UCLOUD优刻得 | intel

云原生技术生态和趋势及英特尔如何助力云原生技术

徐亮厚　英特尔云计算解决方案架构师

UCLOUD优刻得 | intel

1. 云原生产业市场分析
2. 云原生技术生态和趋势
3. 英特尔技术和方案助力云原生技术落地

UCLOUD优刻得 | intel

云计算随云原生技术演进

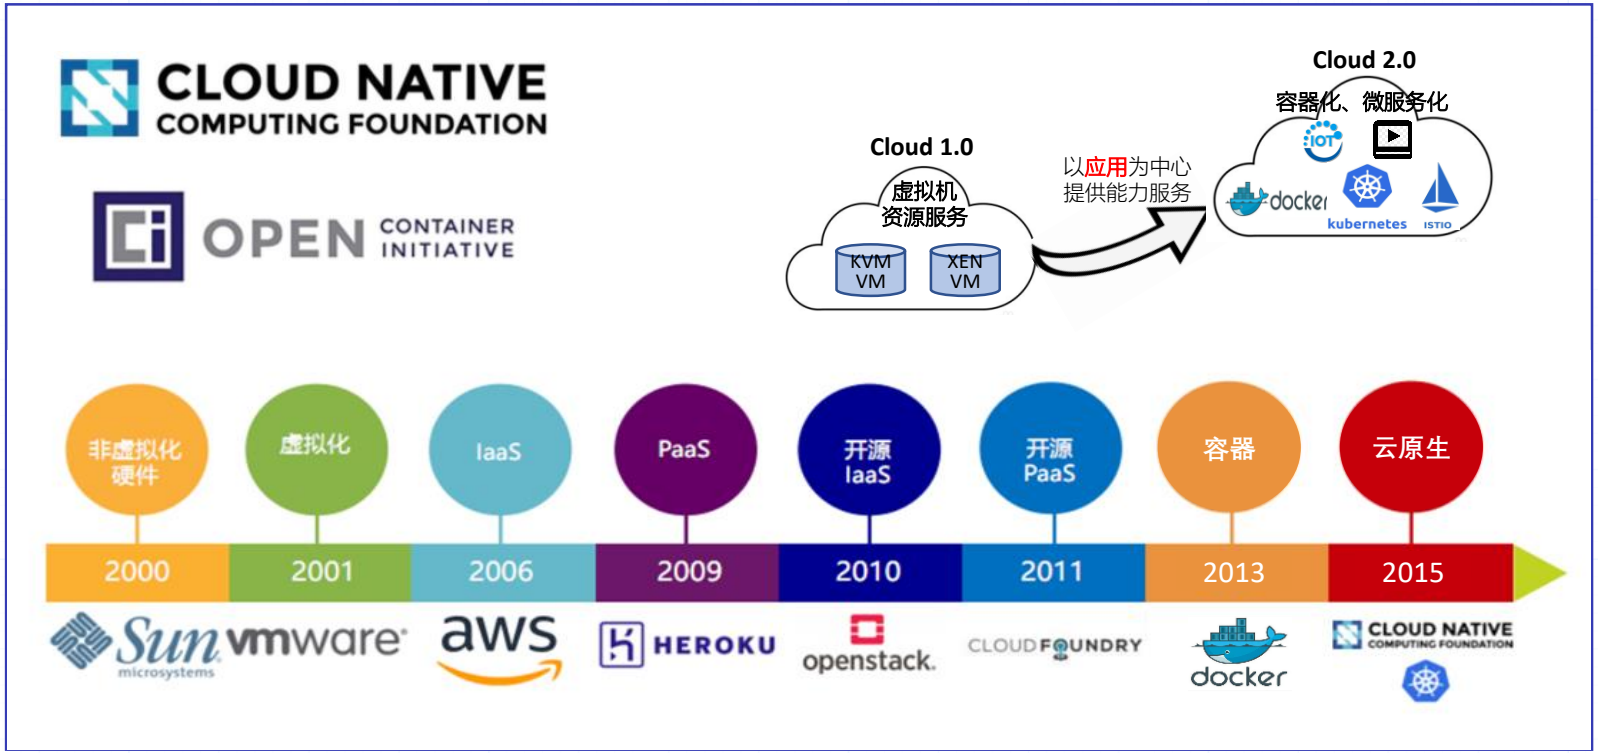

UCLOUD优刻得 | intel

IT基础设施不断演进…

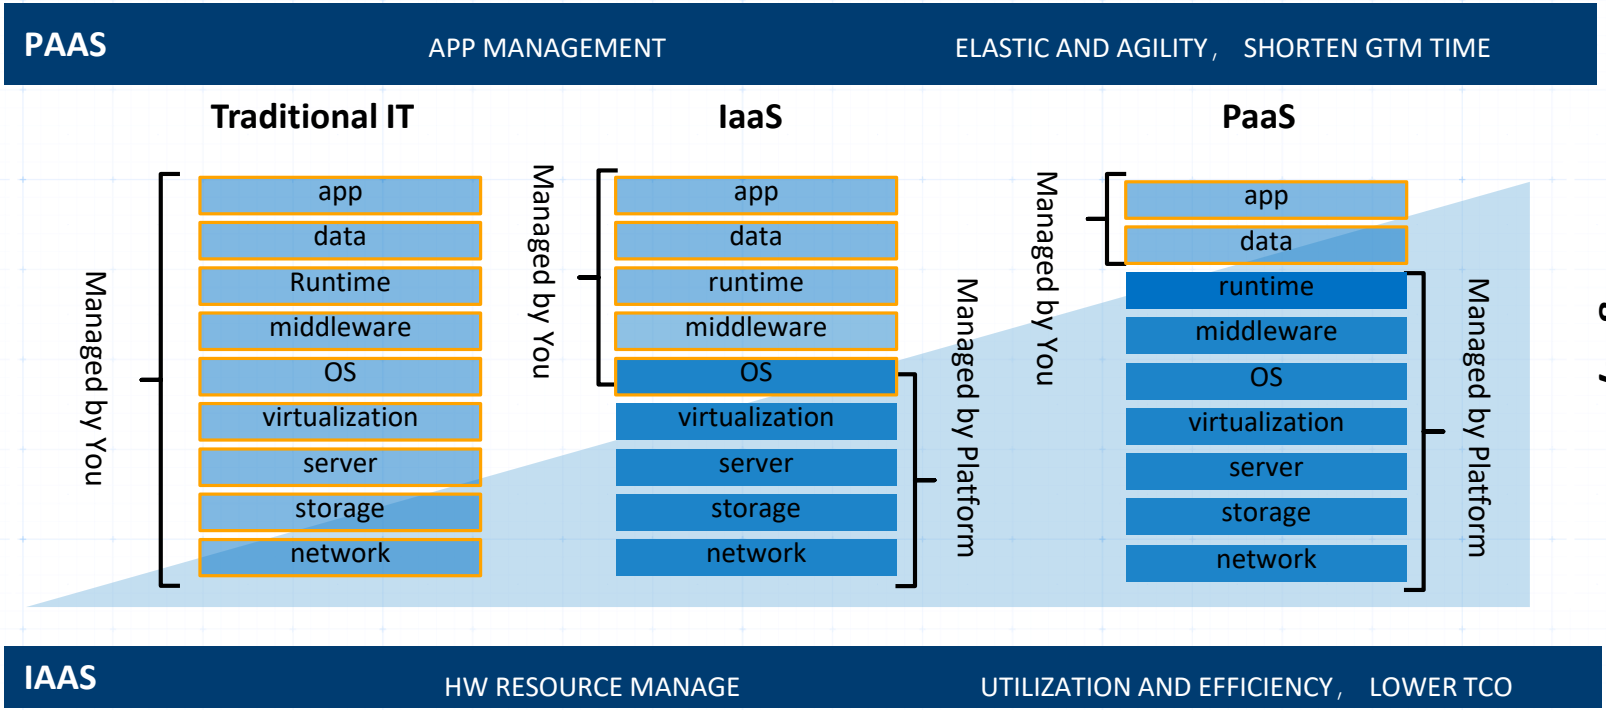

UCLOUD优刻得 | intel

…演进到云原生

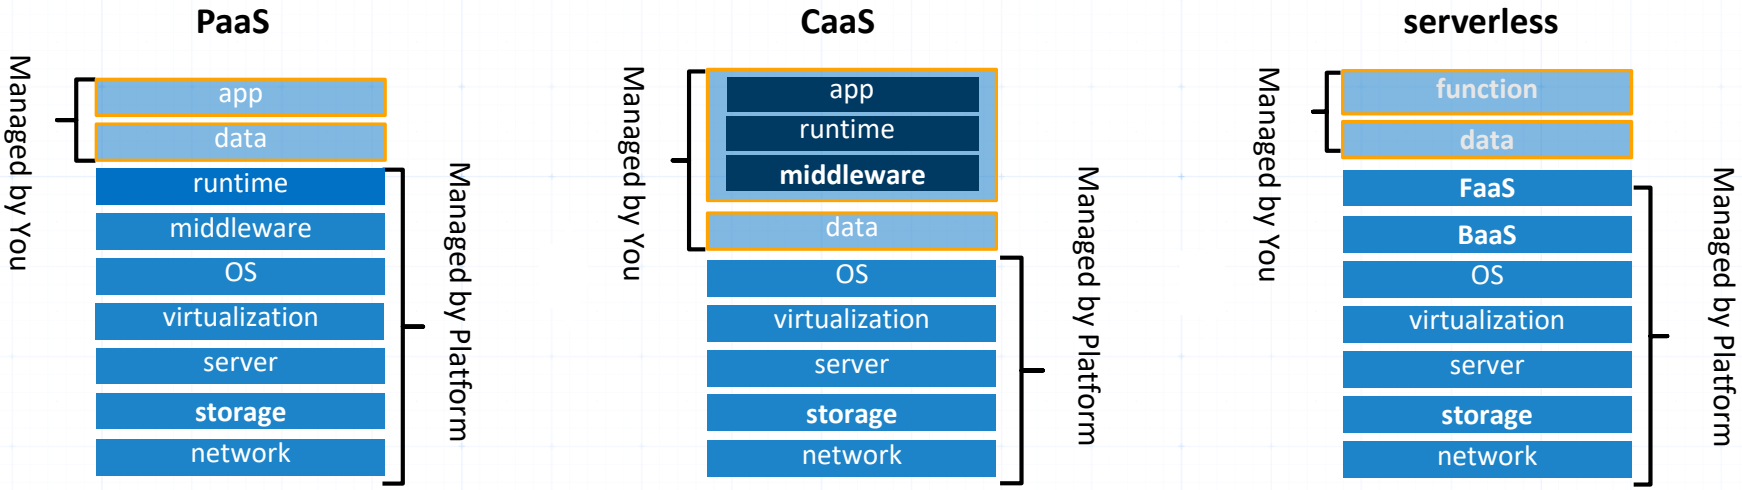

Cloud native

LIGHTWEIGHT

Low resource consume
Fast startup

AGILITY

Repeatable
Build once, run anywhere

DEVOPS

Fast response time
Improve QoS

SCALE

Micro Service
Hyper Scale

UCLOUD优刻得 | intel

应用架构演进使用云原生技术

Bring agility, Fast GTM, Lower TCO

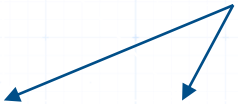

intel xeon PLATINUM 英特尔® 至强® 云创新基石

UCLOUD优刻得 | intel

云原生技术支撑企业数字化转型

企业使用云原生技术在现代、动态的环境（如公有、私有和混合云）中构建和运行可伸缩的应用

intel xeon PLATINUM 英特尔® 至强® 云创新基石

UCLOUD优刻得 | intel

用户正在接受云原生

Enterprise customer starting to adopt cloud native

CHINA LIFE 中国人寿

- Starting from 2017
- Production system
- Running 15 MC system
- 174+ servers, 677+ services, 4827+ containers

一汽-大众 FAW - VOLKSWAGEN

- Staring from 2018
- Partner with Tencent (Alauda)
- 50% new project base on Cloud Native
- 8+ Kubernetes clusters
- 2000+ containers

国家医疗保障局 National Healthcare Security Administration

- Founded in 2018
- Build their platform base on Cloud Native
- Request all venders using Micro Service

中国PaaS产业市场趋势

中国PaaS市场规模及增速，2018H2 - 2020H2

Source: IDC Semiannual Public Cloud Services Tracker, 2020H2

2020年下半年，中国PaaS市场规模达到14.2亿美金，同比增长61%，尽管增速持续放缓，但PaaS仍旧是中国公有云中增速最快的子市场

市场亮点和增长驱动因素

- 厂商加速布局云原生领域，加快完善产品功能和服务
- 低代码开发成为新市场热点，新兴云服务供应商不断涌现，如：简道云、氚云、炎黄盈动
- 传统应用厂商与SaaS厂商持续加强建设PaaS层服务，满足客户对应用系统的灵活、敏捷和个性化的使用需求；当前，与公有云厂商更多的是合作而非竞争
- 2020年以来，疫情加速了企业数字化转型的步伐，企业业务向云迁移和和业务创新的速度加快，拉动了大数据、数据库、中间件、人工智能等服务的需求

中国云原生产业市场趋势

- 容器持续迅猛增长. 68%的机构在生产中使用容器，比去年增长了39%，相比2年前增长240%.
- Kubernetes 已经无处不在，生产中使用 Kubernetes 的比例已从去年的72% 增长到了82%.
- 自动化程度持续提高. 30% 的机构生产使用自动化发布周期，73% 运行 CI/CD 流水线.
- 可观察性工具被广泛使用。95% 的机构使用监控工具，94% 使用日志分析系统，85% 使用分布式追踪.

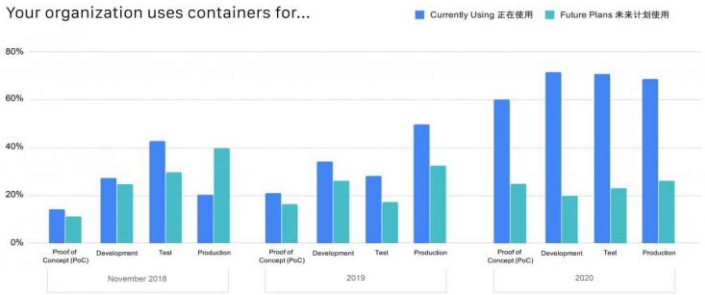

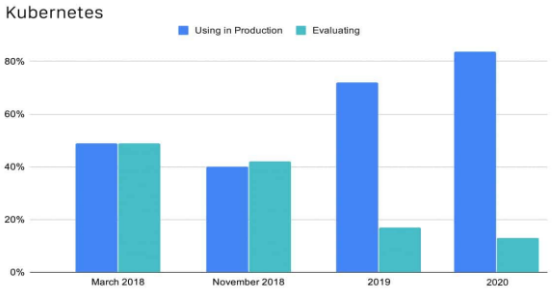

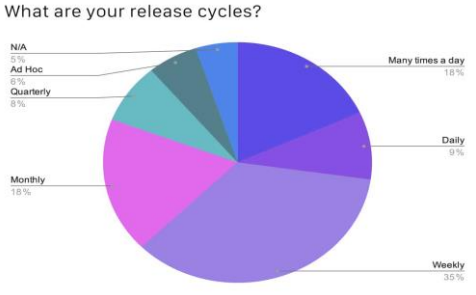

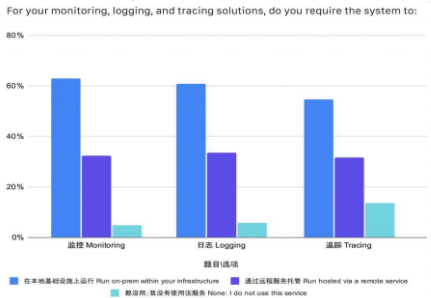

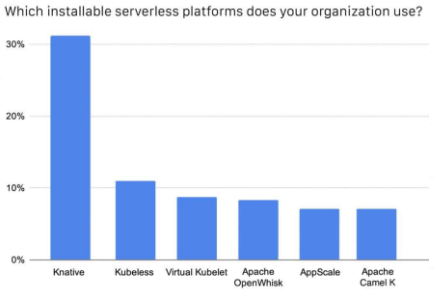

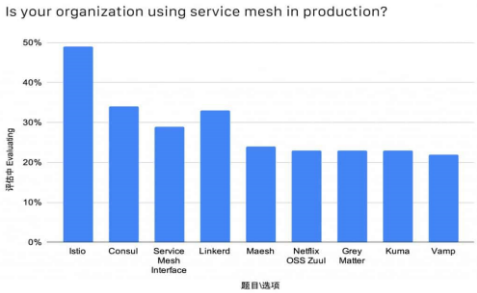

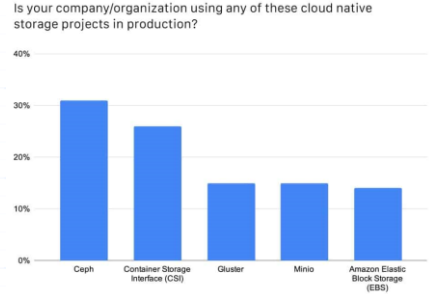

UCLOUD优刻得 | intel

云原生技术生态图景

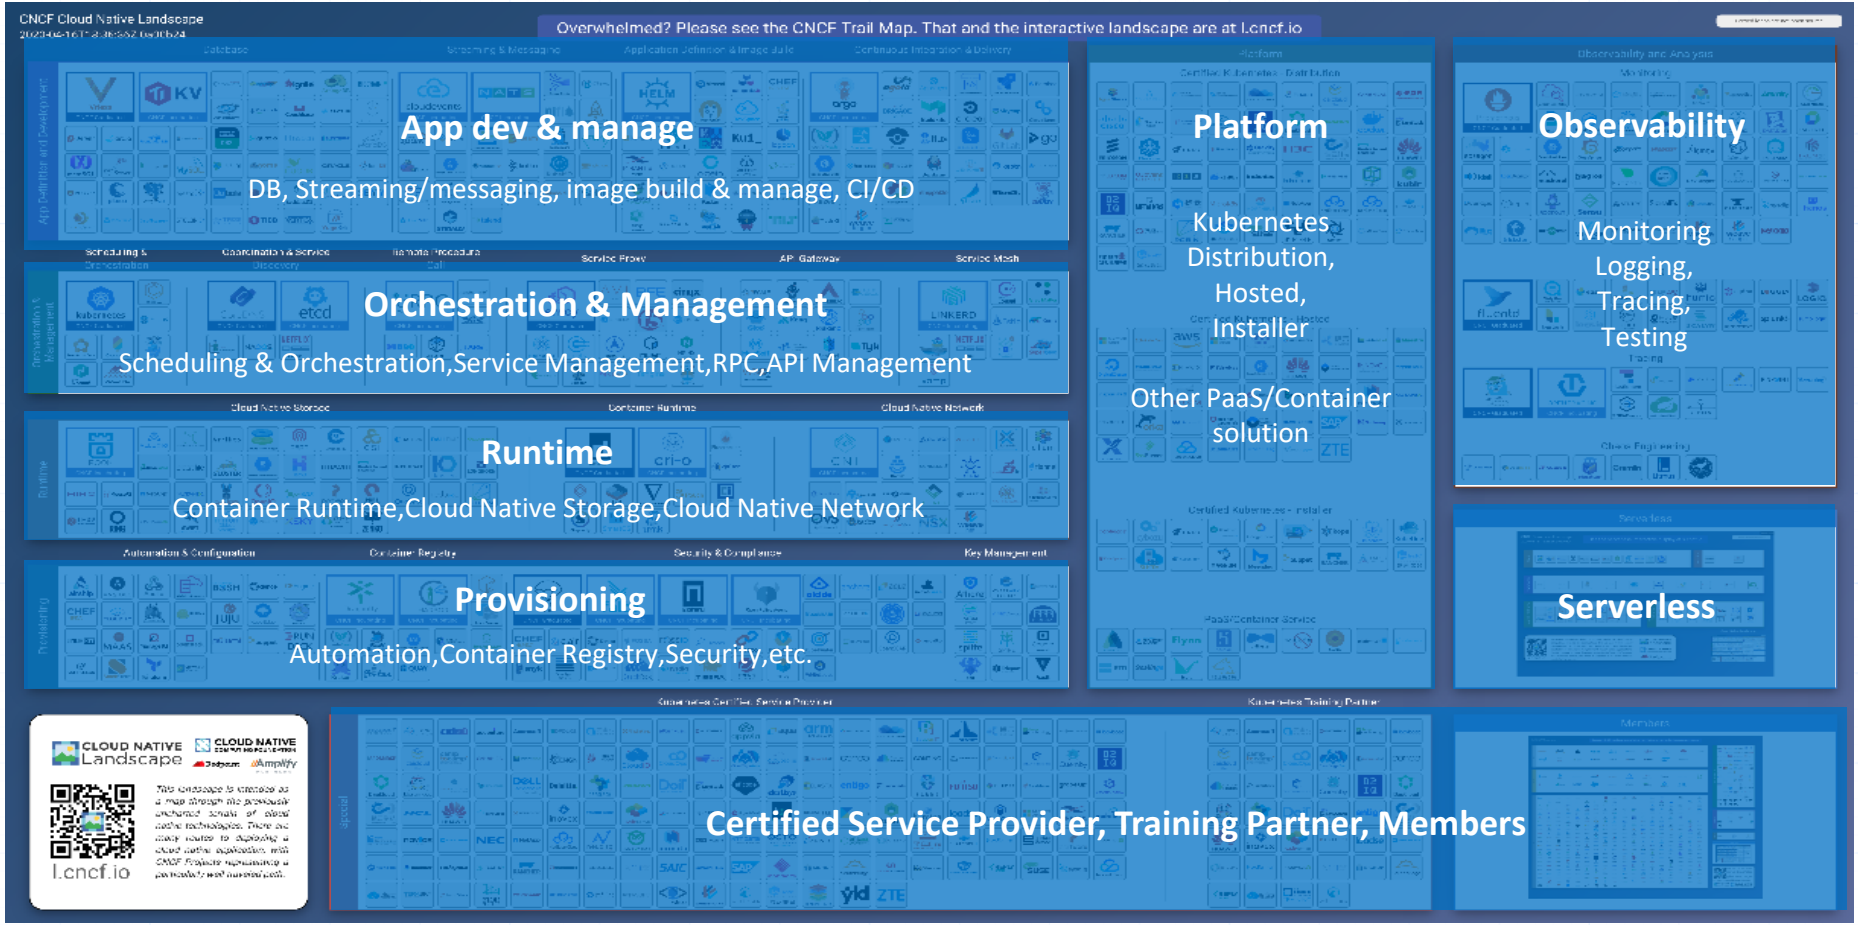

UCLOUD优刻得 | intel

中国云原生技术生态图景

CNIA 云原生产业联盟

[CNIA中国云原生技术生态图景]

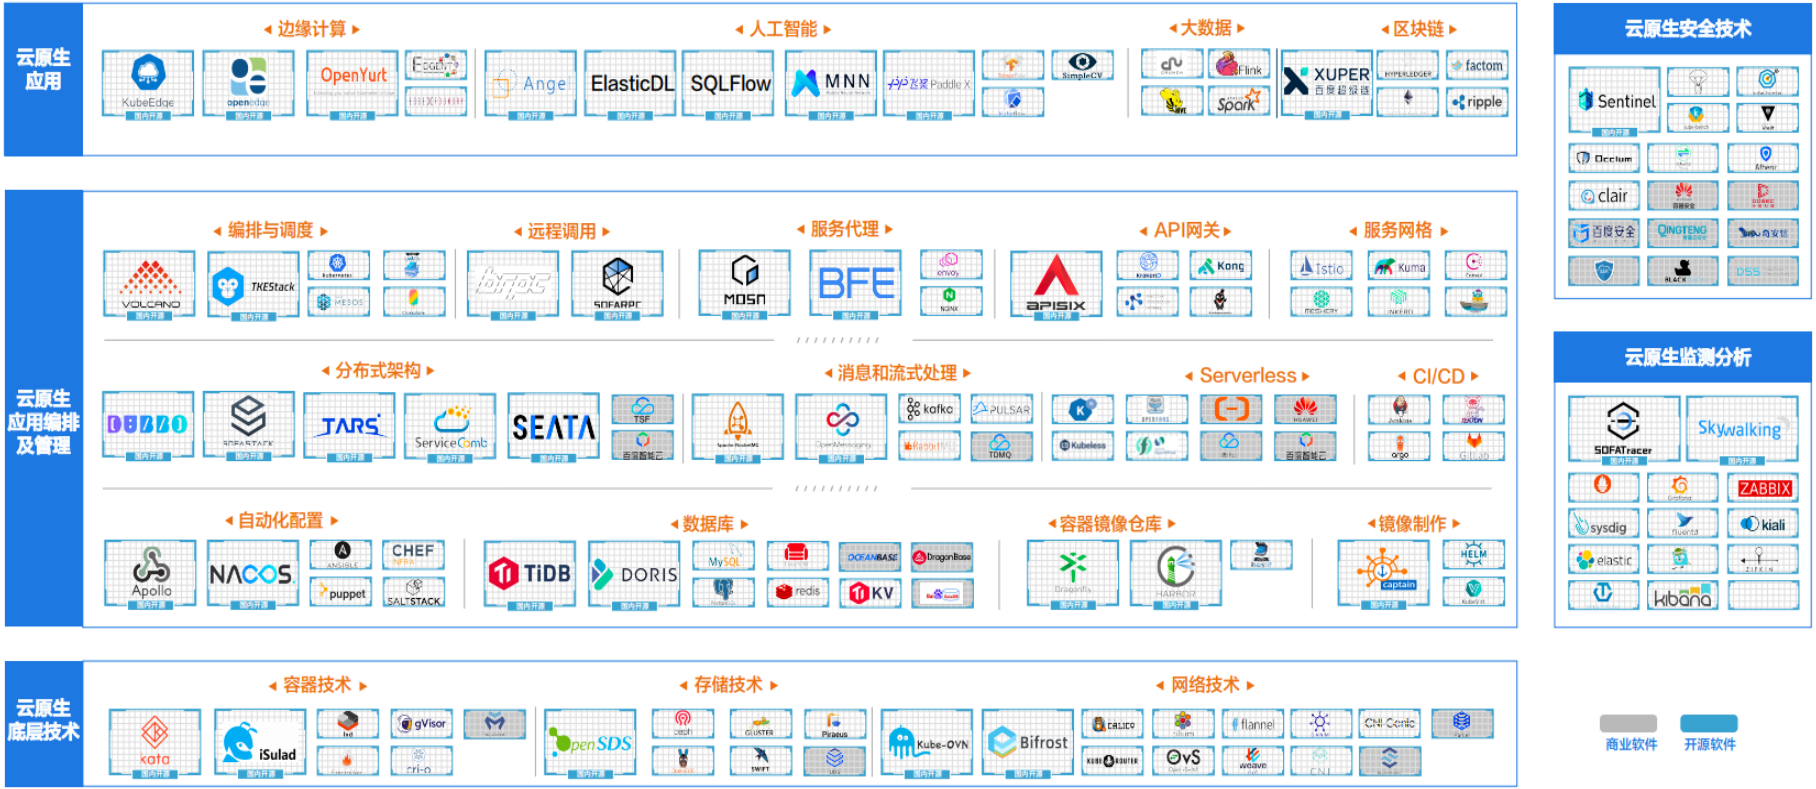

UCLOUD优刻得 | intel

中国云原生主要参与者

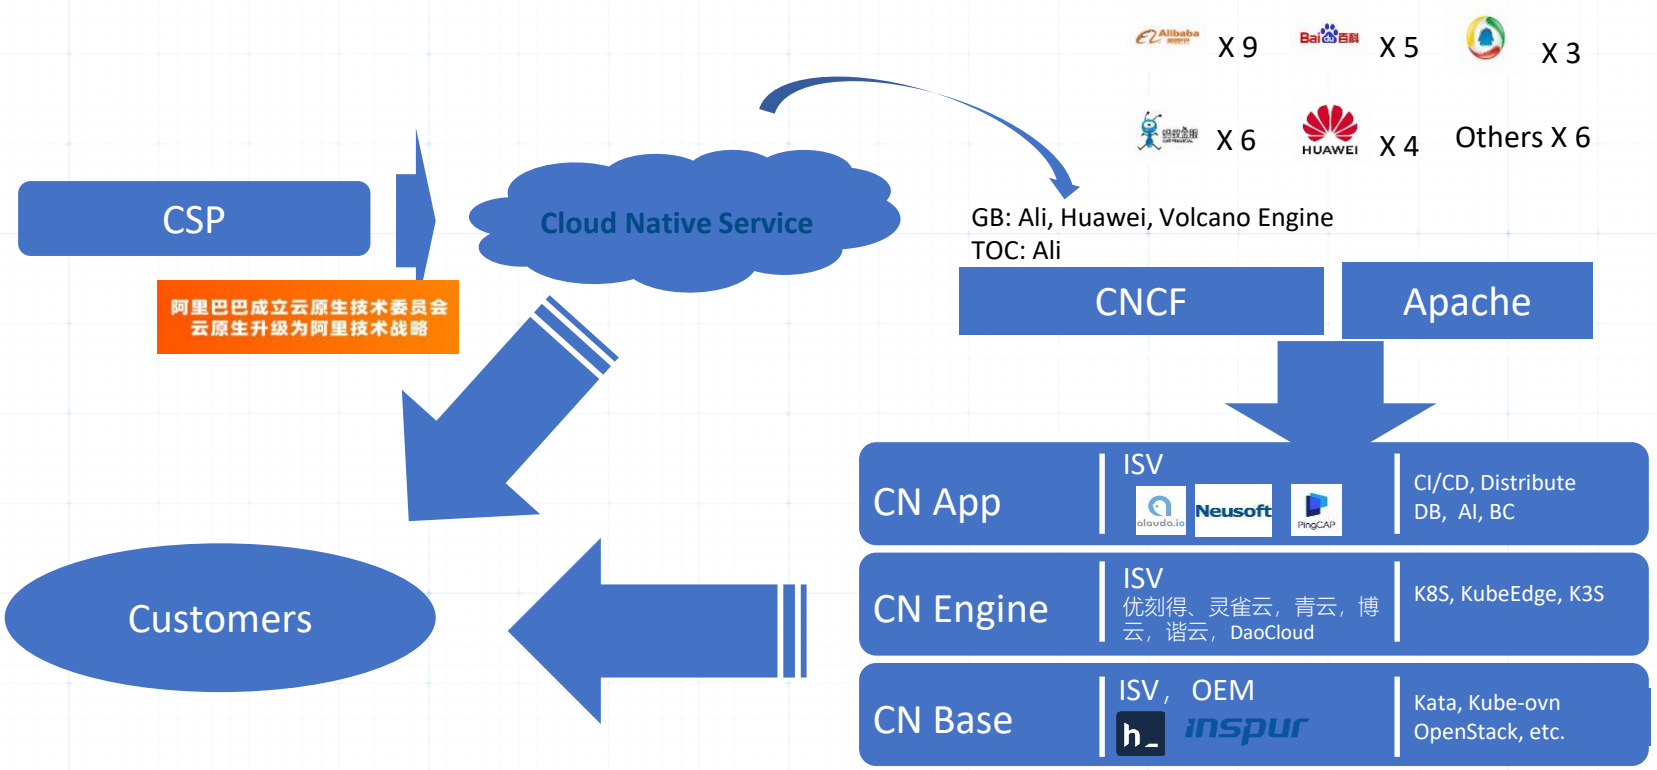

UCLOUD优刻得 | intel

趋势一：以Kubernetes为代表的容器技术已成为云计算新界面

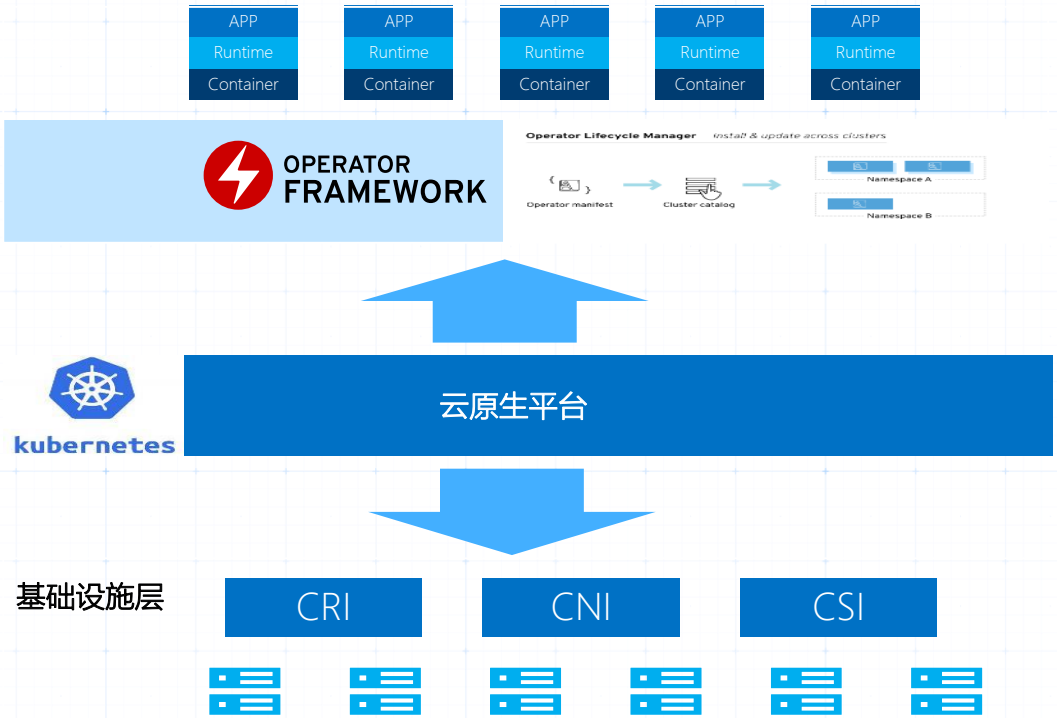

- **向下封装基础设施，屏蔽底层差异性**

  Kubernetes 支持 IaaS 层不同类型的计算、存储、网络等能力，不论是 CPU、GPU、FPGA ...，都可以统一调度、高效使用异构算力的资源

- **向上支撑多种工作负载和分布式架构**

  Kubernetes 向上支持众多开源主流框架构建微服务、数据库、消息中间件、大数据、AI、区块链等各种类型应用。从无状态应用、到企业核心应用、再到数字智能应用，都可基于Kubernetes顺利地自动部署、扩展和管理容器化应用

趋势二：云原生技术关注重心在上移：应用管理/服务网格/Serverless

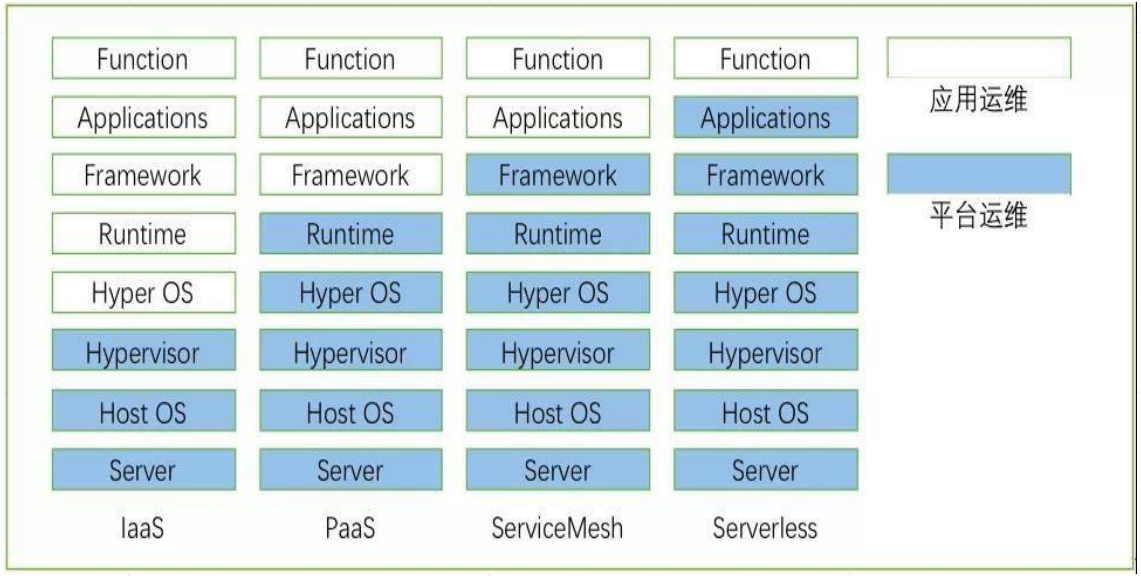

- 应用的定义和配置、发布和线上的自动化运维，成为开发和运维人员关心的核心内容
- 运维下沉，将和业务无关的管理功能和运维工作尽量下沉到基础设施中，应用可聚焦在业务能力的开发和运营

趋势三：边缘计算兴起，云原生从数据中心走向云边一体化

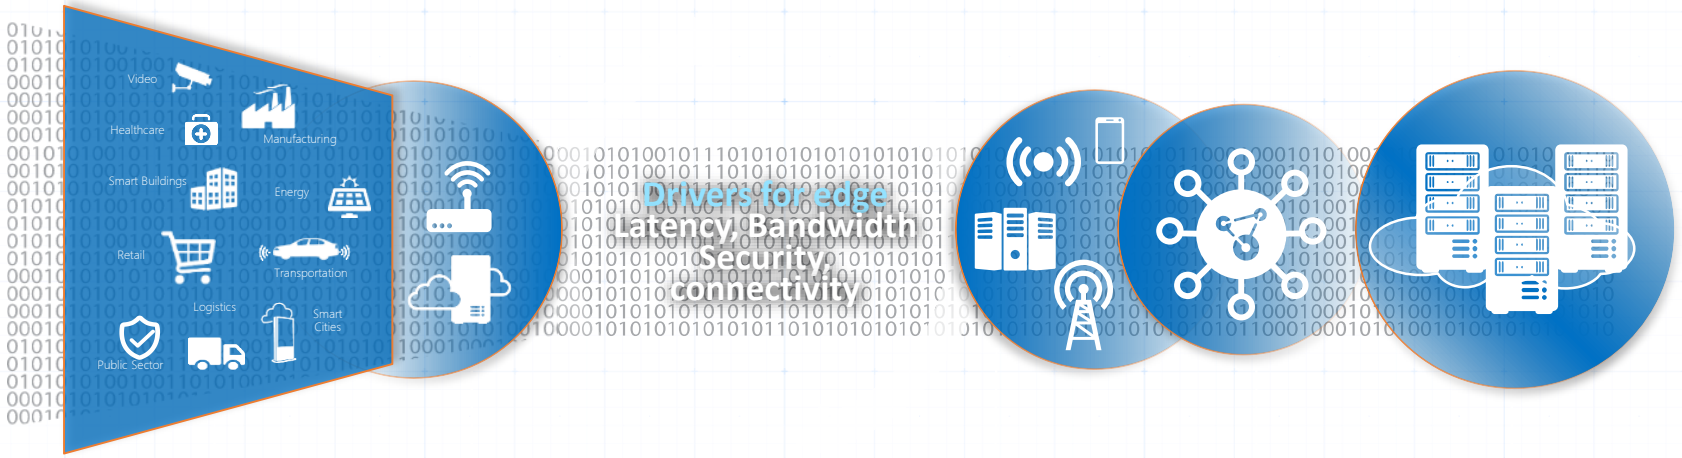

- 云原生技术和边缘计算相结合，可以快速实现云边端一体化的应用分发，解决在海量边、端设备上统一完成大规模应用交付、运维、管控的诉求

UCLOUD优刻得 | intel

云原生平台架构

intel XEON PLATINUM 英特尔® 至强® 云创新基石

UCLOUD优刻得 | intel

云原生平台架构

intel xeon PLATINUM

英特尔® 至强®
云创新基石

UCLOUD优刻得 | intel

英特尔技术优化云原生平台

为工作负载提供最佳性能和安全性，同时简化开发人员体验

计算& 加速	
Intel® QAT Device Plug-in	PC硬件加速卡，加速处理安全、身份验证、压缩和加解密
CPU Manger for K8s	整合CPU Core资源池，将特定的容器化的工作负载绑定到特定类型CPU
Native CPU Manager	绑定CPU Core到特定的工作负载/K8s Pod
Node Feature Discovery	在K8s中启用通用硬件功能发现，支持发现CPUID,SSD,SR-IOV,RDT,IOMMU,.....
Topology Manager for K8s	使其他Kubelet组件能够使用基于拓扑的信息做出资源分配决策
运行时& 语言	
Kata Container	通过虚拟化隔离技术以加强容器的安全可控
网络	
Multus CNI	容器多网络接口设备管理，以支持NFV等多功能网络划分
SR-IOV CNI	支持SR-IOV功能的高性能网络IO技术
Bond CNI	整合多个网络接口捆绑成一个逻辑网络接口，以最大化网络性能
Userspace CNI	高性能用户空间容器网络方案，支持数据平面网络加速的功能
存储	
Persistent Memory CSI Plug-in	K8s对持久化内存支持的插件

intel xeon PLATINUM 英特尔® 至强® 云创新基石

UCLOUD优刻得 | intel

第三代英特尔® 至强® 可扩展平台

MORE CORES, MORE MEMORY, SIMPLE SCALABILITY, AND INTEGRATED TECHNOLOGIES BENEFIT CLOUD NATIVE PLATFORM

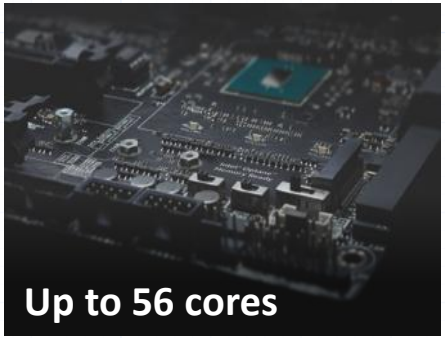

50% more memory than the previous generation and can easily scale to eight sockets in a single chassis.

Intel® QuickAssist Technology helps increase efficiency and enhance data transport and protection across the infrastructure.

Intel® Advanced Vector Extensions can accelerate the performance of mixed workloads.

Intel® Resource Director Technology provides smart resource orchestration.

intel xeon PLATINUM 英特尔® 至强® 云创新基石

UCLOUD优刻得 | intel

优化云原生数据服务

Industry-Standard Server-Based Scale-out Storage Solution

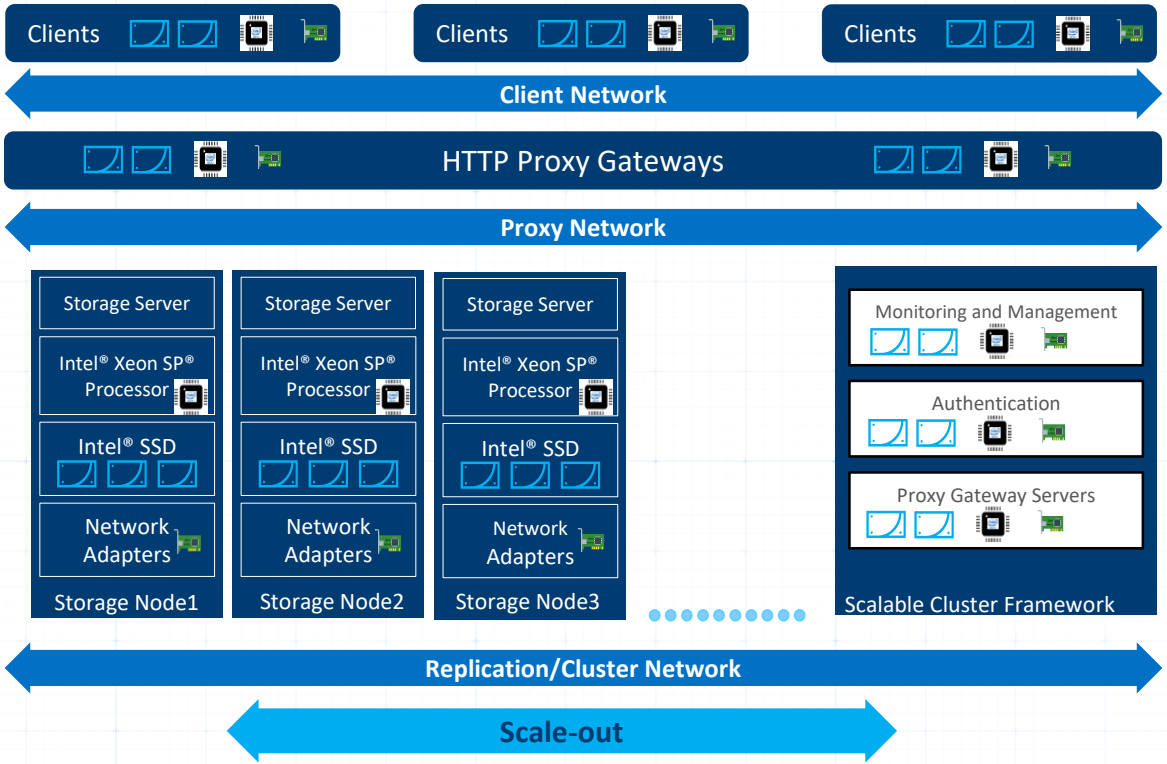

Intel Technology

PCIe, UPI, AVX 512, DDIO, RDT, VMD, RAS, SGX etc.

Erasure Code, Compression, Encryption, CPU offloads

Intel® Optane™ DCPMM

Tiered Cache, Metadata, Performance & Throughput

Intel® SSD Data Center Family

10/25/40/50/100 Gigabit Ethernet Network Adapters

Low Latency

解决Kubernetes网络关键挑战

面临的挑战

- Kubernetes Networking
  - Multiple network interfaces
- Data Plane Acceleration
  - High performance Data Plane (E-W)
  - High performance Data Plane (N-S)
  - Fail-over and high availability of networking
- Resource Management
  - Ability to request/allocate platform capabilities
  - CPU Core-Pinning and isolation for K8s pods
  - Dynamic Huge Page allocation
- Enhance Platform Awareness (EPA)
  - Discovery, Advertise, schedule and manage devices with K8s
  - Guarantee NUMA node resource alignment
- Telemetry
  - Platform telemetry info

解决方案

	SOFTWARE AVAILABILITY*
kubernetes CNI	
MULTUS	Open Source: CNI plug-in – V3.1 Aug '18 Upstream K8s: TBD
CNI USERSPACE CNI	Open Source: CNI plug-in – V1.1 Aug '18
CNI SR-IOV DPDK	Open Source: SR-IOV CNI plugin v1.0.0 6th Dec '18, Next release planned: v3.0 Mar'19
CNI BOND-CNI	Open Source: CNI plug-in – V1.0 Dec '17
Node Feature Discovery	Open Source: v0.3.0 Sep'18. Upstream K8:s Host K8s SIG Mar '19
CPU Manager for Kubernetes	Open Source: CMK v1.3 Sept'18 Upstream K8s: Phase 1 - v1.8 Sep '17
Native Huge page support for Kubernetes	Upstream K8: Beta - v1.10 Dec '17
Device Plugin (SR-IOV, QAT, FPGA, GPU)	Open Source: SR-IOV v2.0.0 Dec '18, QAT v1.0 v1.0 Aug'18, FPGA & GPU May '18
Topology Manager (NUMA)	Upstream K8s: Planned k8s v1.15 Jun'19
collectd	Upstream collectd: v5.7.2 Jun '17, v5.8.0 Nov. '17

Open Source: Available on github https://github.com/intel
NFD at https://github.com/kubernetes-incubator/node-feature-discovery

intel xeon PLATINUM 英特尔® 至强® 云创新基石

Node Feature Discovery

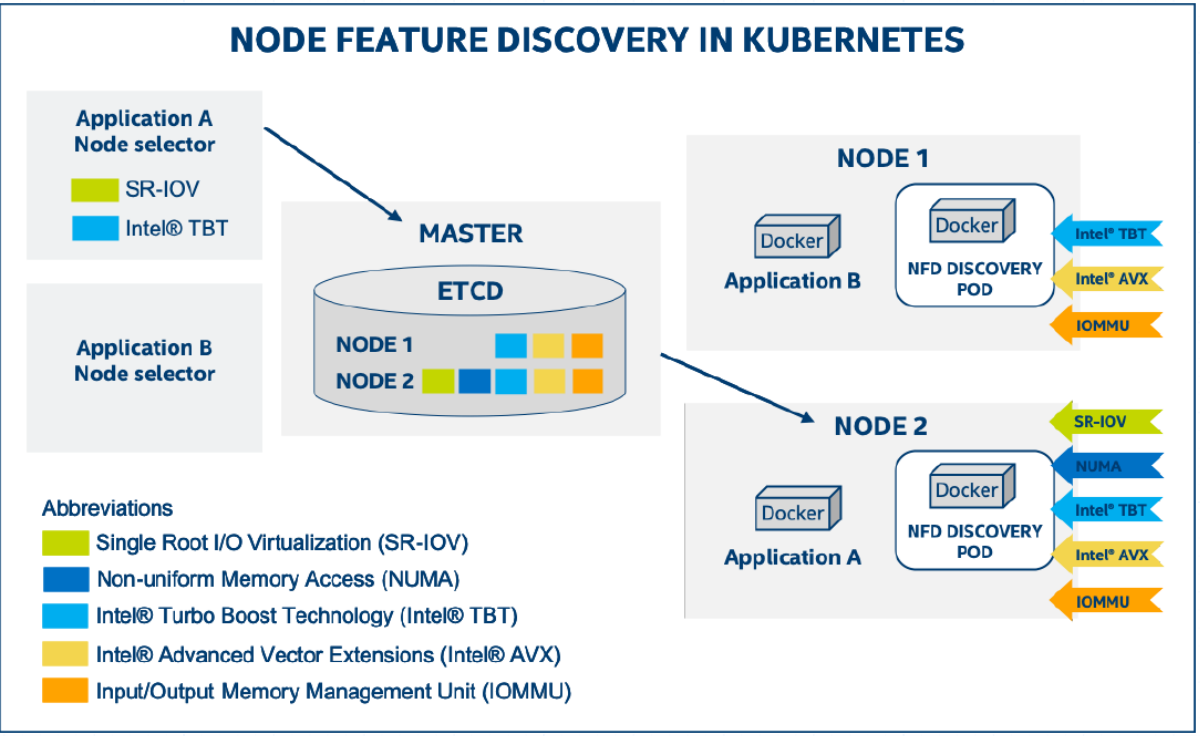

- Node Feature Discovery brings Enhanced Platform Awareness (EPA) in K8s
- NFD detects resources on each node in a K8s cluster and advertises those features
- Allows matching of workload to platform capabilities

UCLOUD优刻得 | intel

Topology Manager for K8s

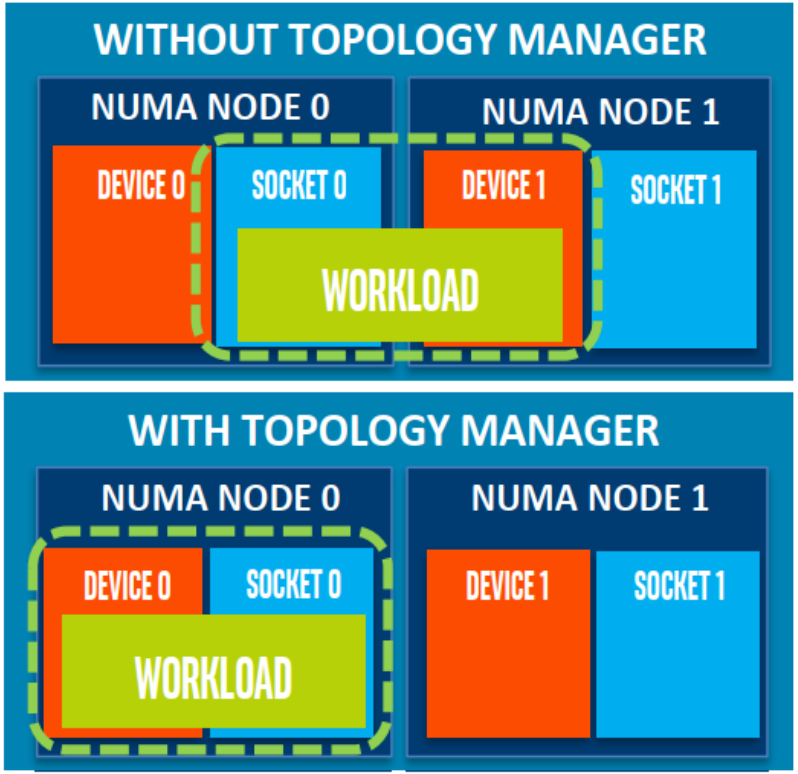

- Topology Manager, a kubelet component, provides an interface to co-ordinate resource assignment on node level
- Manages the resources allocated to workloads in a NUMA topology aware manner

Native CPU Manager & CPU Manager for K8s

- Allocation of exclusive CPUs to certain pod containers
- Enable static policy, configure pod with guaranteed QoS class and request CPU cores
- The CPU manager does not take the isolcpus Linux* kernel boot parameter into account

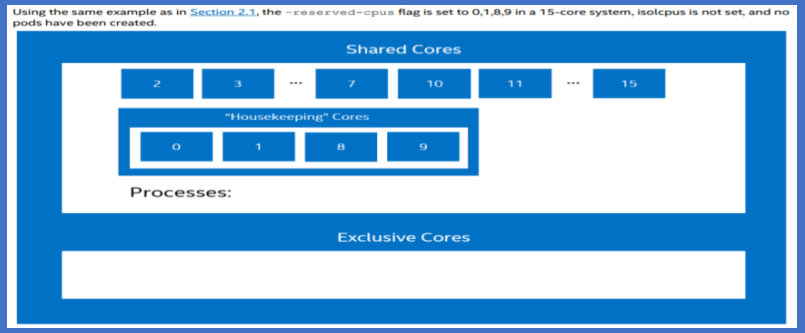

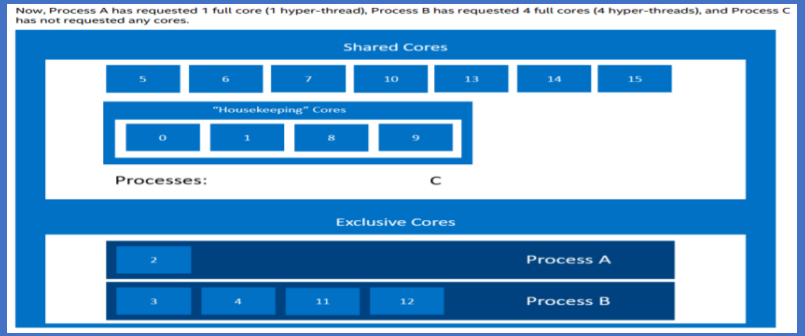

- Guarantees high priority workloads are pinned to exclusive cores
- Gives a performance boost for fastpath workloads & negates noisy neighbour scenario by using cores from exclusive pool (incl. isolcpus)
- Shared cores are assigned for slowpath workloads

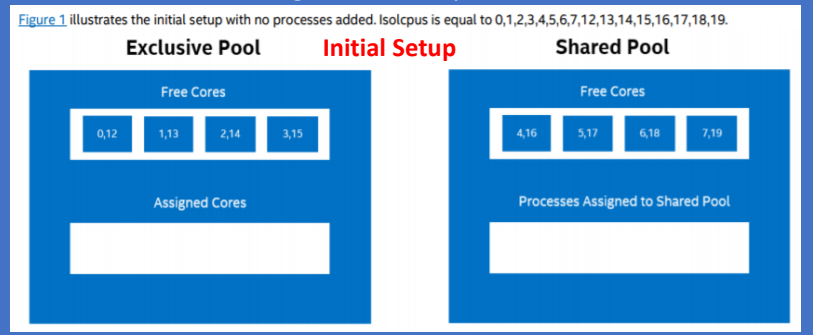

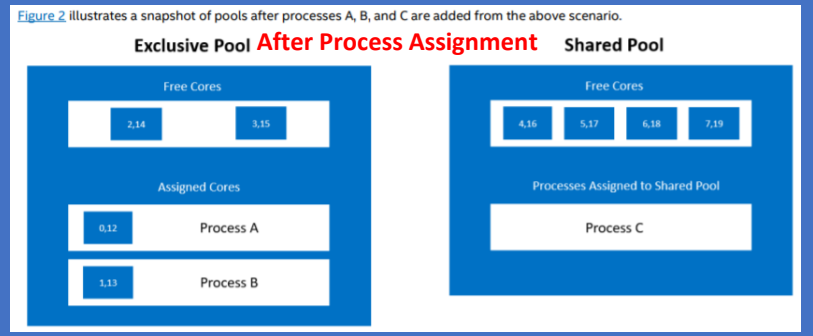

A & B: High priority process / C: Lower priority process

UCLOUD优刻得 | intel

Native CPU Manager vs CPU Manager for K8s

Table 3. Technology Comparison

NATIVE CPU MANAGER	CPU MANAGER FOR KUBERNETES
K8s* code base. Beta since 1.10	Kubernetes* integration using K8s* external APIs
Updates container cgroups to provide pinning	Wrapper program that runs before workload and performs taskset command for pinning
Unaware of isolcpus	Uses isolcpus
Pod level isolation guaranteed	Gentleman's agreement for isolation
Resource account done via K8s* first class resource CPU	Resource accounting done via host file system and extended resources
Pod spec contains CPU requests	Resource accounting done via host file system and extended resources
3 CPU Pools – Shared, Reserved & Exclusive Allocations	4 CPU Pools – Exclusive, Shared, Infra & Exclusive-non-isolcpus (optional)
Shared & Exclusive pools grow and shrink dynamically as requests come in	All pools are static after deployment
NUMA Alignment with Topology Manager	NUMA alignment manual
Deployment done via K8s* release	Deployment done via a set of K8s* Pods

Multus CNI

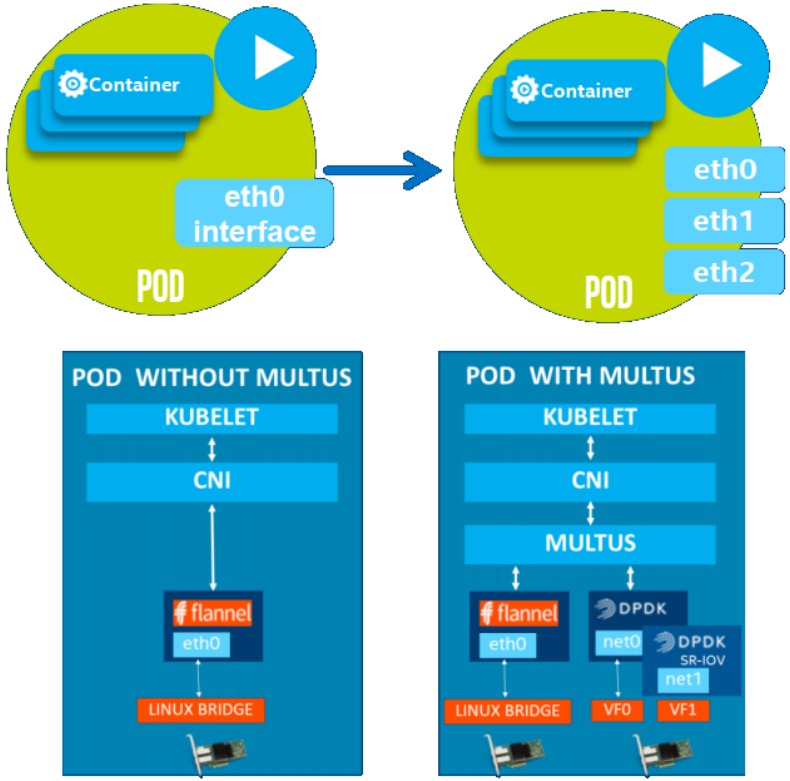

- Multus is a CNI proxy and arbiter of other CNI plugins
- Multus invokes other CNI plugins for network interface creation
  - Master plugin is identified to manage the primary network interface (eth0) for the pod
  - Other CNI minion plugins (SR-IOV, vHost CNI, etc.) can create additional pod interfaces (net0, net1, etc.) during their normal instantiation process

SR-IOV CNI & SR-IOV Device Plugin

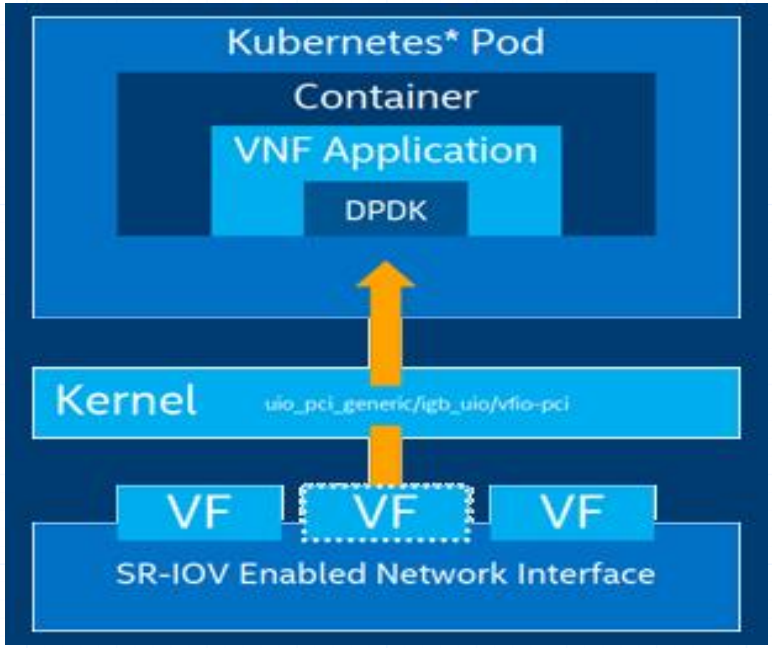

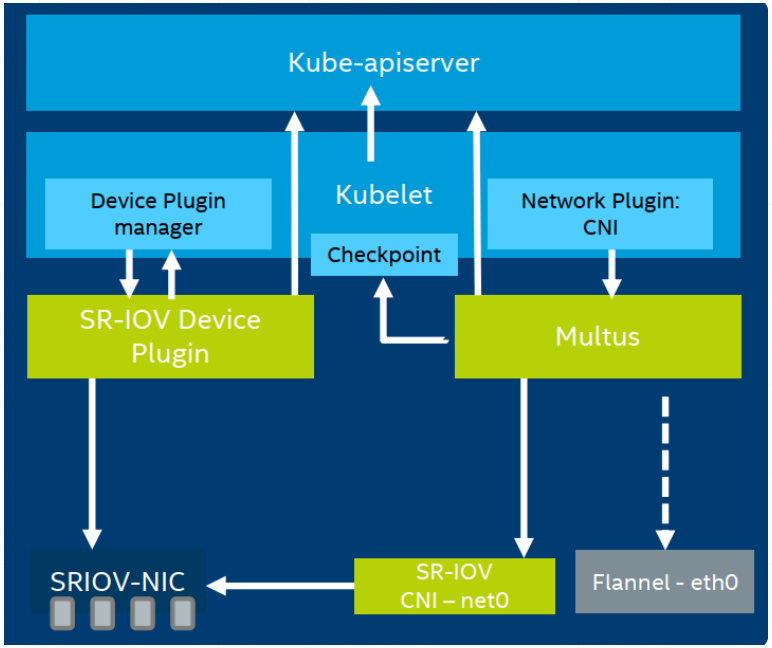

- Allows SR-IOV support in Kubernetes via a CNI plugin
- SR-IOV Network device plugin provides a mechanism to advertise SR-IOV network capabilities of a host
- VNF application that demands high performance NICs as resource requests then can be scheduled to this host by K8s Scheduler
- SR-IOV device plugin allocates required VFs meeting NUMA alignment

云原生上的安全和隔离

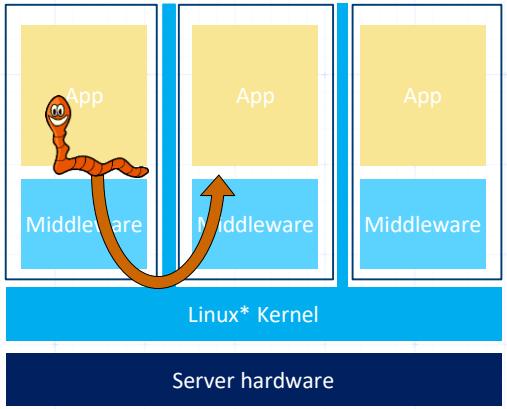

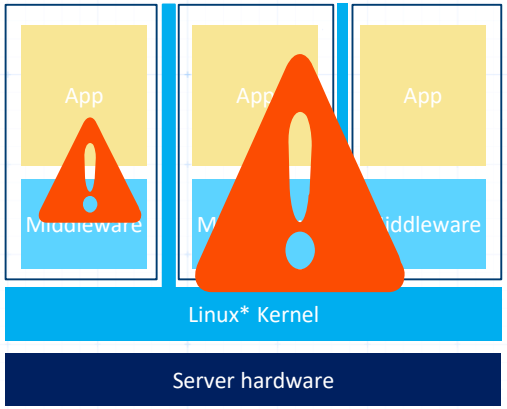

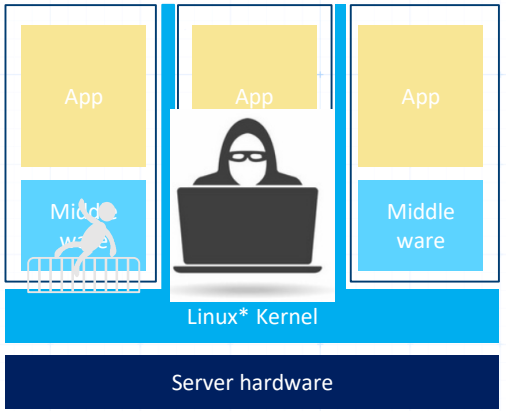

CVE-2019-5736

CVE-2019-14271

UCLOUD优刻得 | intel

基于虚拟机隔离的Kata容器

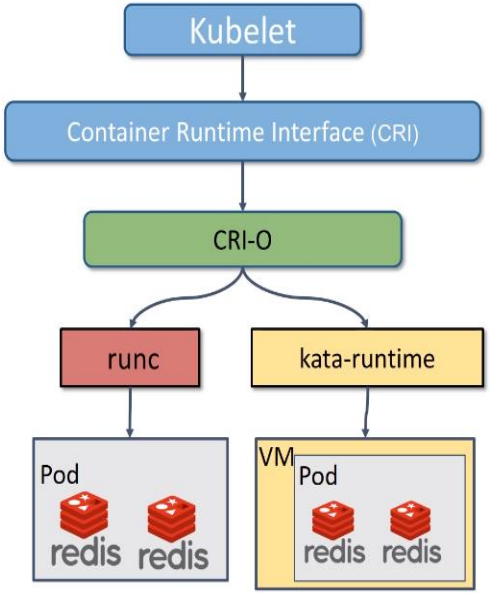

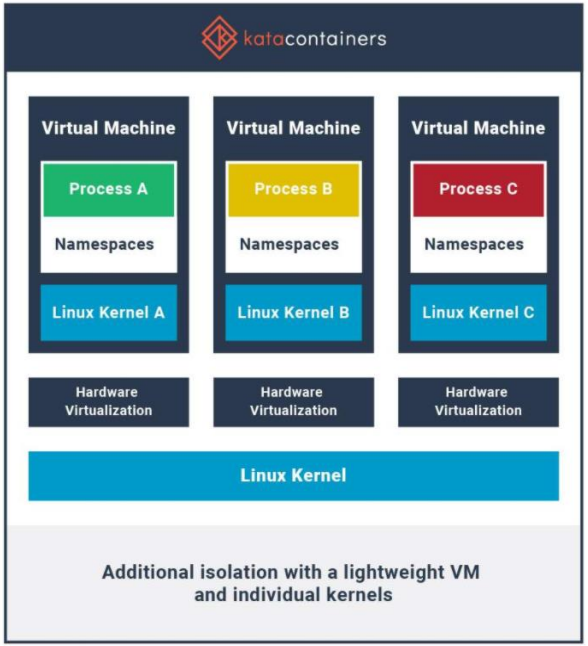

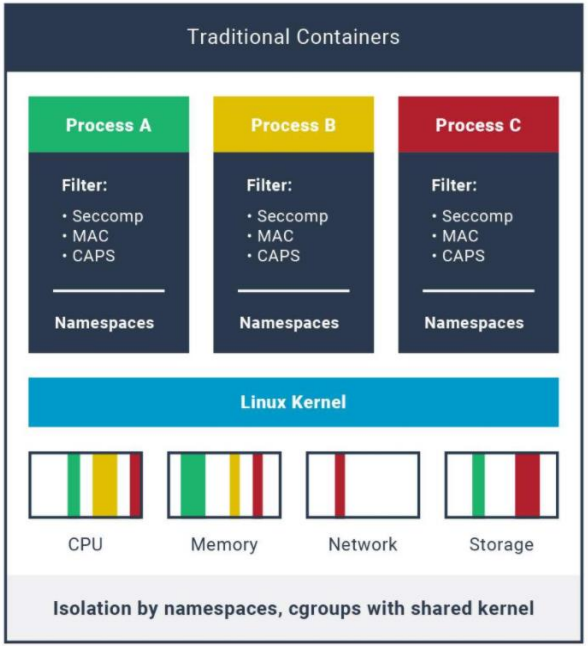

UCLOUD优刻得 | intel

Intel 助力客户云原生技术落地

Technology Insight

Abundant Cloud Solutions

Global Use Cases

Vendor Neutral

Broad Ecosystem

68% IMPROVED HPC FSI PERFORMANCE

30x IMPROVEMED AI INFERENCE PERFORMANCE

7.8x IMPROVED STORAGE COMPRESSION

2.5x IMPROVED CLOUD-NATIVE CONTAINER OPENSSL PERFORMANCE

Move Faster

Intel® Ethernet

Intel® Silicon Photonics

Intel® Network Switches

BAREFOOT NETWORKS | an Intel company

Store More

intel OPTANE PERSISTENT MEMORY

intel OPTANE SSD

Process Everything

intel XEON PLATINUM

intel ATOM

intel AGILEX

intel eASIC

Xe

intel MOVIDIUS

UCLOUD优刻得 | intel

总结

1 云原生是不可阻挡的，它给客户带来价值的同时也带来挑战，这对云合作伙伴来说是机遇。

2 Intel 拥有广泛地产品组合来增强云原生，在云化方面也有丰富地经验。

3 从最简单地步骤开始，制定您的云化战略，与 Intel 合作利用云原生的价值。

UCLOUD优刻得 | intel

THANKS!