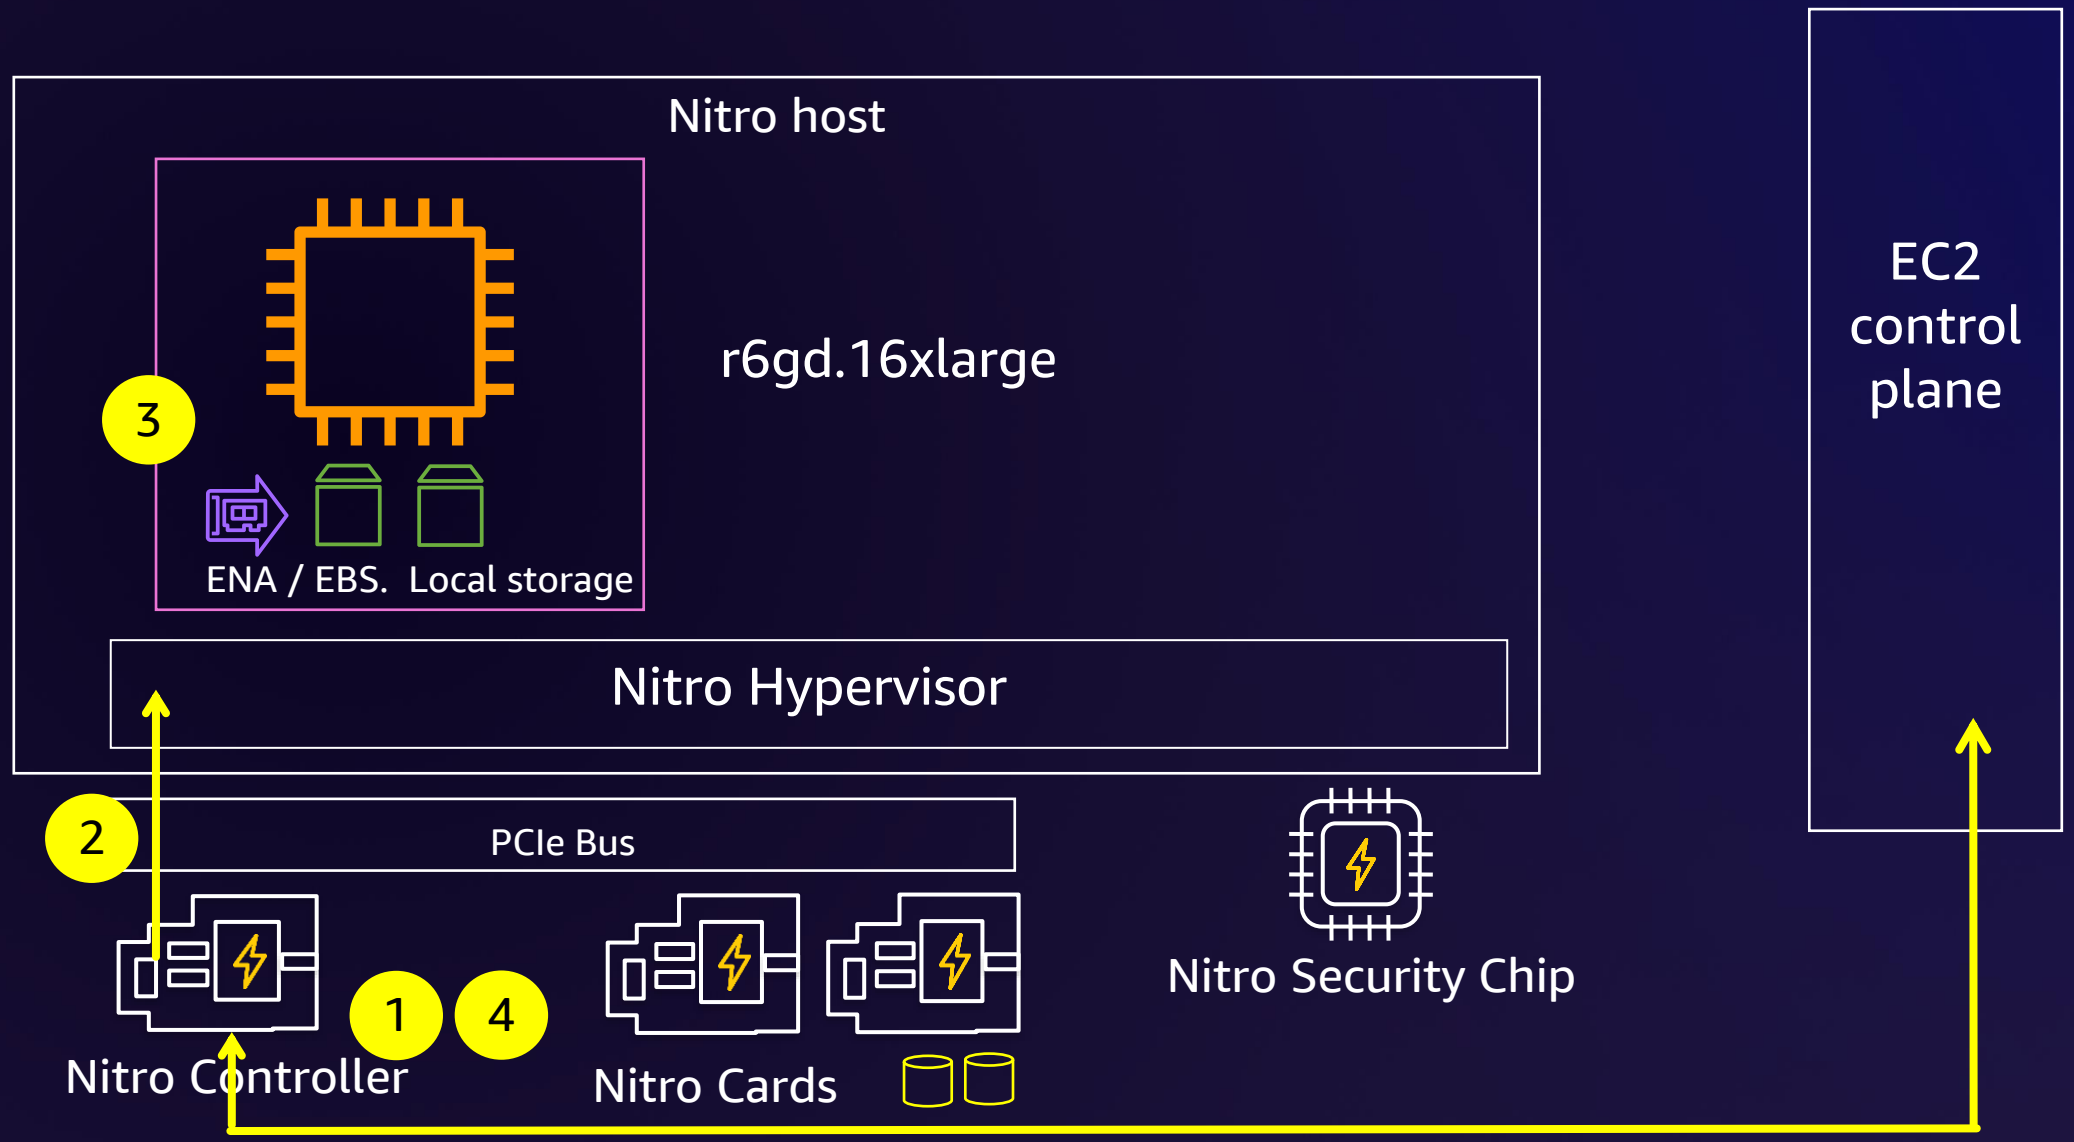

The Nitro system: Putting it together

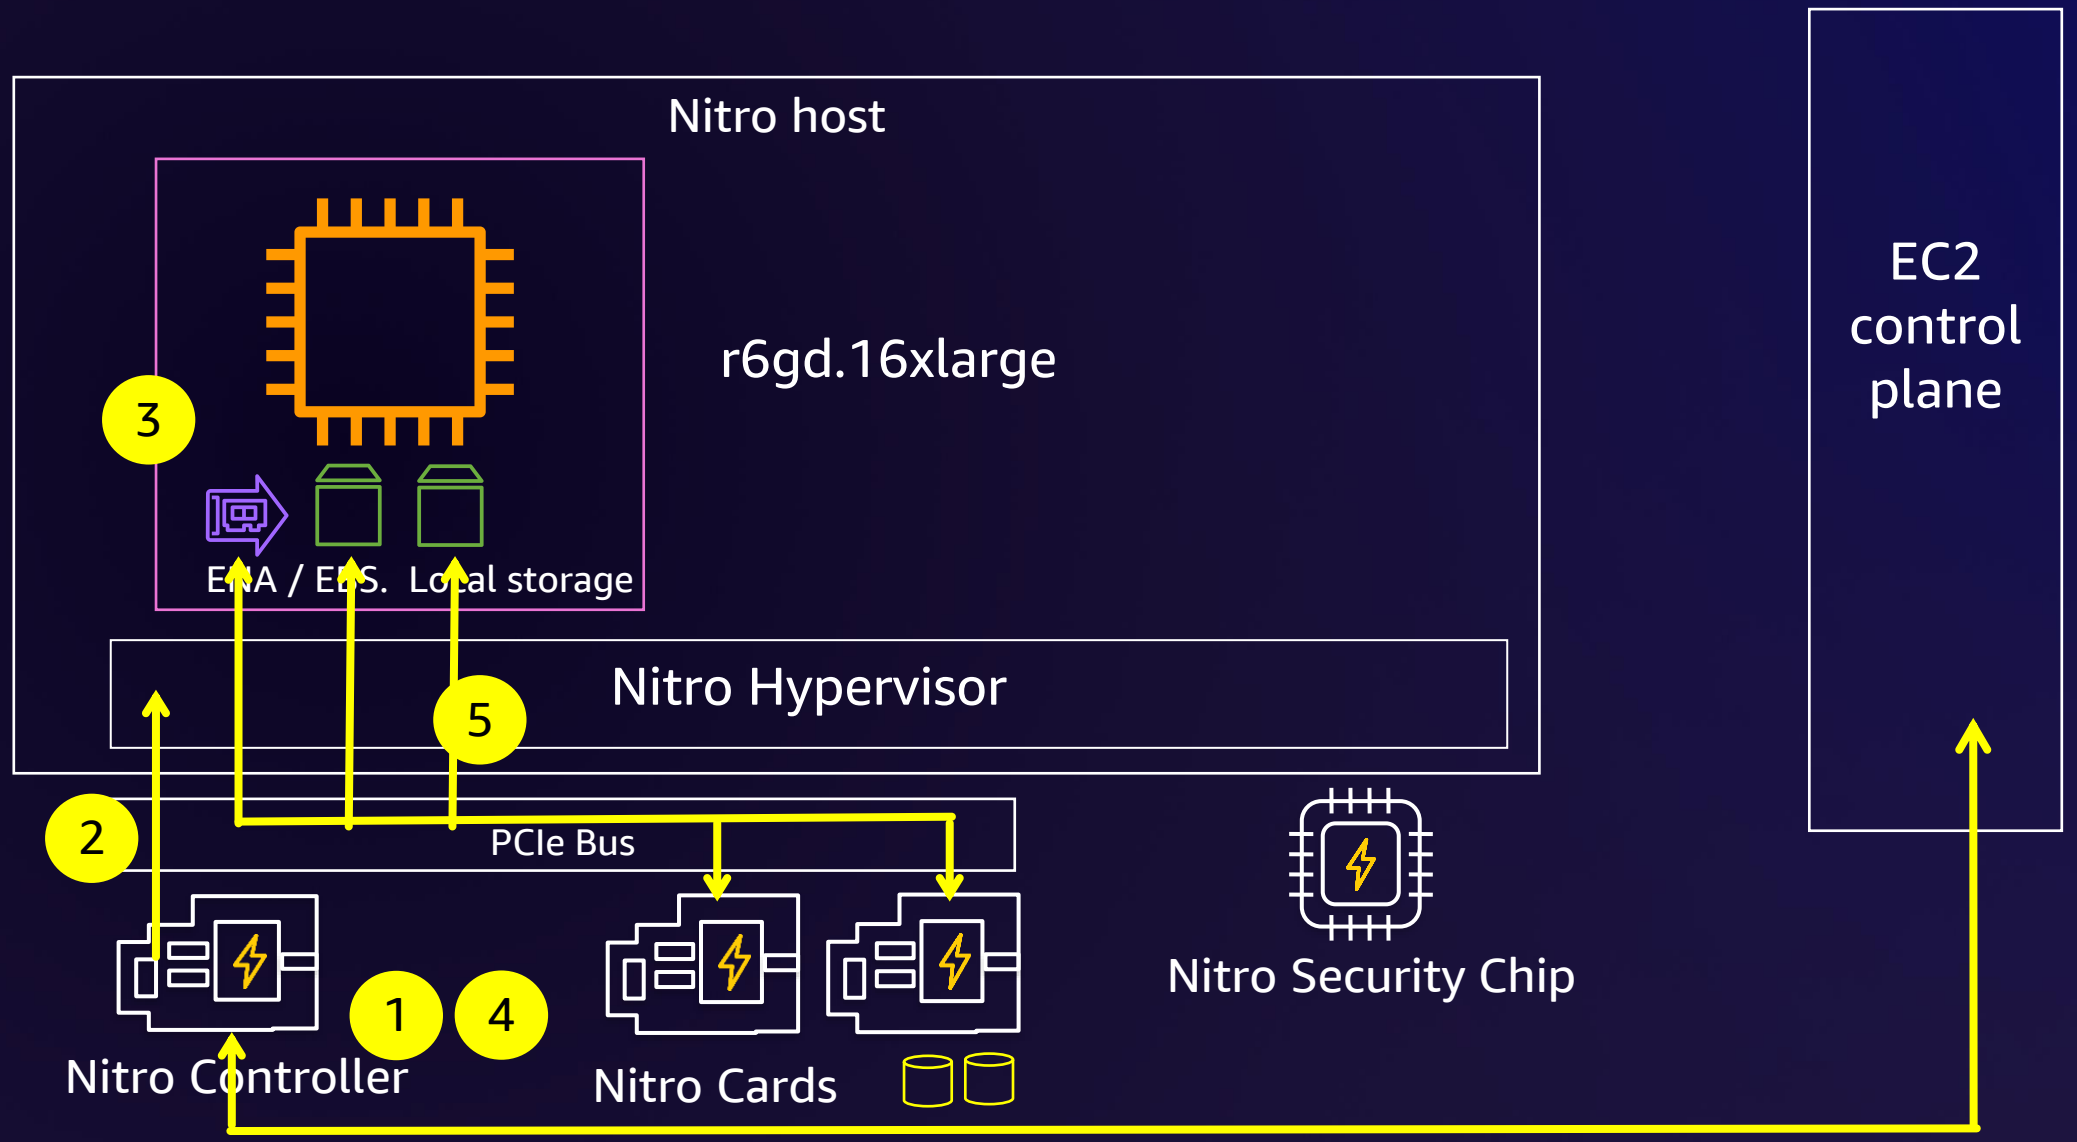

The Nitro system: Putting it together

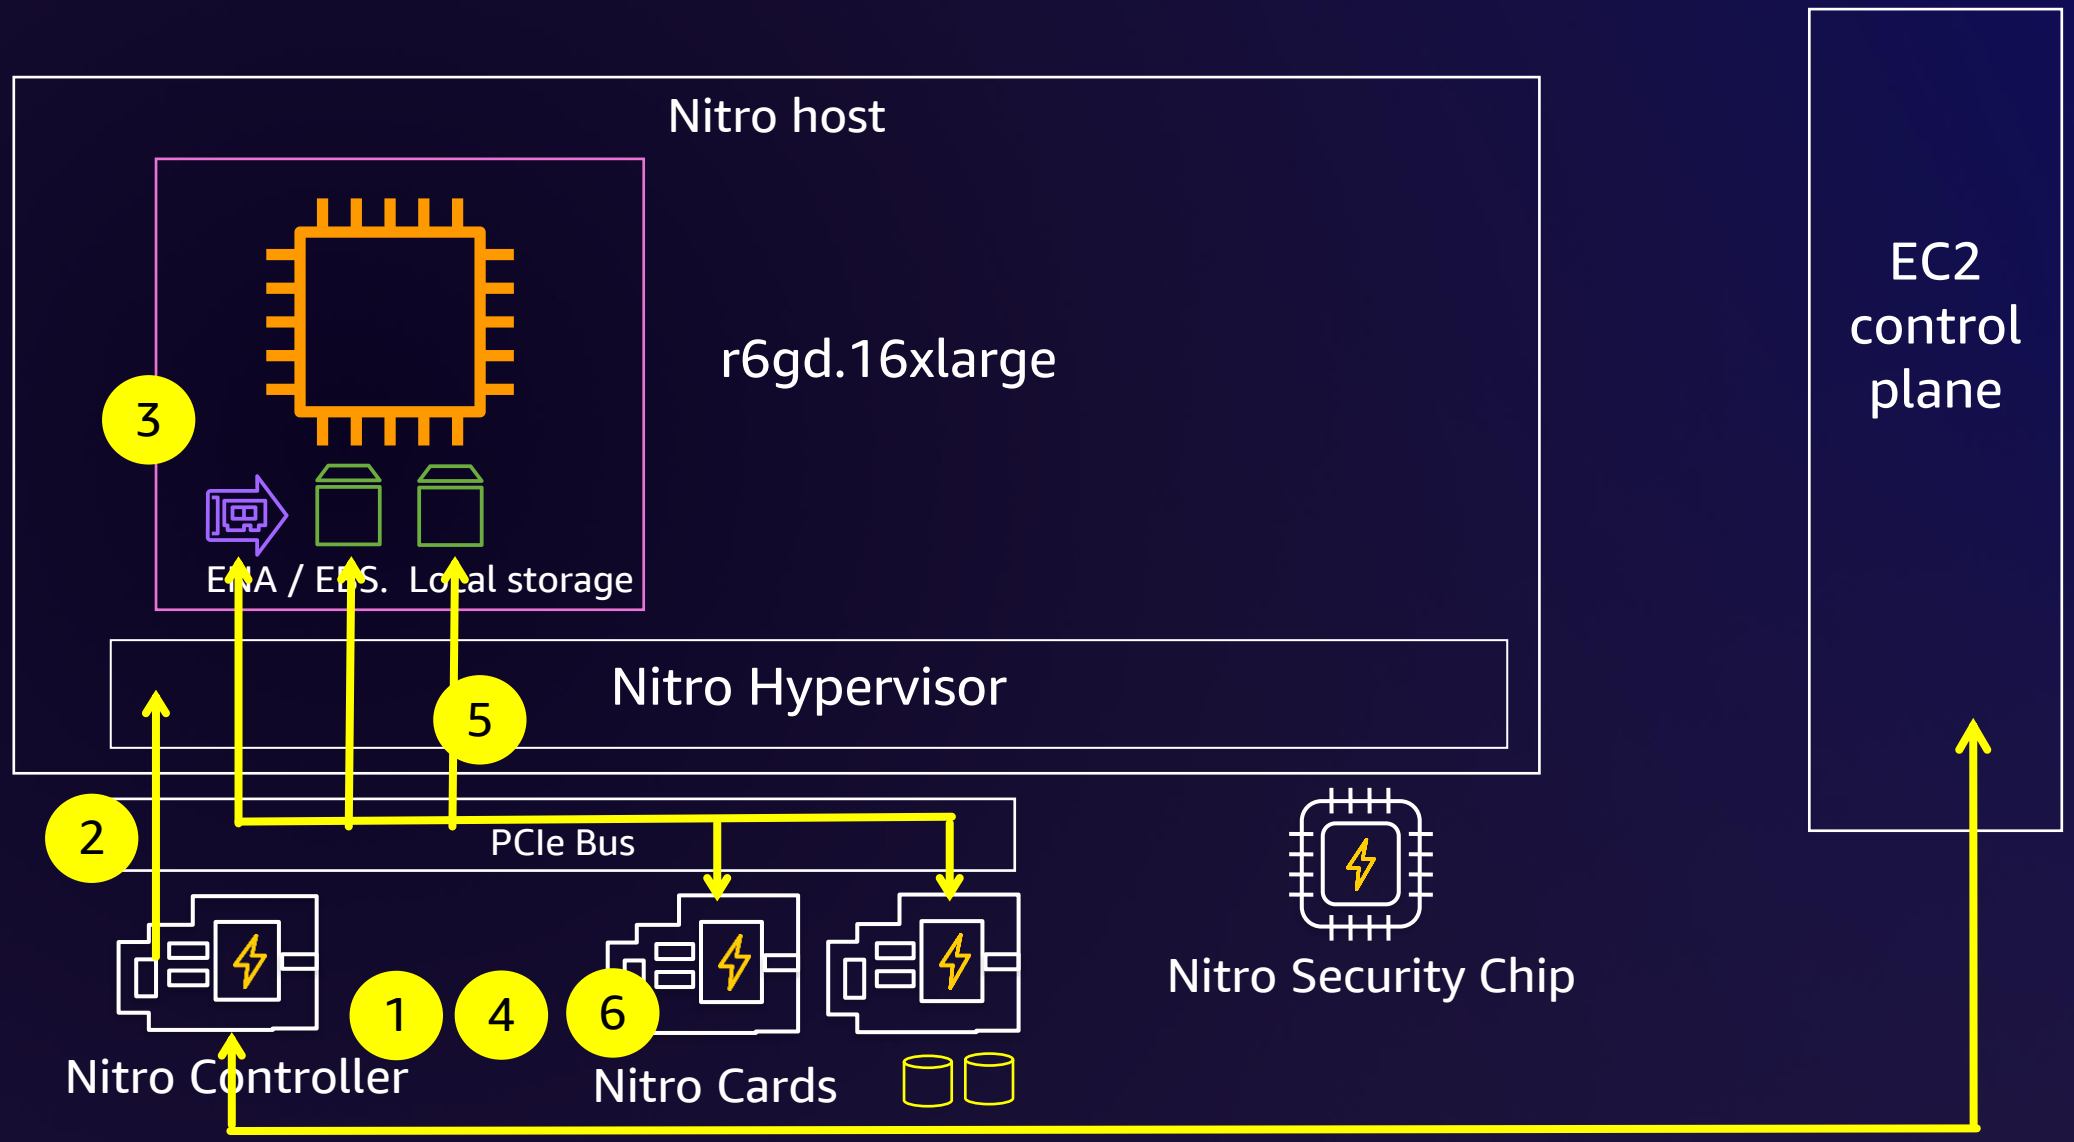

The Nitro system: Putting it together

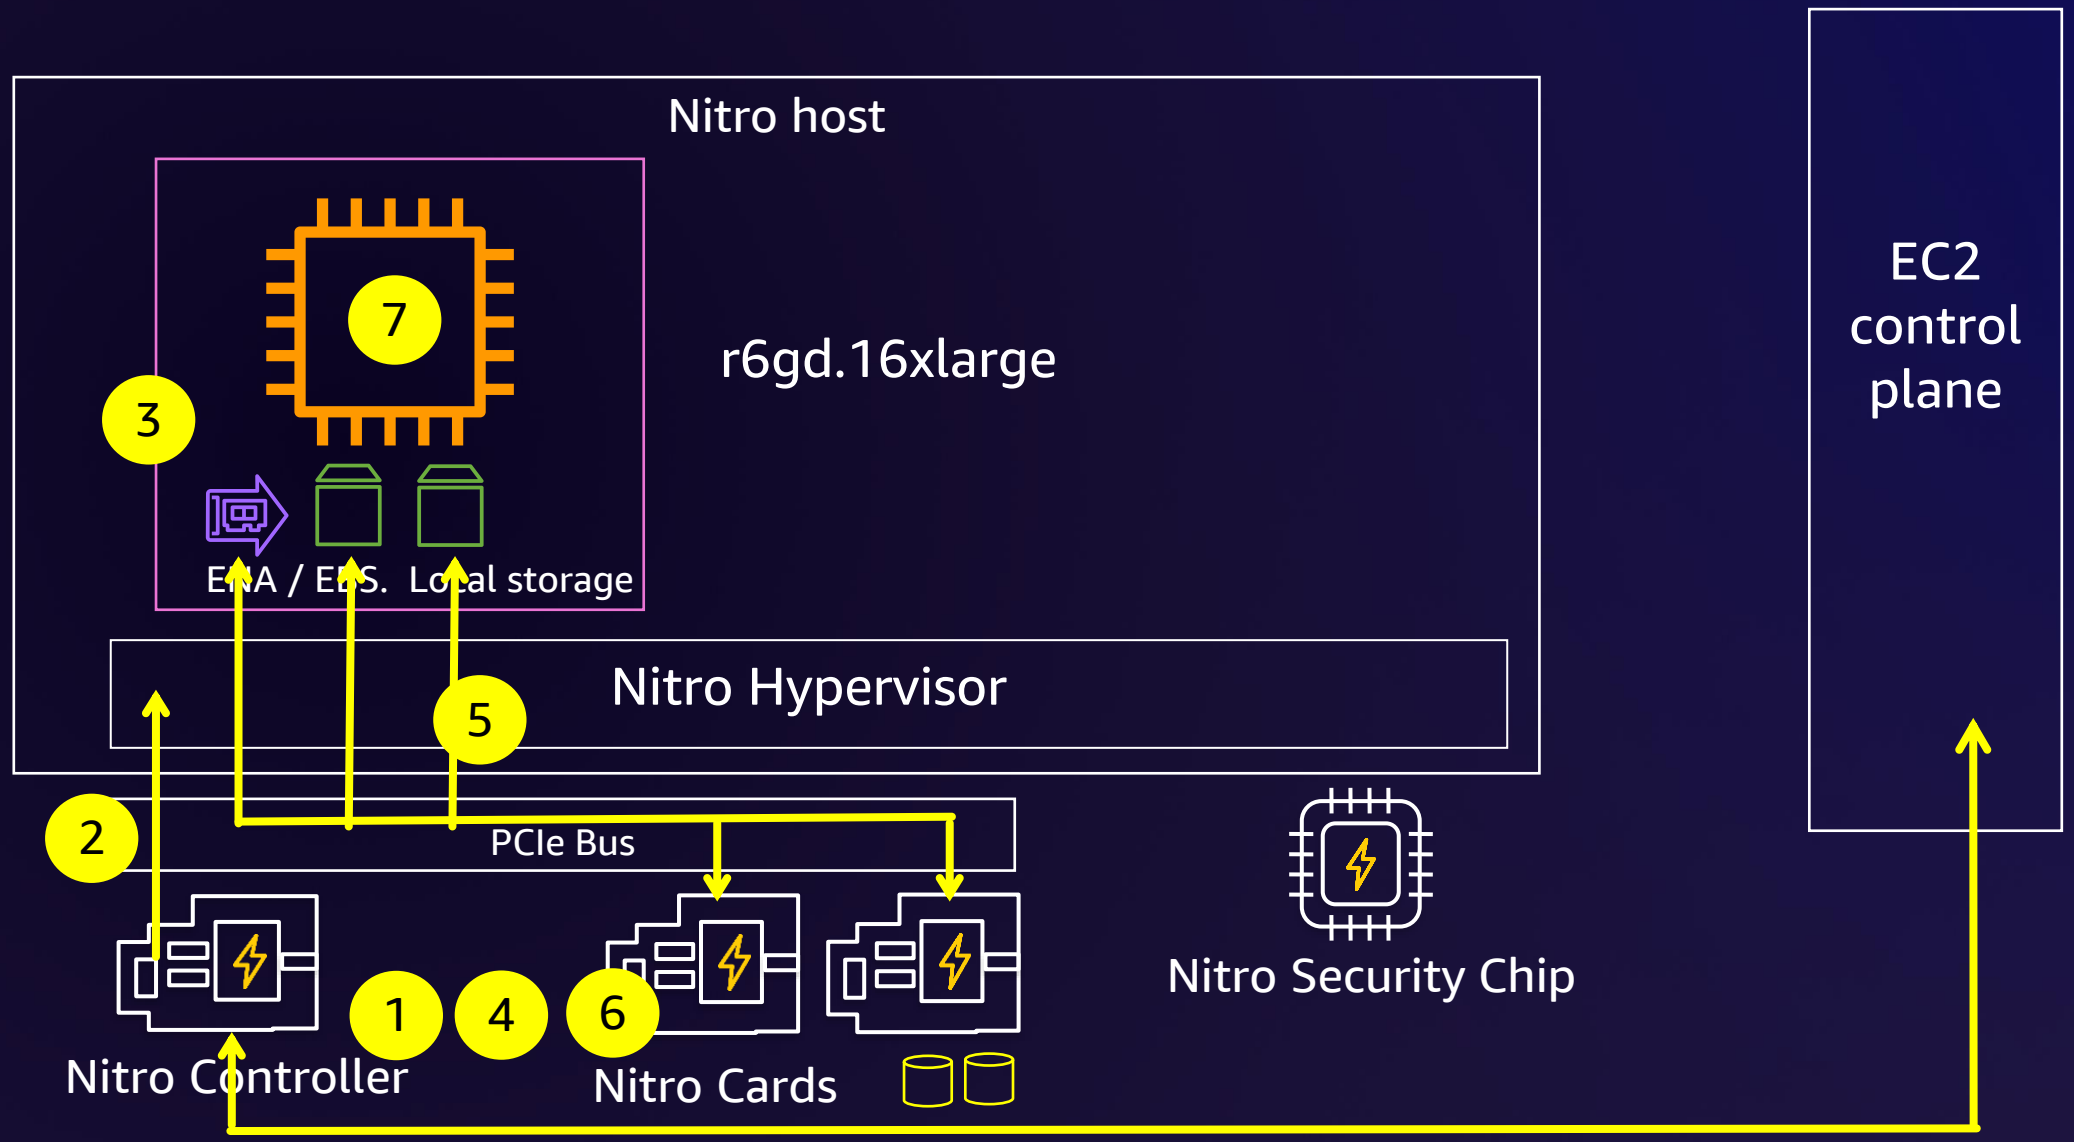

Nitro System EBS encryption

1. Customer uses IAM credentials authorized both for the EBS volume and for the AWS KMS key to request a volume attachment
2. … (see previous slide)
3. The Nitro Card received an encrypted copy of the volume key which is encrypted by a key present only in AWS KMS
4. An IAM forward access session is used to make a time-bound call to AWS KMS on behalf of the customer principal to obtain a narrowly scoped AWS KMS grant for the key

Nitro System EBS encryption

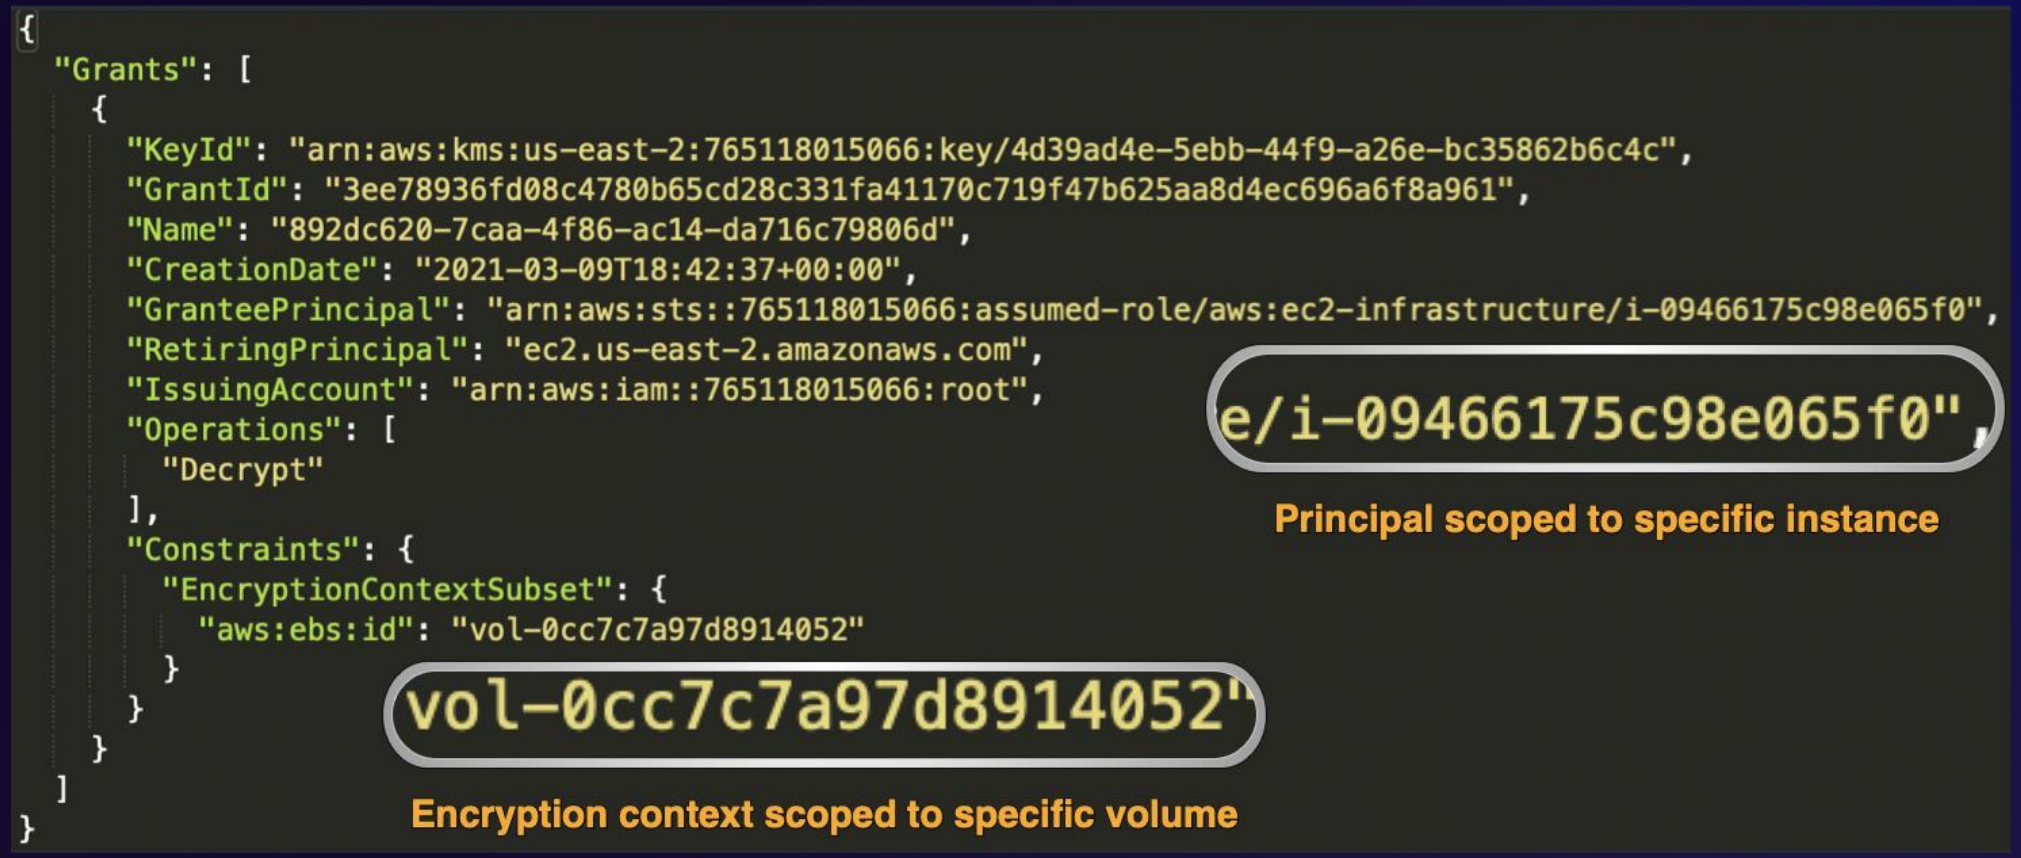

Nitro System EBS encryption

5. The Nitro Card uses the scoped down grant to send the encrypted volume key to AWS KMS for decryption
6. AWS KMS decrypts the volume key, re-encrypts the key using public key provided by the Nitro Card for which the corresponding private key is available only to that specific device
7. AWS KMS transmits the re-encrypted volume key back to the Nitro Card using TLS for a second defense-in-depth layer of encryption in transit
8. The Nitro Card terminates the TLS connection, decrypts the volume key and holds the plaintext key in its protected volatile memory
9. …

No AWS operator access to the Nitro System

- There is no operator access mechanism in the Nitro System design
- No SSH or general purpose access of any kind
- All Nitro operations are performed via secure, authenticated, authorized, logged & audited administrative APIs
- No APIs provide access to customer data

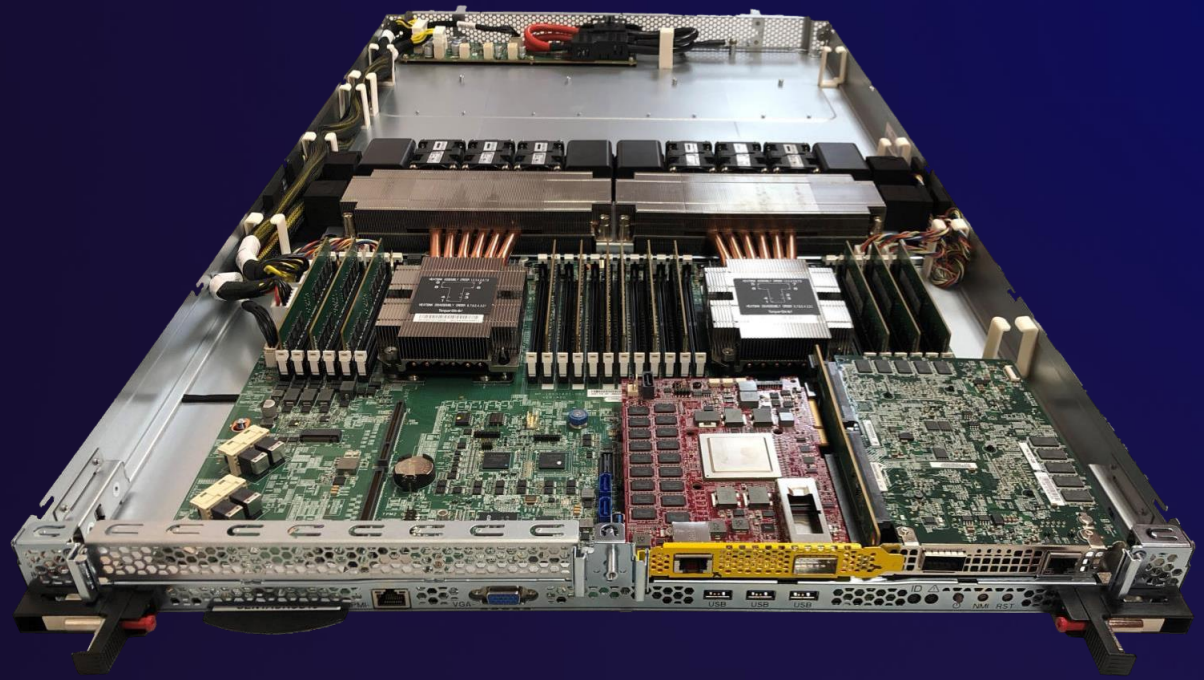

Nitro-based EC2 server

Additional assurance for the Nitro System

1. Nitro system security whitepaper

2. Independent third-party architecture review

3. Updated AWS Service Terms

Nitro security design documentation

New!

The Security Design of the AWS Nitro System

Publication date: **November 18, 2022** (*Document revisions* (p. 27))

Abstract

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides secure, resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers. The AWS Nitro System is the underlying platform for all modern EC2 instances. This whitepaper provides a detailed description of the security design of the Nitro System to assist you in evaluating EC2 for your sensitive workloads.

https://a.co/hYWhsH9

- Detailed review of the security design the three primary components of the AWS Nitro System
- Deep dive on the AWS Nitro System integrity protections, tenant isolation model, and no operator access design

Independent architecture review of the Nitro System

New!

- Assessment conducted by NCC Group in 2023
- The NCC report focuses on AWS claims on zero operator access in the Nitro System
- Outcome: "*no gaps in the Nitro System that would compromise these security claims*"

Updated AWS Service Terms to include Nitro System

New!

Our AWS Service Terms now include the following on the Nitro System:

"AWS personnel do not have access to Your Content on AWS Nitro System EC2 instances. There are no technical means or APIs available to AWS personnel to read, copy, extract, modify, or otherwise access Your Content on an AWS Nitro System EC2 instance or encrypted-EBS volume attached to an AWS Nitro System EC2 instance. Access to AWS Nitro System EC2 instance APIs – which enable AWS personnel to operate the system without access to Your Content – is always logged, and always requires authentication and authorization."

Live update of the Nitro System

Update processes do not relax security protocols or defenses

Both the Nitro Card firmware and the hypervisor are designed to be live-updatable (zero downtime for customer instances)

Eliminates the need for carefully balanced tradeoffs around updates yielding improved security outcomes

Change management for the Nitro System

- No operator, no matter how privileged, can push code to a production Nitro Host
- Multi-party review and approval, and staged rollouts in both testing and production environments
- Code reviews, security reviews, automated testing
- Then and only then are the binaries cryptographically signed by a private key which is only accessible through the automated pipeline and which logs all key signing activity
- We deploy conservatively and monitor for perturbations and automate any rollbacks as needed

Passive communications principal

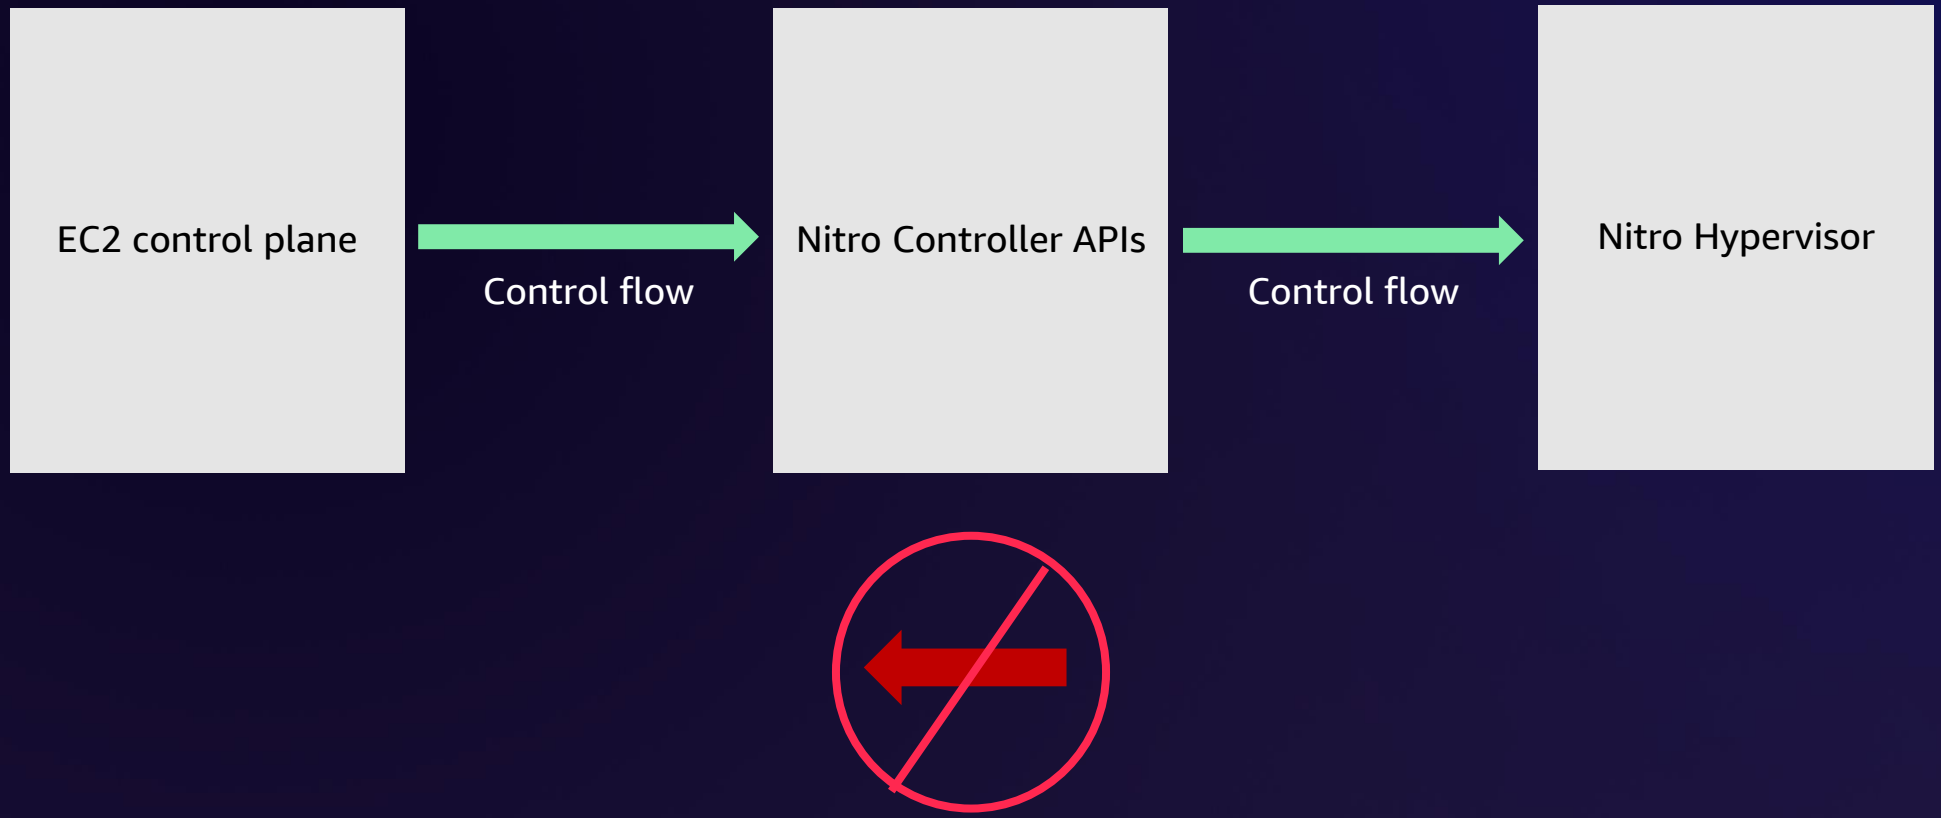

EC2 side channel protections

Always:

- No memory page coalescing
- No simultaneous sharing of threads on a CPU or low-level core attributes

In general:

- Instances are provided with CPU cores and memory which are pinned to that instance throughout its lifetime

Exception:

- Burstable instances are carefully designed to meet our high security bar for isolation

Confidential computing

PROTECTING CUSTOMER CODE AND SENSITIVE DATA IN USE

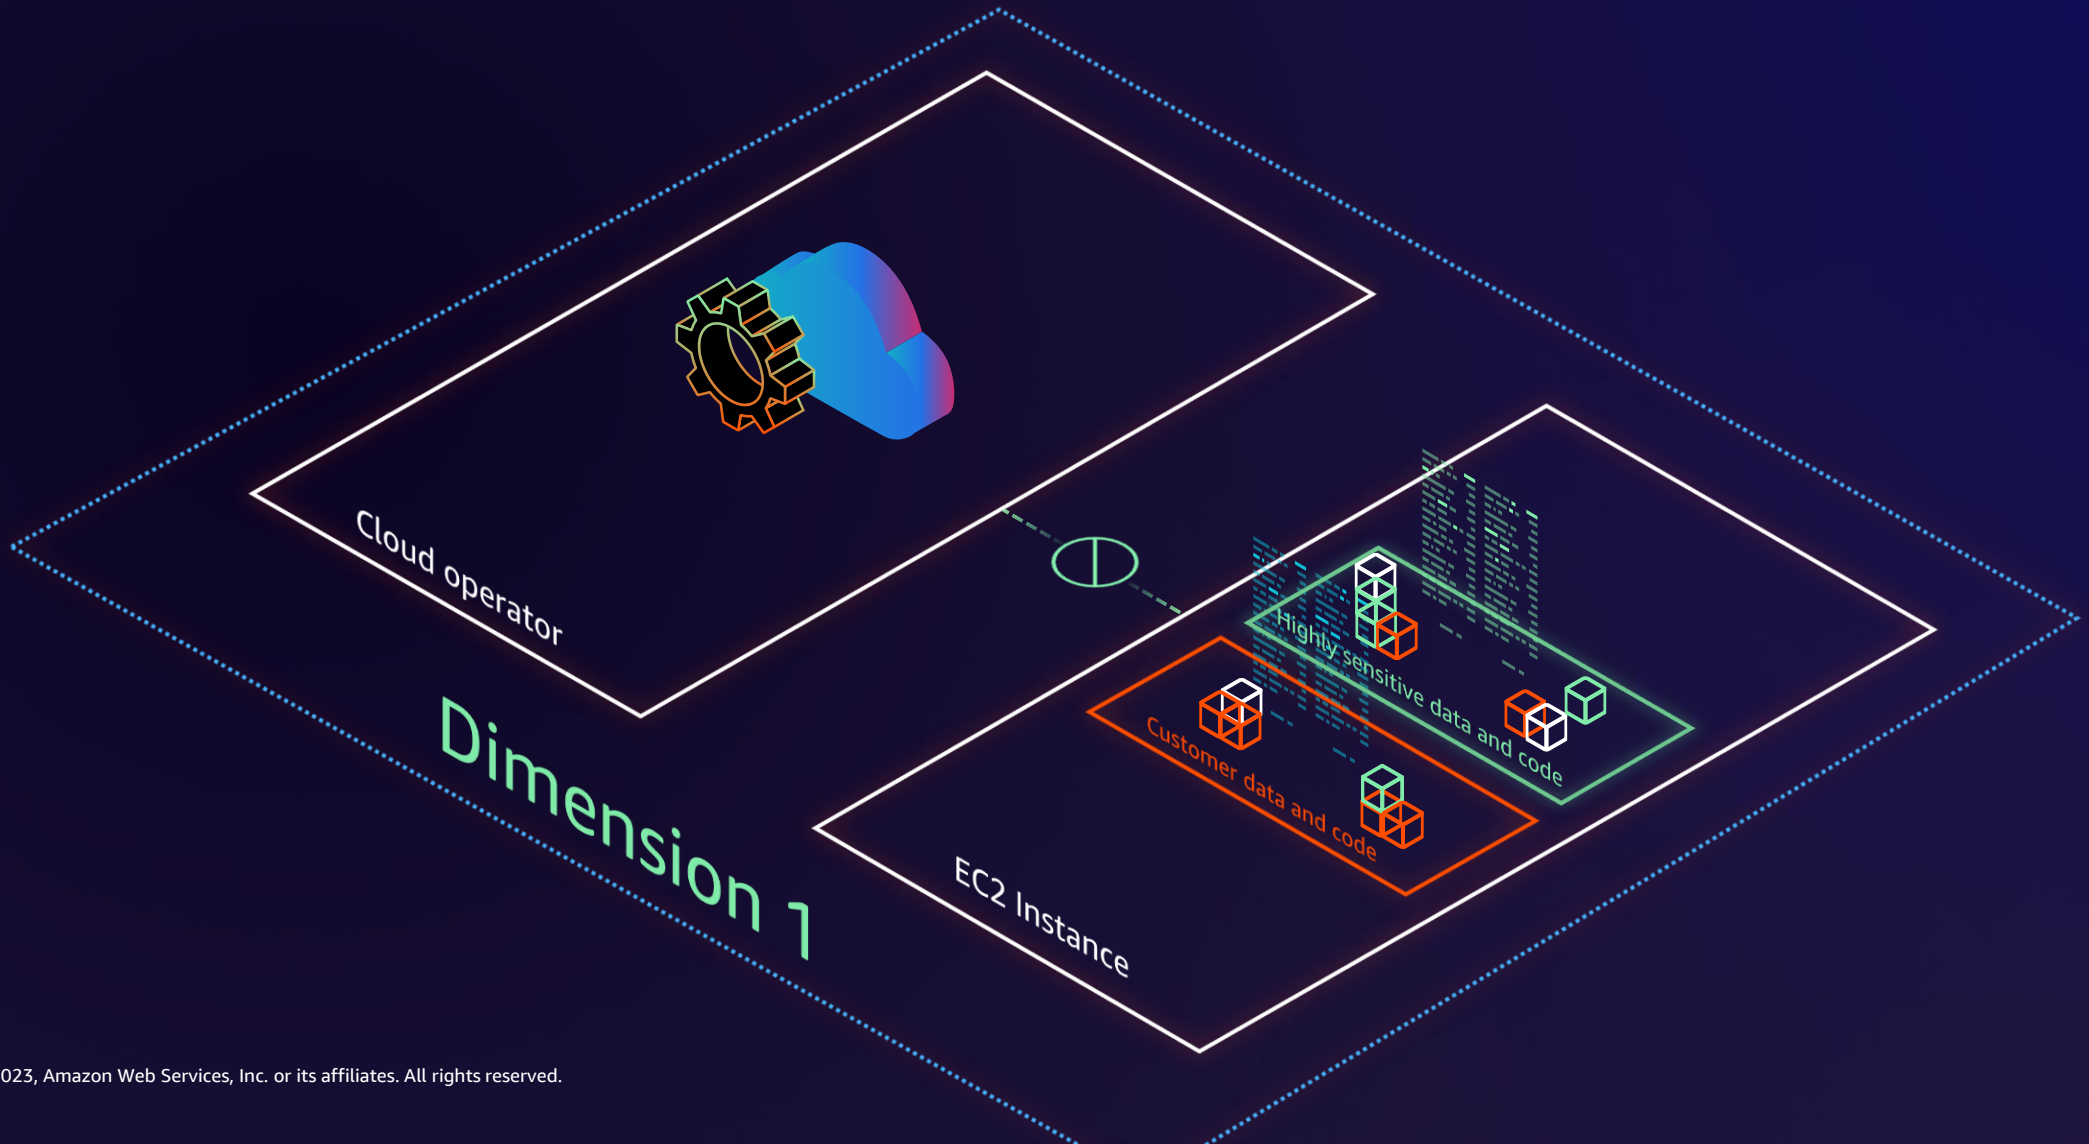

Confidential computing

PROTECTING CUSTOMER CODE AND SENSITIVE DATA IN USE

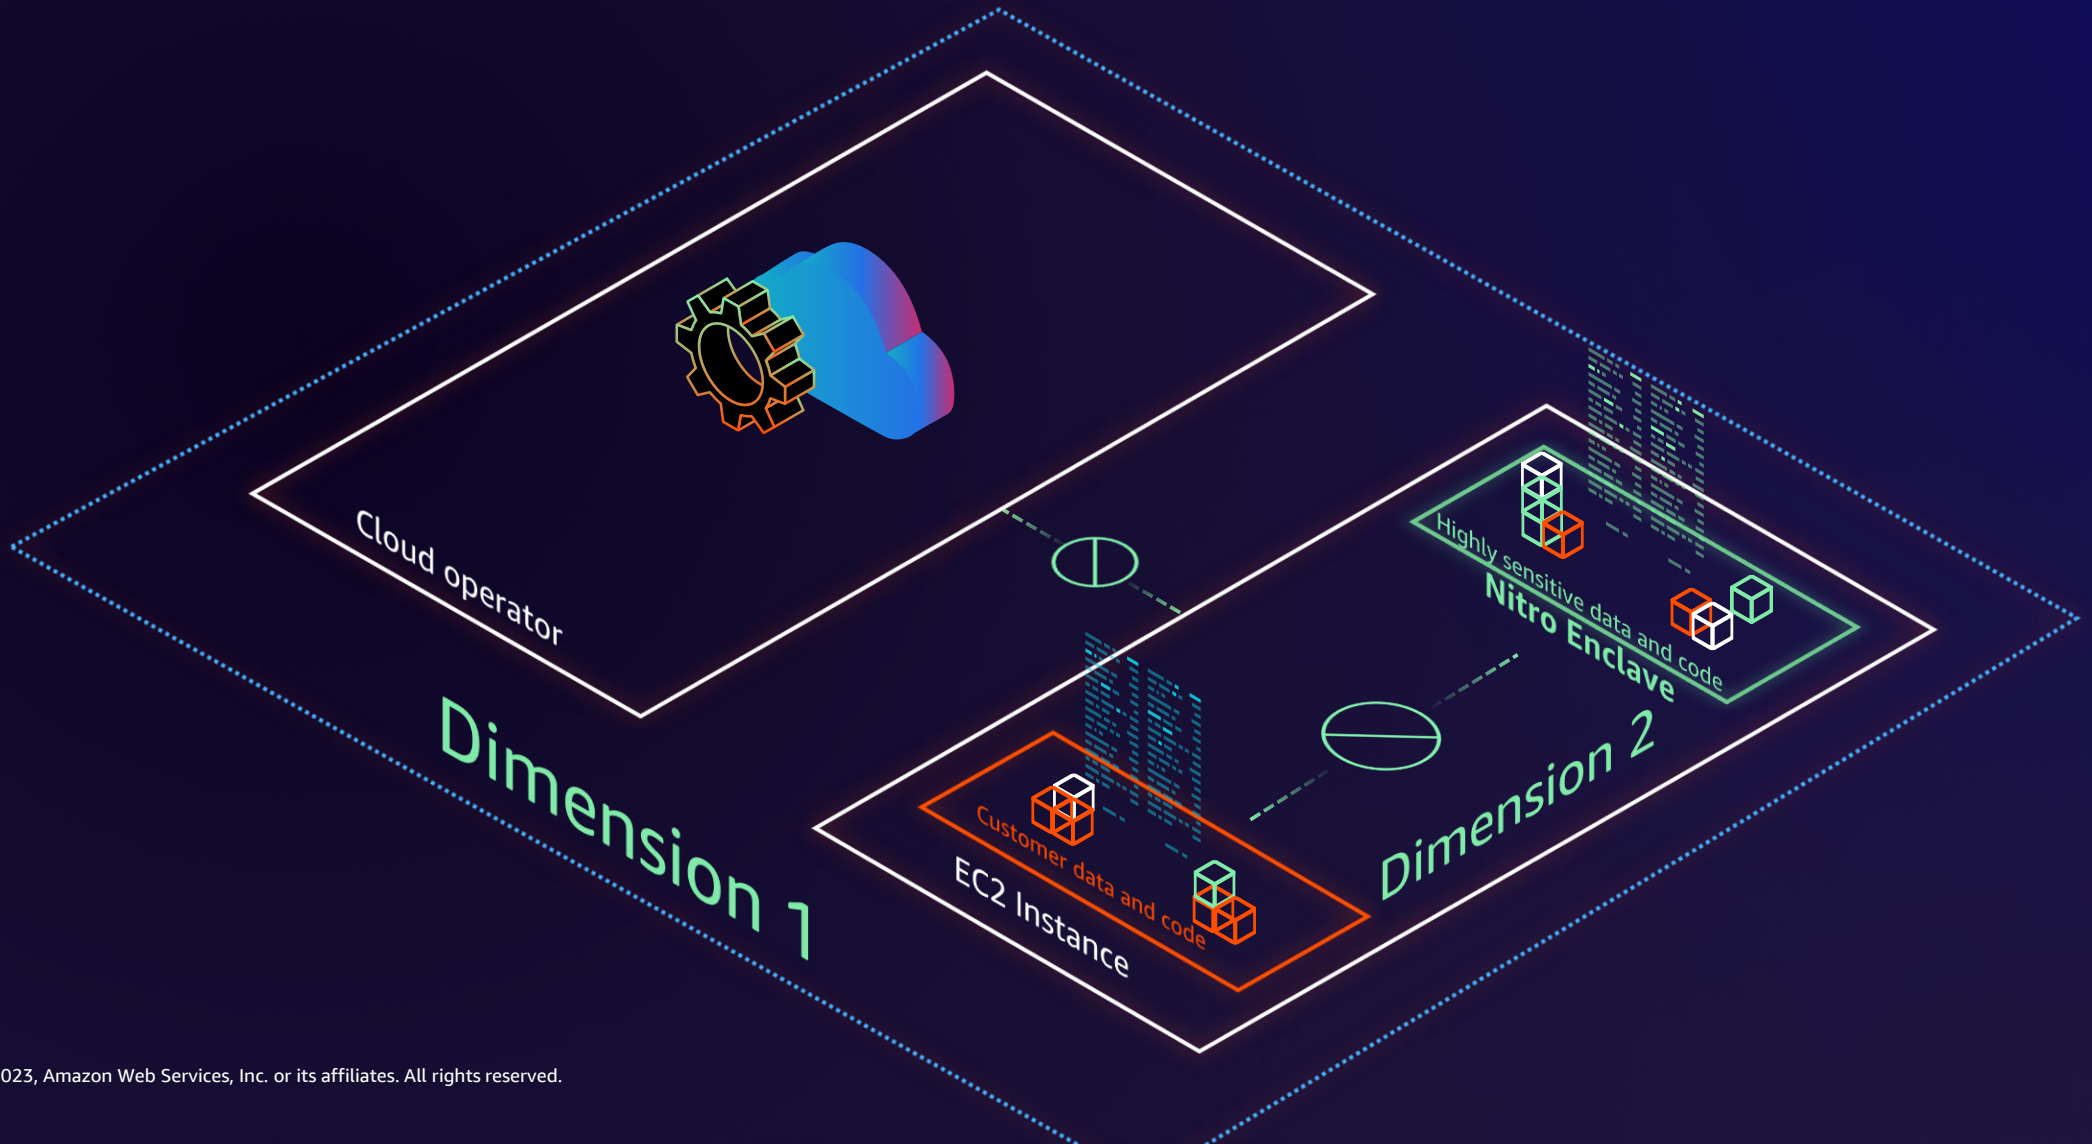

Confidential computing is the use of specialized hardware and associated firmware to protect customer code and data while in use

1. Protection of customer code and data from the operator of the underlying cloud infrastructure

AWS Nitro System

2. Protection and isolation among **related, cooperating workload components** – including protection of customer code and data from the customer's own operators

AWS Nitro Enclaves

AWS Nitro Enclaves

DIMENSION 2

For customers who want to take confidentiality a step further and isolate their highly sensitive data from the users, applications, and libraries on their EC2 instance in an isolated and cryptographically attested environment

Multi-party collaboration

TWO OR MORE PARTIES PROCESS SENSITIVE DATA WITHOUT GIVING ACCESS TO EACH OTHER

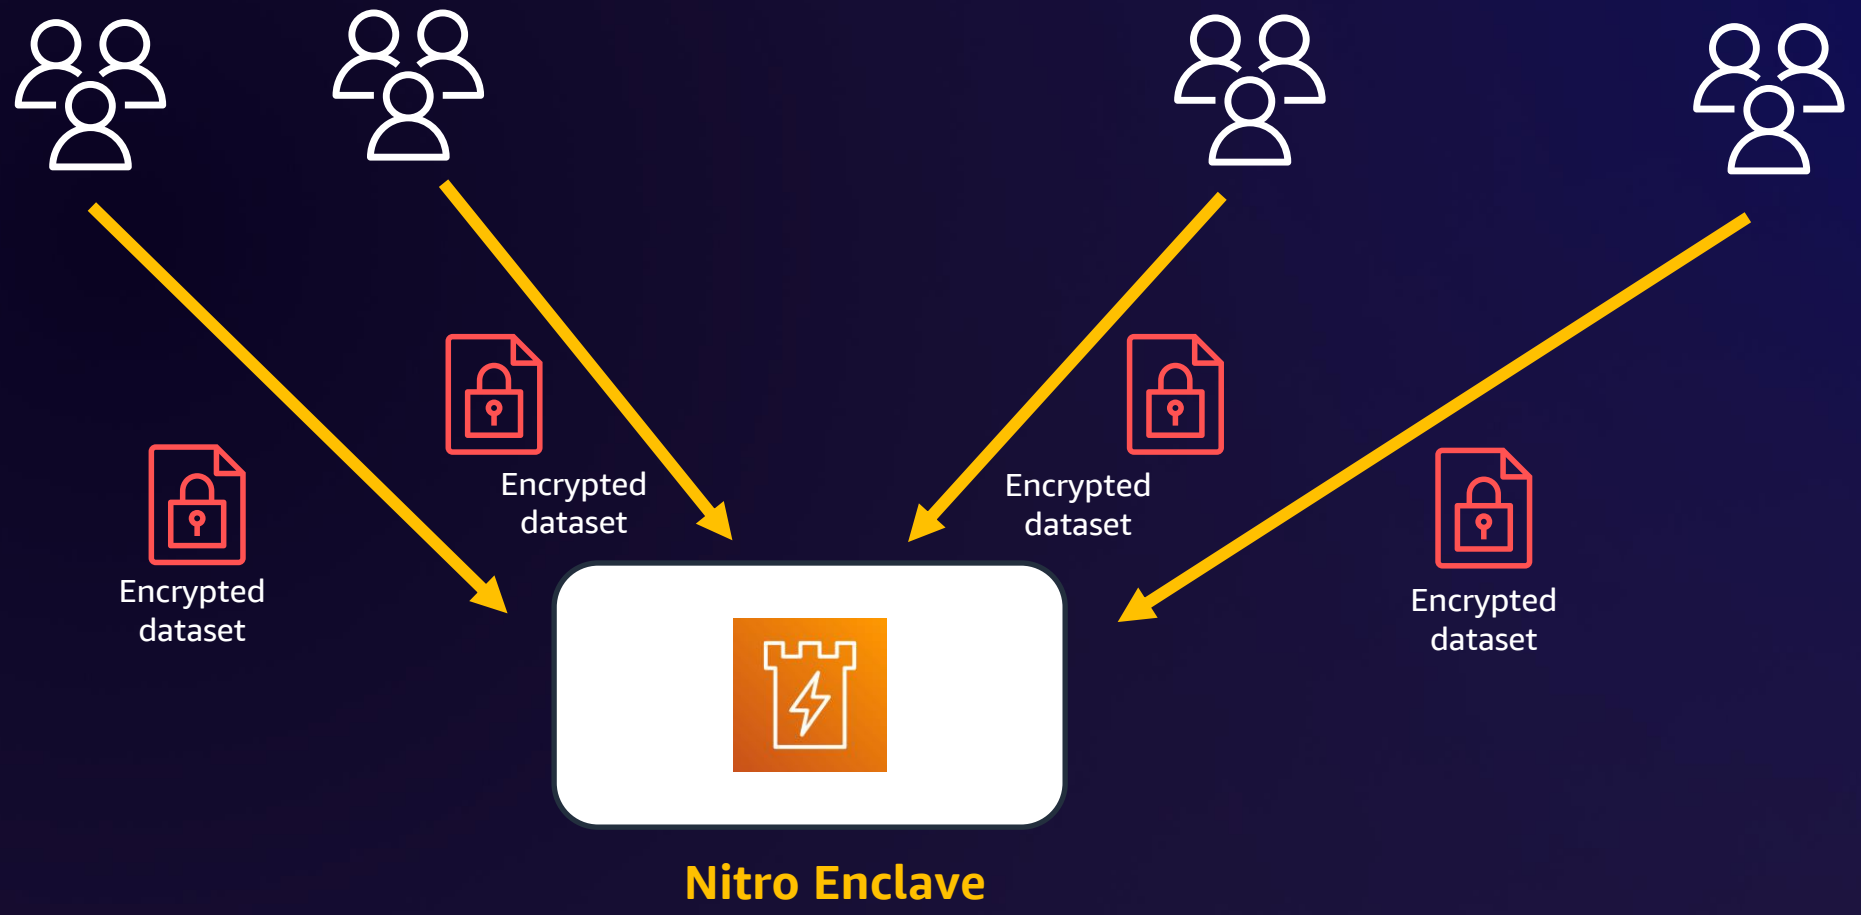

Powered by Nitro System confidential computing

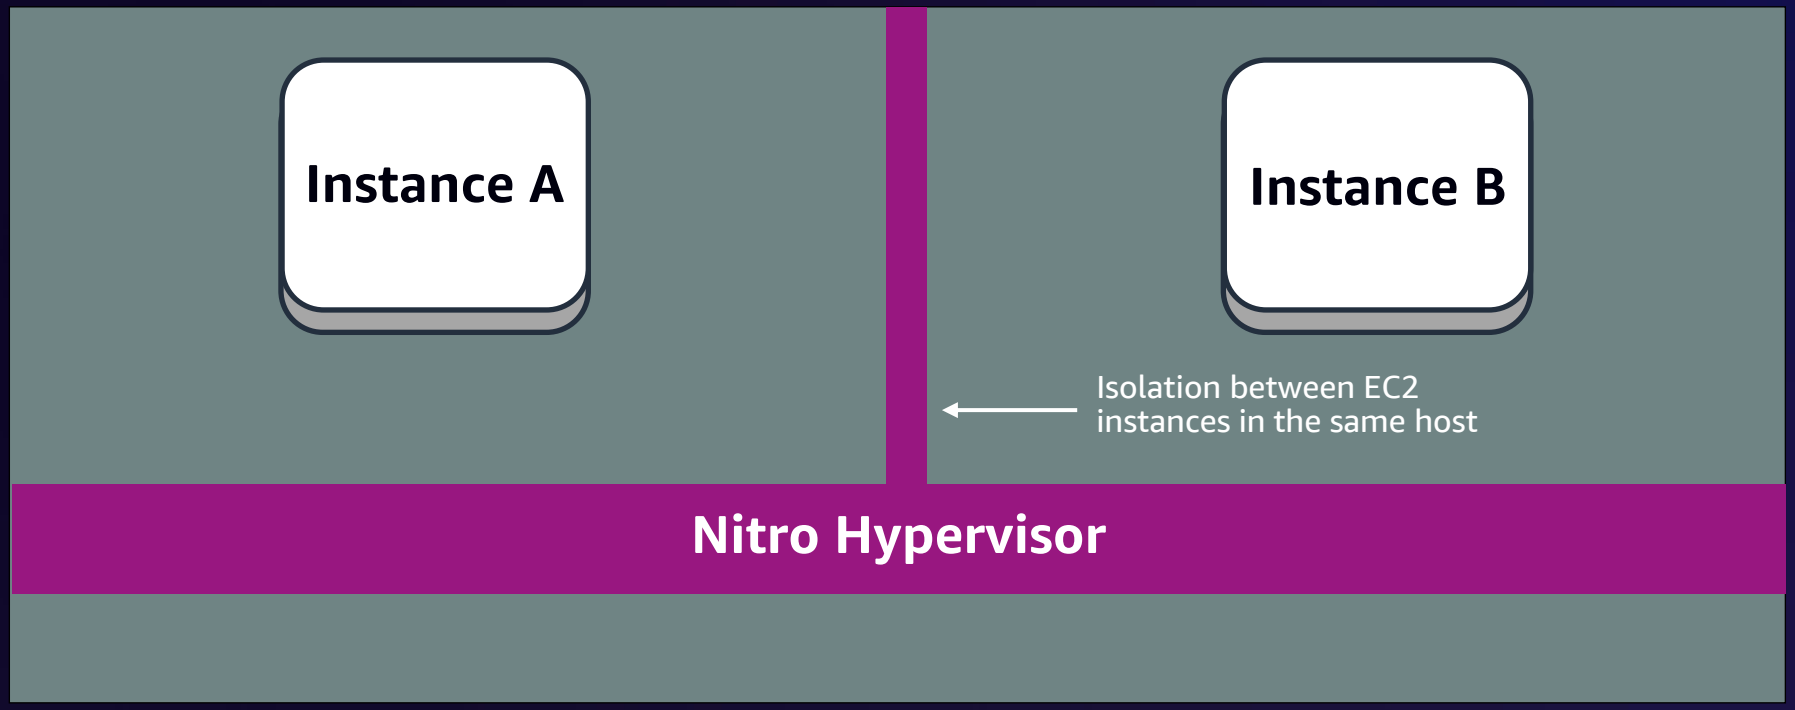

EC2 host

Powered by Nitro System confidential computing

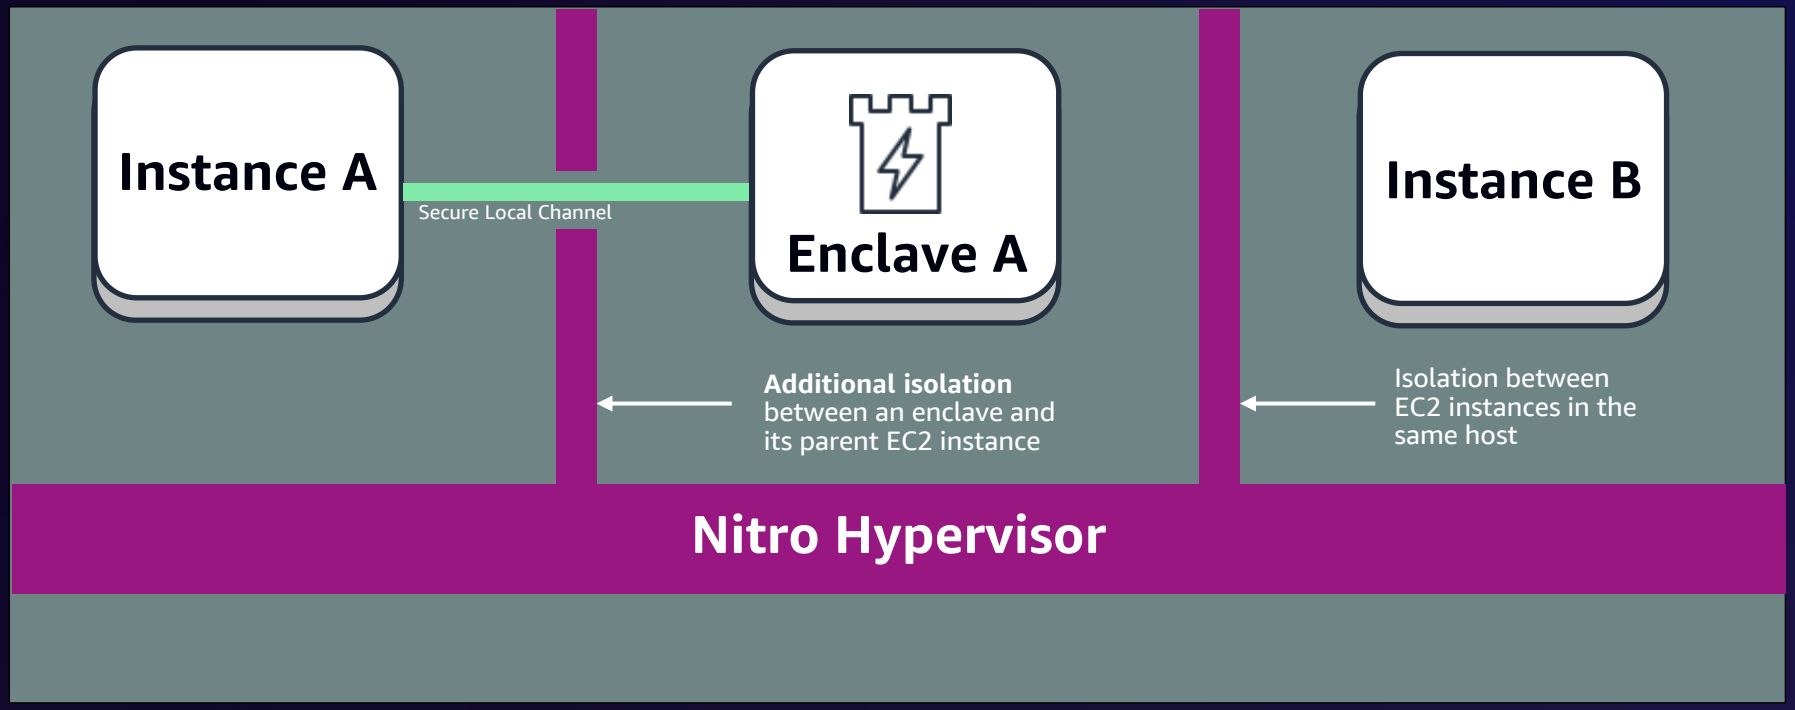

Broadest and deepest platform choice

Categories

General purpose + burstable

Compute-optimized

Memory-optimized

Storage-optimized

Accelerated computing

HPC-optimized

Capabilities

Choice of processor
(AWS, Intel, AMD, Apple)

Fast processors
(up to 4.5 GHz)

High memory footprint
(up to 24 TiB)

Instance storage
(HDD and NVMe)

Accelerated computing
(GPUs, FPGA, ASIC)

Networking
(up to 1600 Gbps)

Bare Metal

Size
(<1vCPU to 448 vCPU)

Options

Amazon EBS

Windows, Linux, UNIX, macOS

More than

600

Instance types

for virtually every workload and business need

Increased pace of innovation

AMAZON EC2 INSTANCES TIMELINE

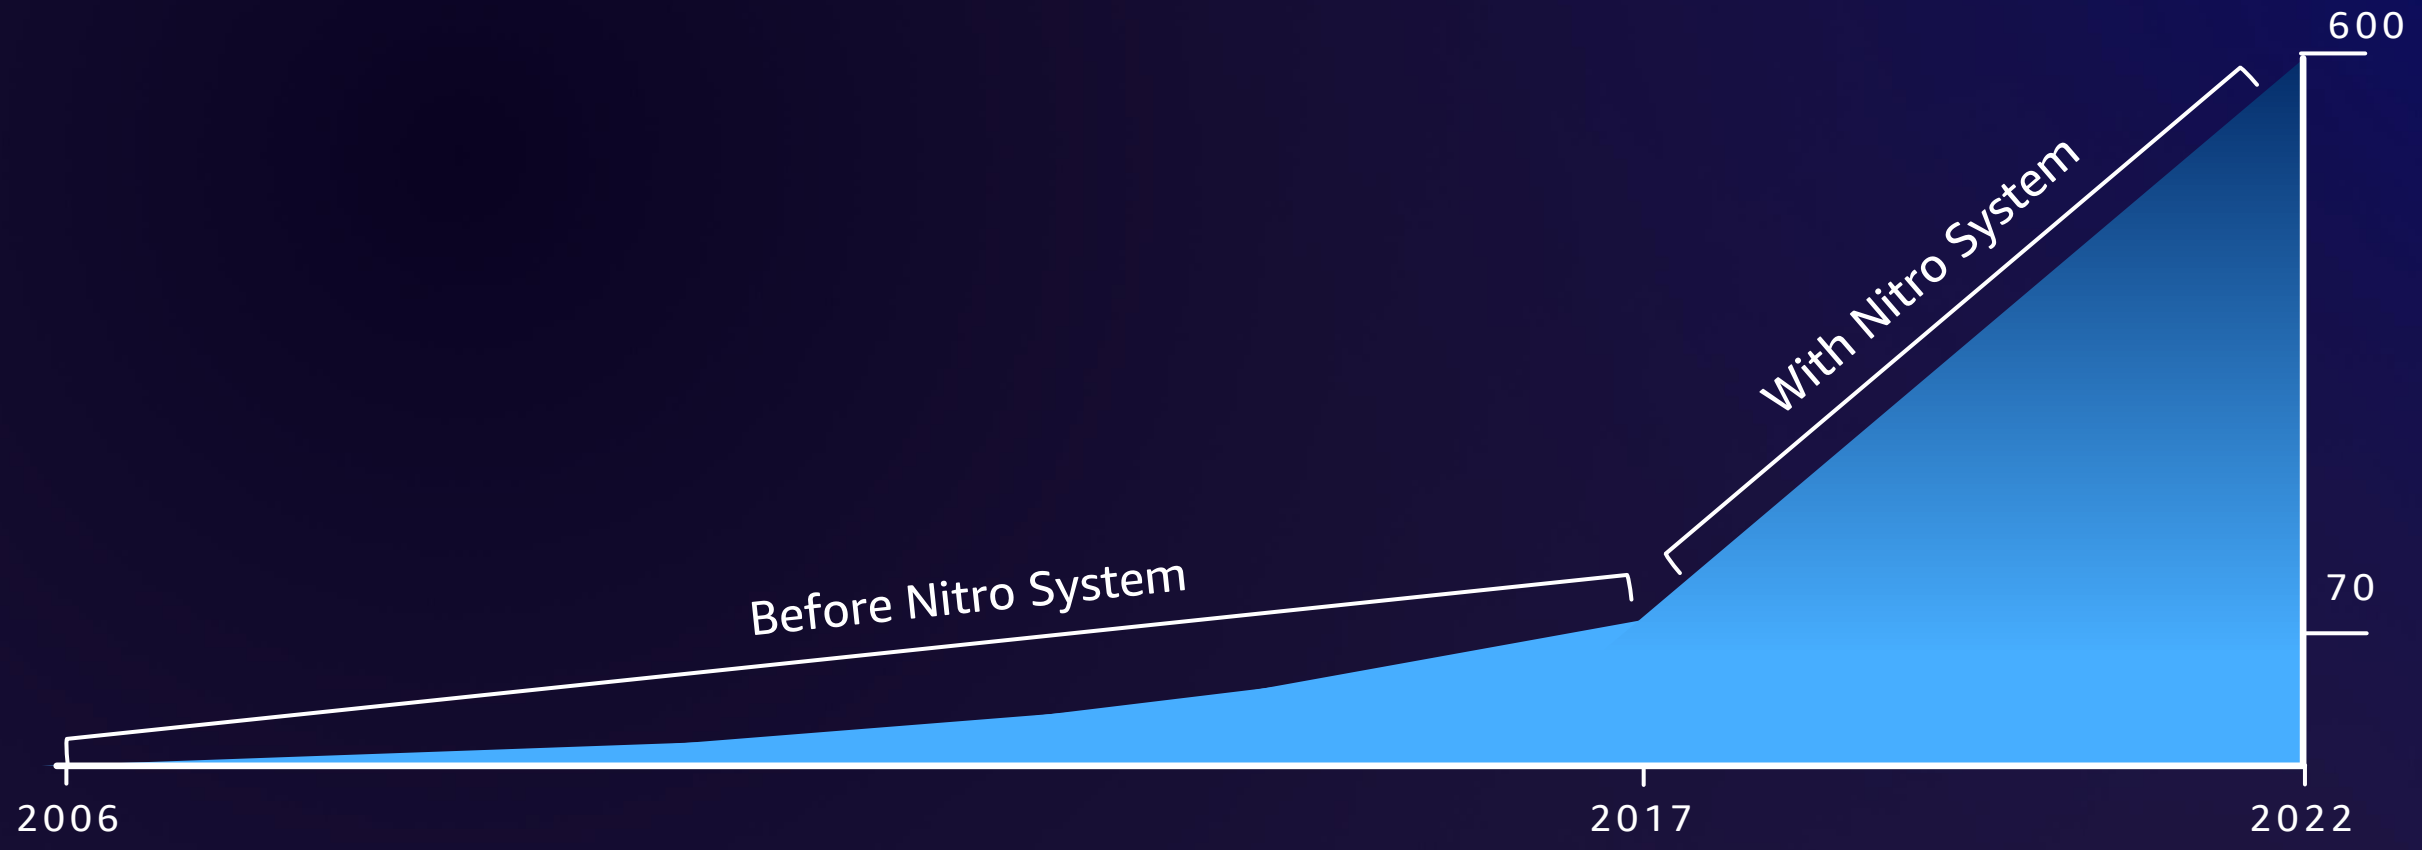

Firecracker

Firecracker is an open source virtualization technology that is purpose-built for creating and managing secure, multi-tenant container and function-based services. It supports workloads in services such as AWS Lambda, AWS Fargate, and Amazon Athena.

Security from the ground up
KVM-based virtualization

Speed by design
<125 ms to launch
150 microVMs per second/host

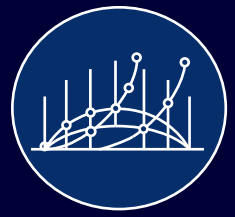

Scale and efficiency
<5 MB memory footprint per microVM

UEFI Secure Boot

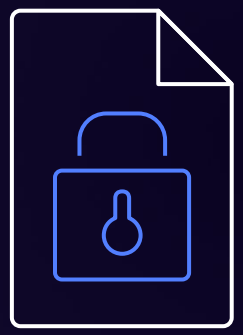

UEFI Secure Boot flow ensures that the bootloader is properly signed by a known authority

Validate the signed bootloader (for example, Grub2) against certificates stored in UEFI

Fall back to backup bootloader or stop if validation fails

NitroTPM

A trusted platform module.

Conforms to the industry standard TPM 2.0 specification.

Software compatibility. Makes it easier for customers to migrate applications which require TPMs to EC2.

Provides capabilities like attestation of system state, storing and generating cryptographic data, and proving platform identity to your EC2 instance.

Nitro System security

NITRO IS THE FOUNDATION FOR INNOVATIONS IN CONFIDENTIALITY AND PRIVACY

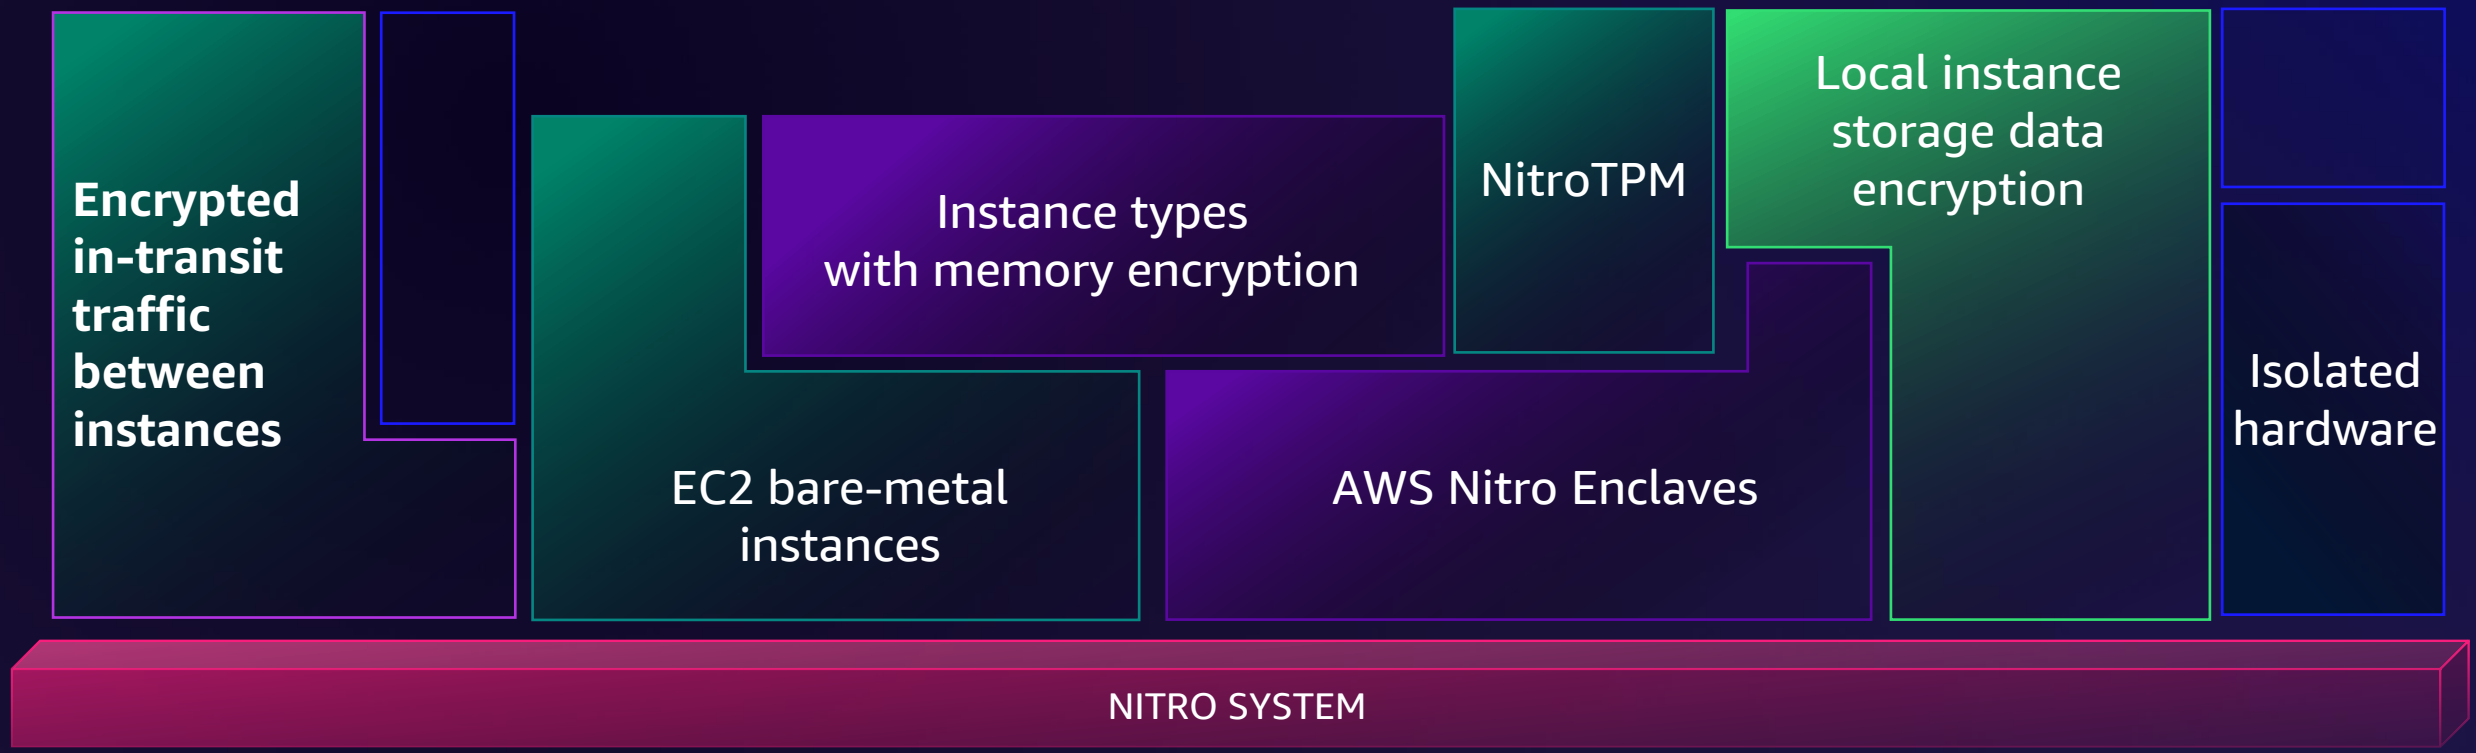

The Nitro System is the foundation for AWS

RAISING THE BAR ON SECURITY THROUGH PURPOSE-BUILT HARDWARE, FIRMWARE, AND SOFTWARE

- All EC2 instance types released since 2018 are powered by the Nitro System – On its 5th generation of custom silicon
- Secure boot process based on a hardware root of trust ensures that every component is signed and every operation is pre-vetted for safety
- Every critical element of Nitro system can be live-updated without impacting customer workloads
- Enables hardware-accelerated transparent encryption of storage, networking, and memory
- Zero operator access through confidential computing
- Robust and conservative tenant isolation
- Mainboard is isolated from the physical network
- An enabler of security innovation

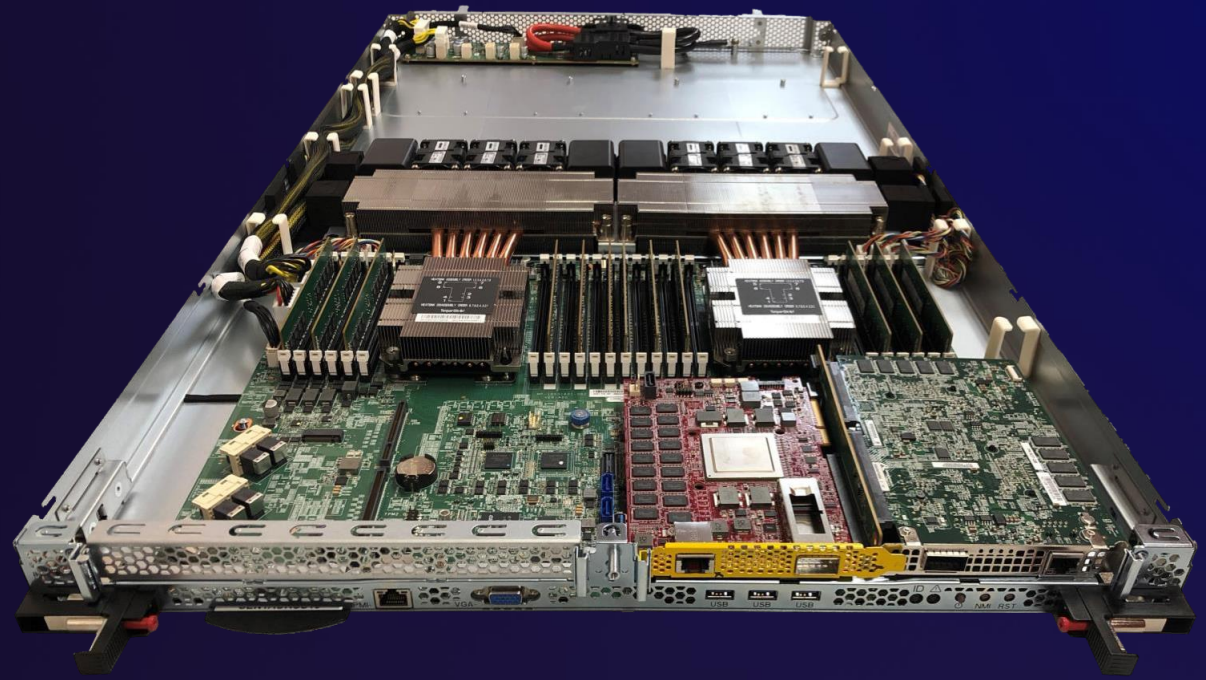

Nitro-based EC2 server

Thank you!

Please complete the session survey in the mobile app

J.D. Bean

Principal Security Architect, EC2

AWS