

# 扩大Ceph的影响力

Tushar Gohad, 首席工程师

英特尔数据中心事业部



Ceph中国社区

IT大咖说  
知识共享平台

# CEPH & DATA 愿景

## 未来

数据量持续不断的增加  
人们保留并挖掘数据以获取信息/洞察力  
**90 %**的数据在过去**2年**中生成  
到**2025年**共会产生**165ZB**数据(来自IDC报告)

## Ceph 已实现

## 最新进展

开放式解决方案  
支持基于开放解决方案的业务流程  
使我们能够容纳大量的数据

数据块和文件发展迅猛  
电信运营商开始应用Ceph解决方案  
效率提高**50 %**  
硬盘和固态硬盘的发展迅猛



Ceph中国社区

IT大咖说  
知识共享平台

# 数据增长出现在各个领域



**10 GB**

普通互联网用户

**3000 GB**

智能医院

**4000 GB**

自动驾驶

**40000 GB**

飞机数据

**1 PB**

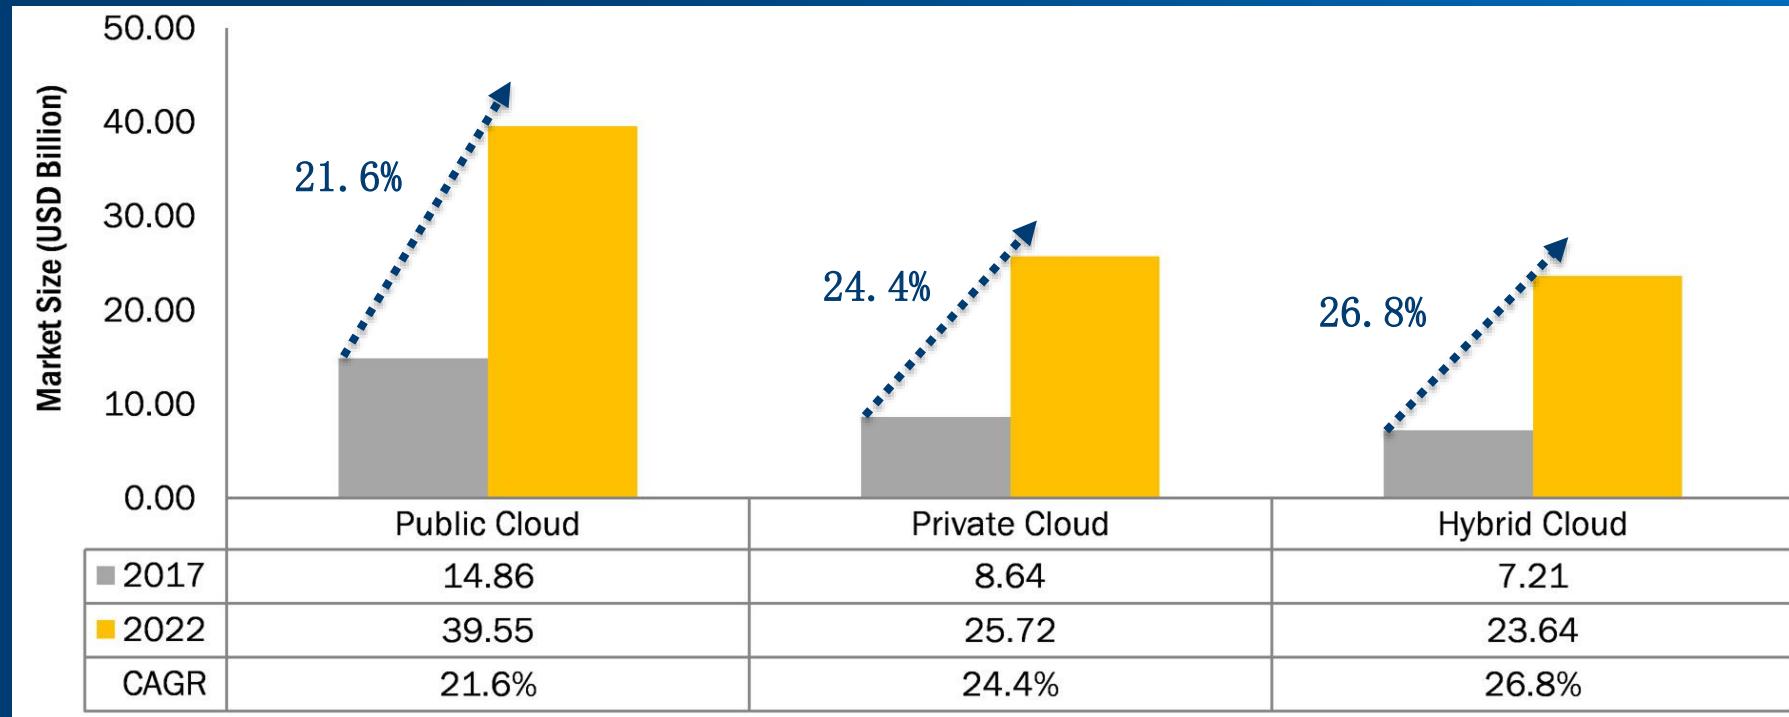
智能工厂

1. 来源: <http://www.cisco.com/c/en/us/solutions/service-provider/vni-network-traffic-forecast/infographic.html>

2. 来源: [http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud\\_Index\\_White\\_Paper.html](http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.html)

3. 来源: <https://datafloq.com/read/self-driving-cars-create-2-petabytes-data-annually/172>

# Ceph将推动云存储革命



# Ceph全球用户



## 电信行业



## CSP/Web



## 企业



## 金融领域



## 媒体行业



Level(3)  
SONY  
Bloomberg



## 能源领域



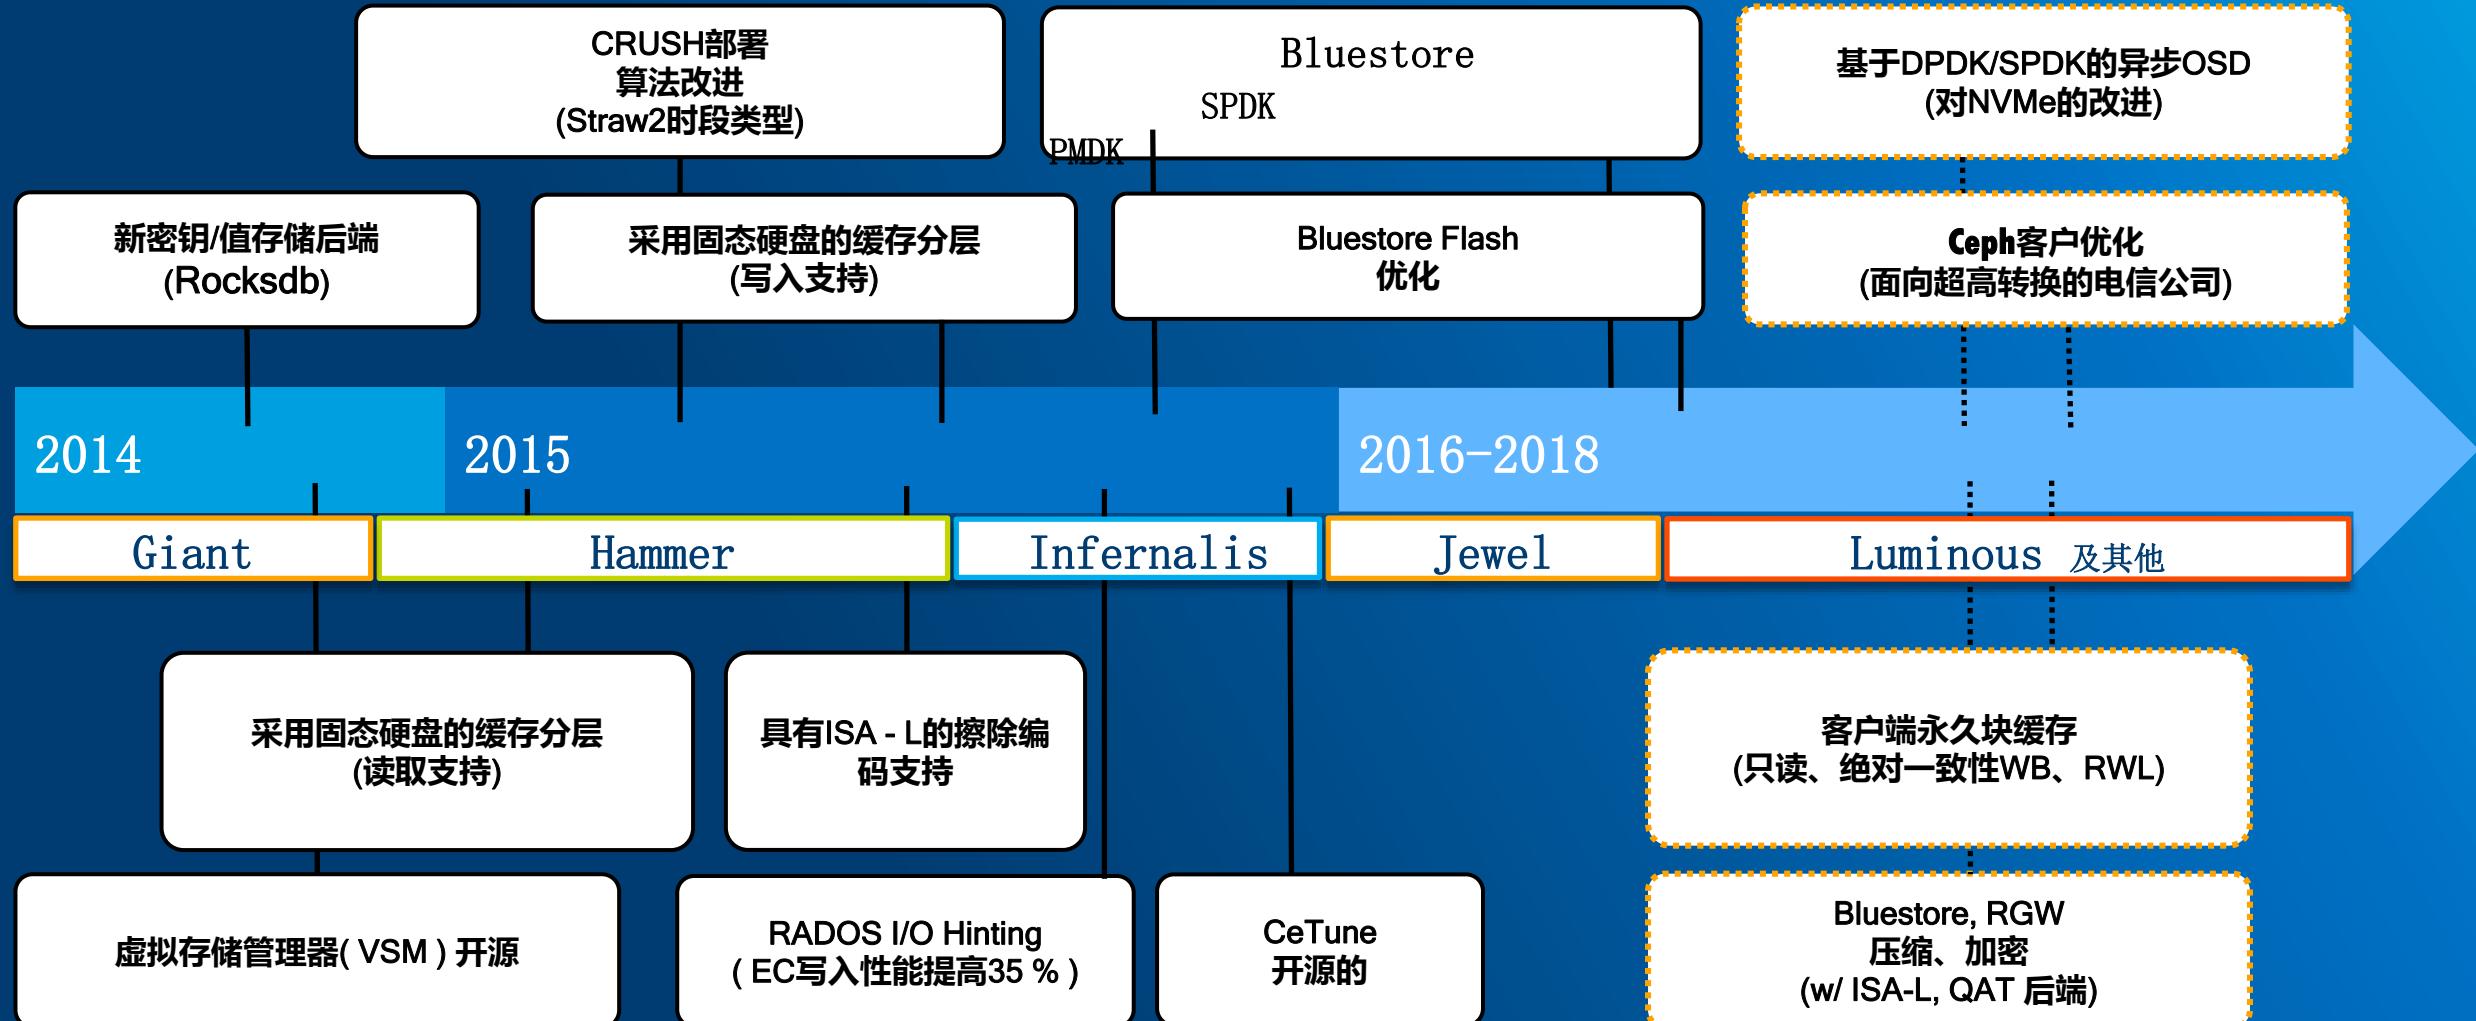


Ceph中国社区

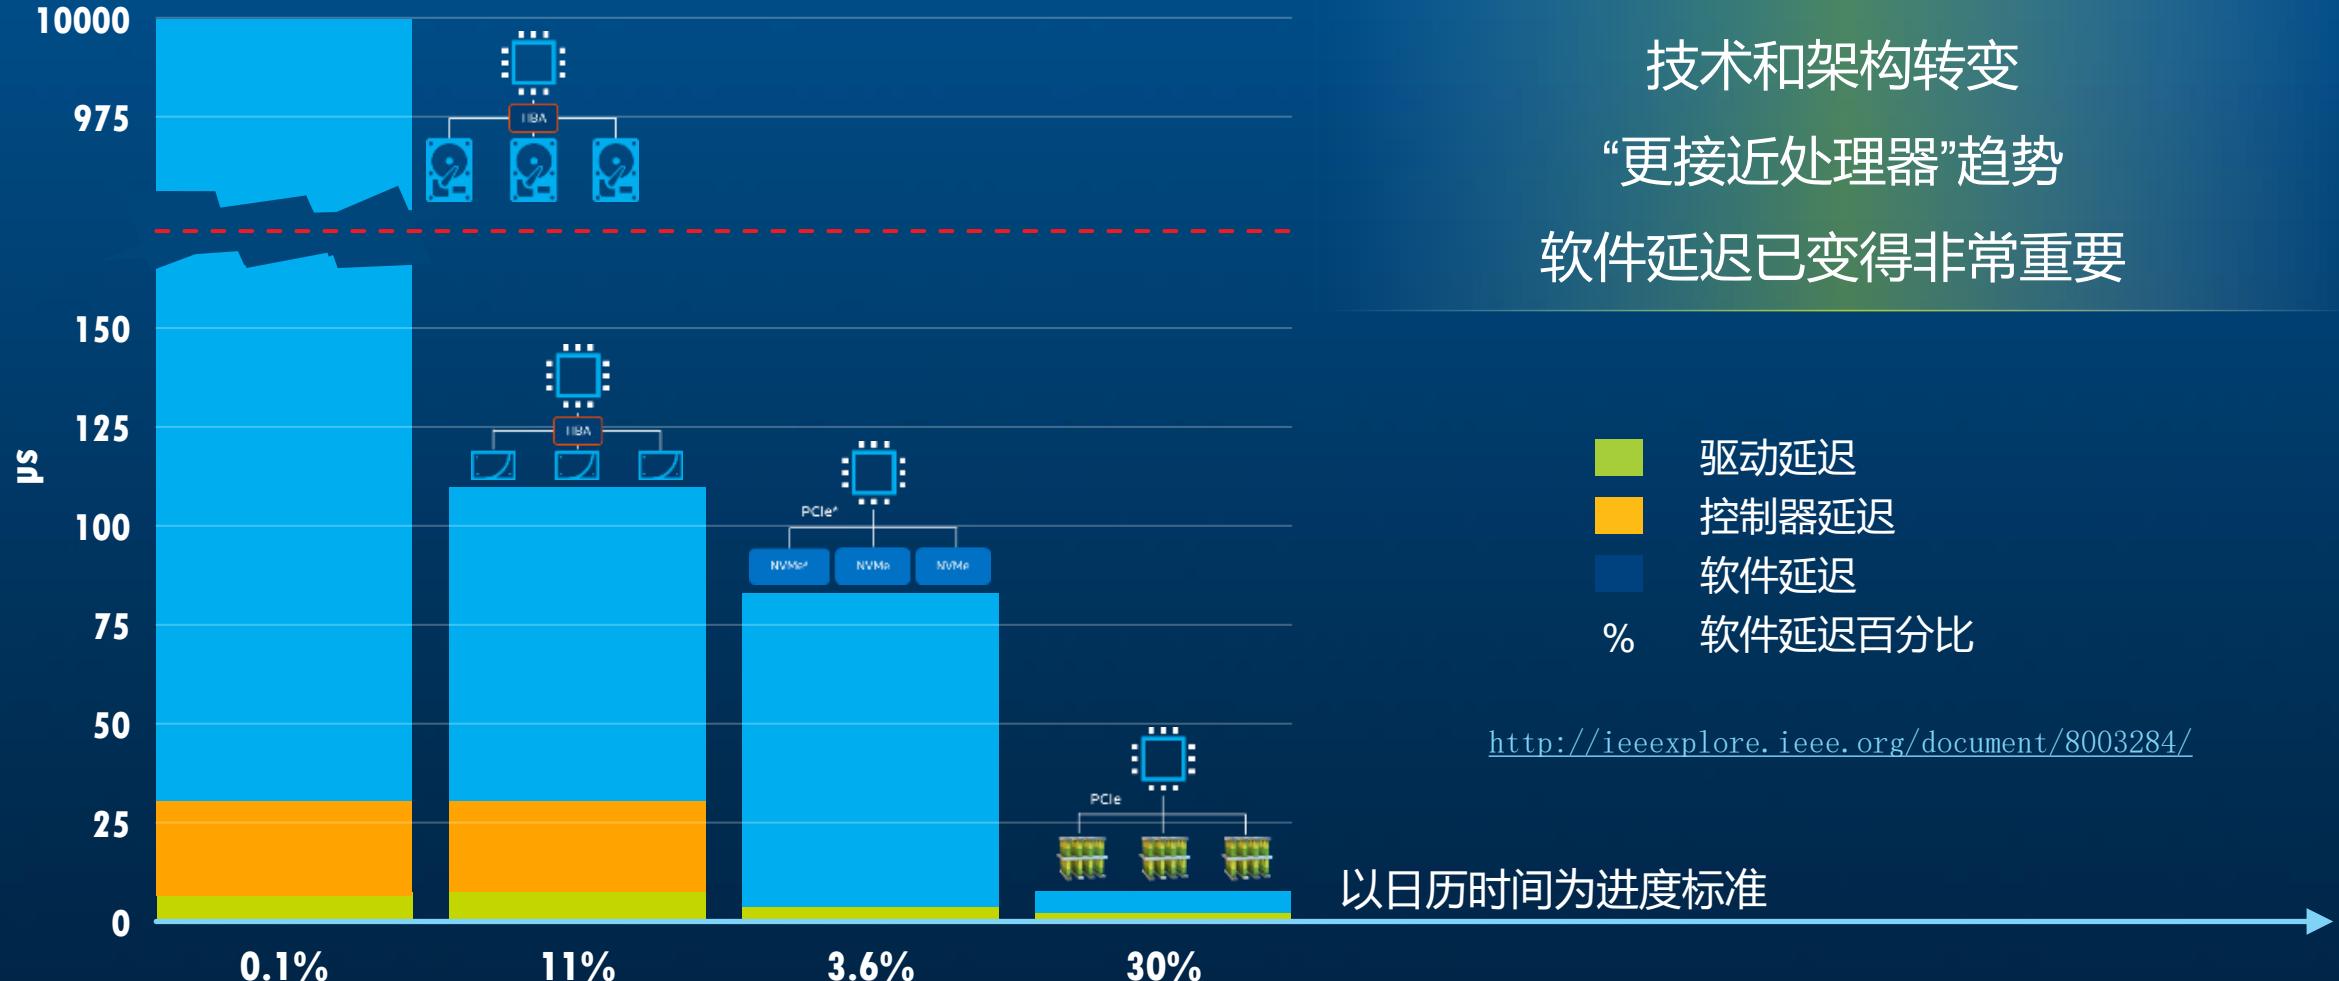
IT大咖说  
知识共享平台

# 技术趋势

# 英特尔在Ceph社区贡献历程



# 非易失性存储器和延迟



# 存储加速 软件



## 数据平面 开发套件 (DPDK)

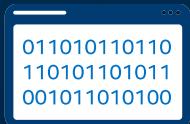
building blocks for fast packet processing  
async, polled mode, lockless, userspace  
[dpdk.org](http://dpdk.org)



## 存储性能开发套件 (SPDK)

high performance storage building blocks  
async, polled mode, lockless, userspace  
[spdk.io](http://spdk.io)

## 英特尔®智能存储加速库 (Intel® ISA-L)



optimized low-level functions for storage apps  
EC, CRC, RAID, compression, crypto, hash  
[github.com/01org/isa-l](https://github.com/01org/isa-l), [isa-l\\_crypto](https://github.com/01org/isa-l_crypto)

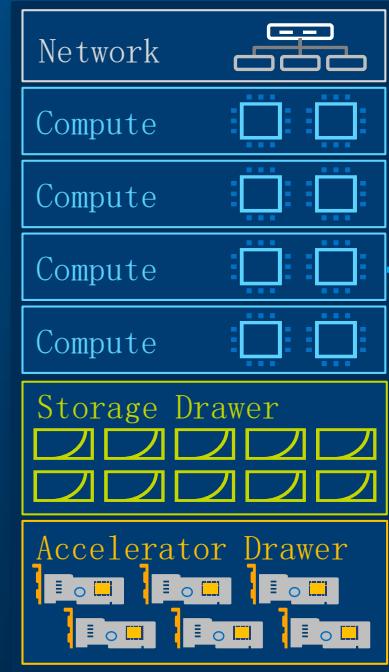


## 永久存储器 开发工具包 (PMDK)

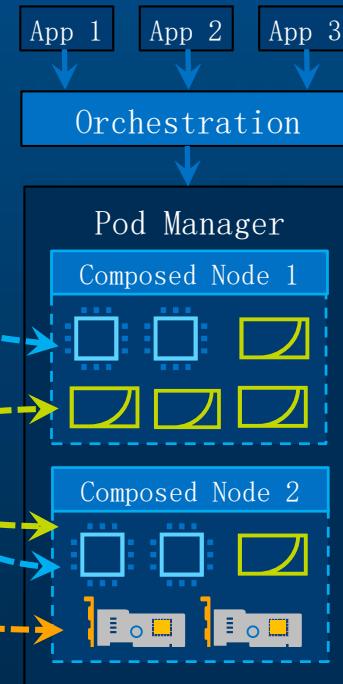
libraries for using persistent memory  
SNIA NVM programming model  
[pmem.io/pmdk](http://pmem.io/pmdk)

# 解耦：可扩展性的关键

解耦

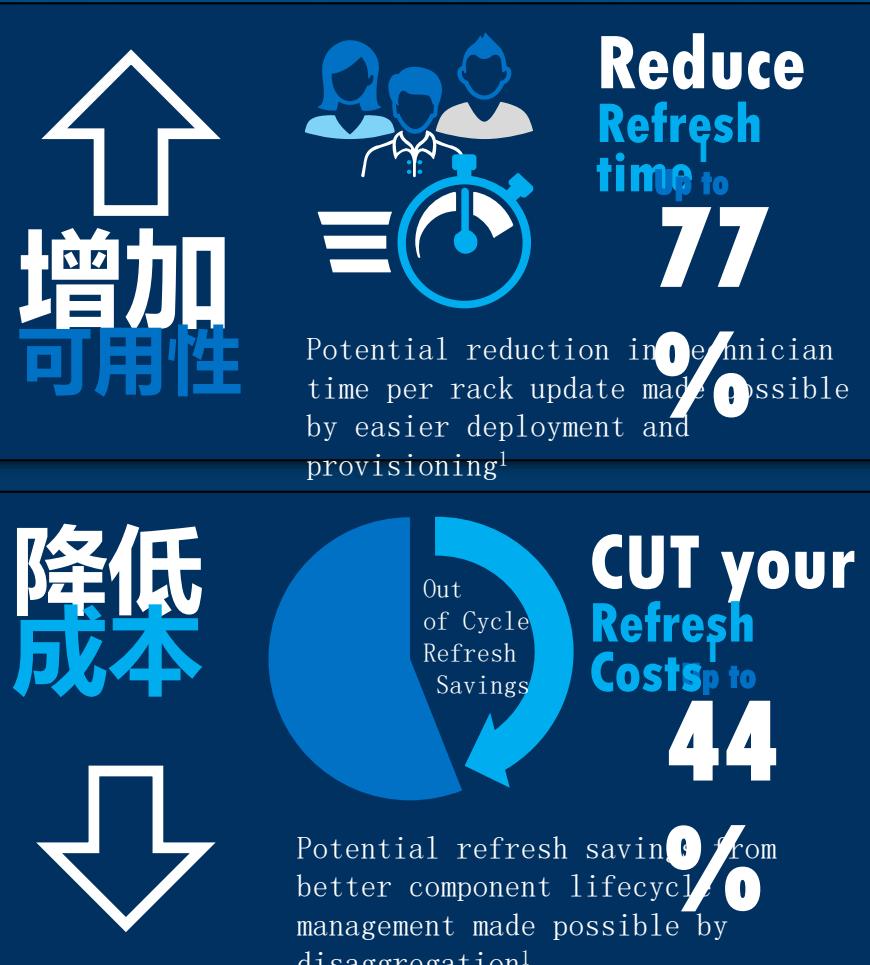


重组



减少前期支出  
并逐渐减少总体支出

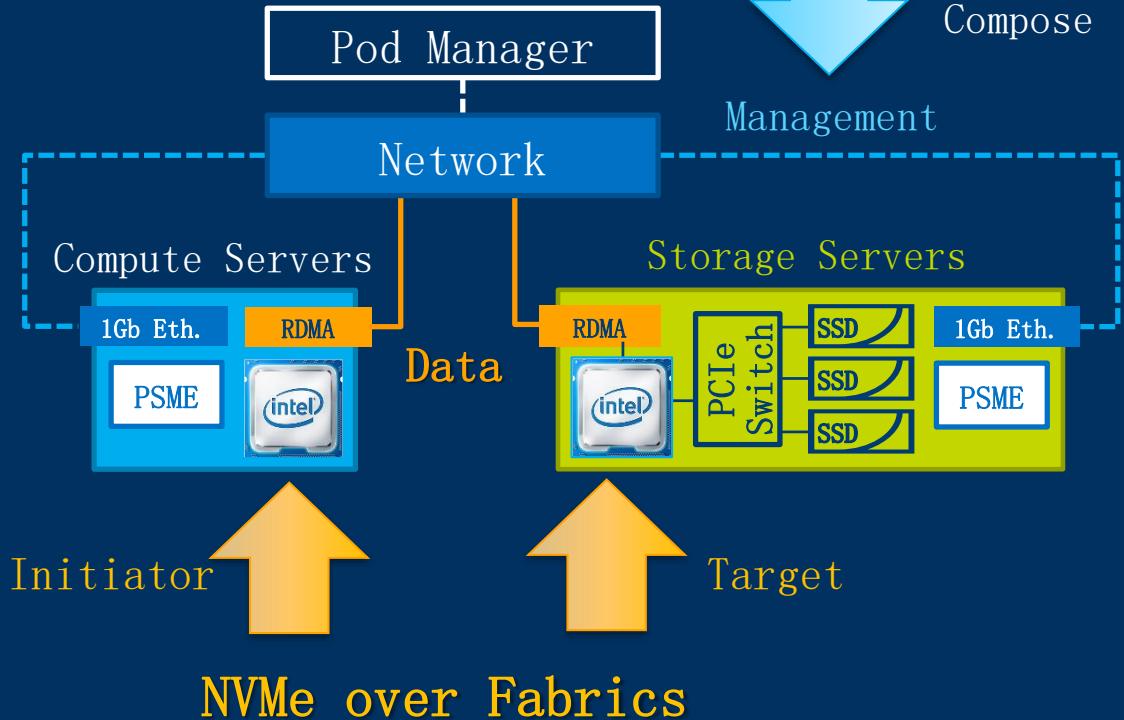
在运行中编写硬件资源



# 解耦：NVMe over Fabrics



Built on modern,  
**OPEN standards**



**“Swordfish”**  
Storage Management  
Configure &  
Compose



- Configure Network and Storage for Management
- Compose Storage Resources

bringing **multiple standards** together under one framework



- Discover Targets
- Connect with Initiators
- Manage Connection

# 可管理性

- 需要以更标准化的方式将通用工作流自动化
  - Provisioning
  - 运维和监控
  - 诊断和故障排除
- 整合云管理的架构 (K8S)
- OPNFV (Open Platform for NFV)
  - 让Ceph更容易地应用到电信行业





Ceph中国社区

IT大咖说  
知识共享平台

让世界 变得更美好

# 失踪和被拐卖儿童国家中心 ( NCMEC )

其任务是帮助寻找失踪儿童、减少对儿童的性虐待和防止儿童受害

面向家庭、受害者、私人组织、执法机构和公众的国家信息中心和资源中心



NATIONAL CENTER FOR  
**MISSING &  
EXPLOITED**  
CHILDREN®

# Ceph在ncmec

- Ceph通过以下方式提供了帮助:
  - 数据库备份和恢复时间从**12小时减少到10分钟**
  - 转换到“按需”服务器购买模式
  - 对托管于CephFS的大型数据集进行生产分析
- 考虑采用Ceph的用户应当.....
  - 亲自尝试DEMO！以周为单位衡量投资回报率
  - 从朋友或顾问处获取有关Ceph的入门知识
  - 利用Ceph API—与Sensu一起使用可大大改进监控效果
- 开发者社区留言:

“NCMEC依赖于Ceph等开源工具以及每天支持它们的社区。没有你们的付出，我们不可能取得现在的成绩。谢谢！”

- Russ Johnson , NCMEC**信息技术总监**



# 持续创新!

- 保证工作负载交易在10毫秒之内 (databases, AI)
- Non-volatile memory 编程 模式
- Disaggregation with NVMe-over-Fabrics
- 可管理型 - 降低用户使用门槛



# 钟忻

海航易建科技技术副总裁兼云服务事业群总经理

# 海航eKingCloud基础云

海航eKingCloud基础云覆盖七个国家十三个大区域，形成全球化战略布局，承载海航整体业务系统部署，满足业务运行稳定性要求，业务运营创新性需求，航空和金融相关监管要求，同时实现多地容灾备份和数据保护，助力海航国际化发展

海航业务系统截至2018 H1云化比例将达**80%**以上，实现**700+**套业务系统上云



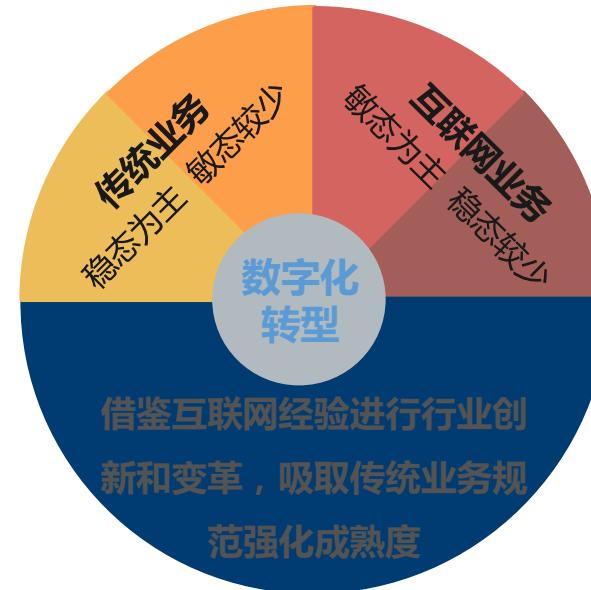
# 海航eKingCloud基础云创新理念

采用“**OMEGA**”架构设计理念，推进海航eKingCloud基础云产业化落地  
统一的技术栈具备通用性和独立性，可快速交付并灵活扩展，推进创新速度

互联网行业是开源技术的践行者

随着云计算和大数据的快速发展

采用开源技术构建技术架构已经成为主流选择



- ✓ 开放的架构设计
- ✓ 良好的生态系统
- ✓ 活跃友好的社区



应用



演进

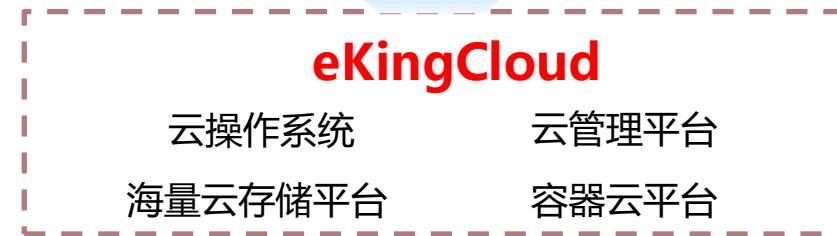


生态



# 海航eKingCloud产品体系

面向大型政企规模化部署



面向中小企业、远程和分支机构部署

超融合一体机

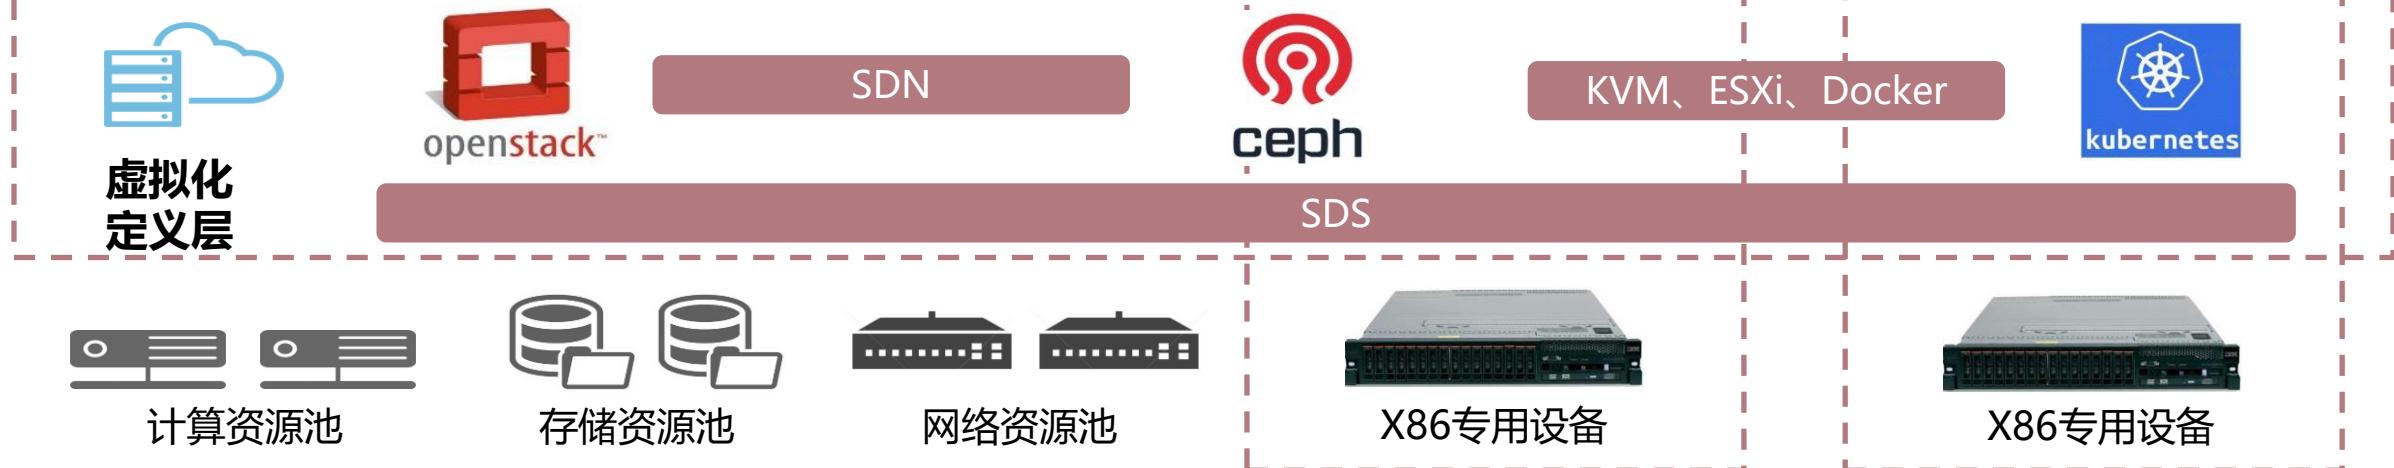
以Ceph分布式  
统一存储技术为基础

基于KVM做深度二次开  
发的计算虚拟化专用设备

存储一体机

基于Ceph做深度二次开  
发的分布式存储设备

提供对象存储、块存储和  
文件接口



# 海航eKingCloud与Intel深入合作

## ✓ 云存储数据加密

降低CPU开销 加快加密速度 提升服务性能

QAT技术，实现数据加密算法中的加密运算和数据压缩/解压运算

## ✓ EC集群

提高存储吞吐量 减少空间使用量 安全灵活

通过在EC纠删码环境下对ISA-L使用，实现存储加速，对Ceph的存储数据可恢复性、数据完整性、数据安全性以及加速数据压缩等方面提供了帮助



## ✓ 全闪存

性能测试调优 SPEC验证

eKingCloud验证了 Intel P3500、P3700、P3520、S3500四类SSD在Ceph下的性能

# Notices and Disclaimers

Performance estimates were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system.

Software and workloads used in performance tests may have been optimized for performance only on Intel® microprocessors. Performance tests, such as SYSmark\* and MobileMark\*, are measured using specific computer systems, components, software, operations, and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [intel.com/performance](http://intel.com/performance).

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804

Intel® technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at [intel.com](http://intel.com).

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document. The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software, or configuration may affect your actual performance.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system.

Copyright © 2018 Intel Corporation. All rights reserved. 3D XPoint, Intel, the Intel logo, Intel Optane, VTune and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

# Notices and Disclaimers

Test Configuration: Intel® Xeon® Processor E5-2600v3, E5-2650v3, 10C, 2.3 GHz, M1, Aztec City CRB, 4x8 GB DDR4 2133 MT/s ECC RDIMM. Intel® Xeon® Processor E5-2600v4, E5-2650v4, 12C, 2.2 GHz, M0, Aztec City CRB, 4x8 GB DDR4 2400 MT/s ECC RDIMM. Intel® Xeon® Processor Scalable Family, Platinum 8180 Processor, 28C, 2.5 GHz, H0, Neon City CRB, 6x16 GB DDR4 2666 MT/s ECC RDIMM. BIOS Configuration: P-States: Disabled, Turbo: Disabled, Speed Step: Disabled, C-States: Disabled, Power Performance Tuning: Disabled, ENERGY\_PERF\_BIAS\_CFG: PERF, Isochronous: Disabled, Memory Power Savings: Disabled ISA-L 2.19