

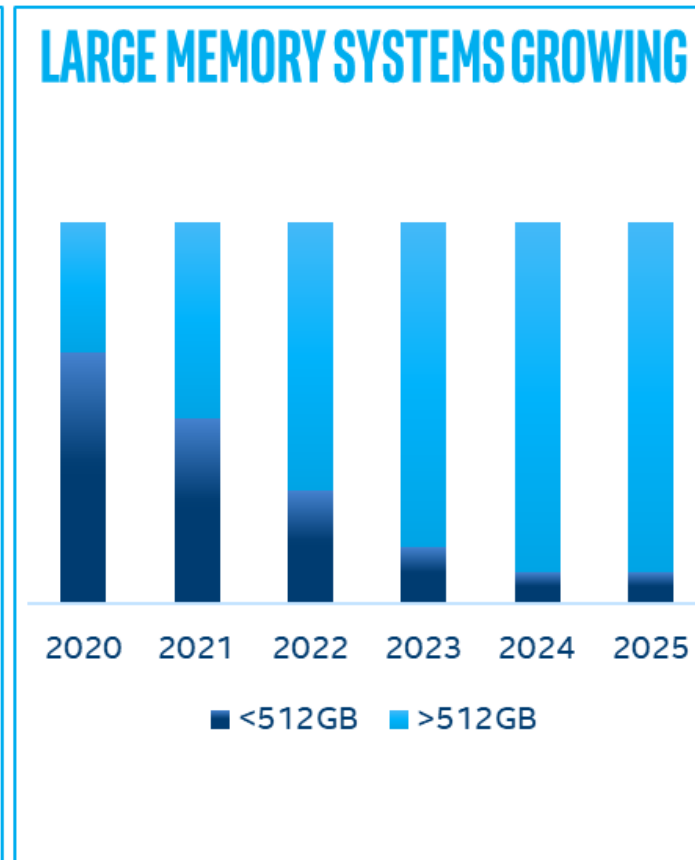


# 英特尔®傲腾™持久内存方案

Wu, Dennis PRC PMem CoE Engineering Manager



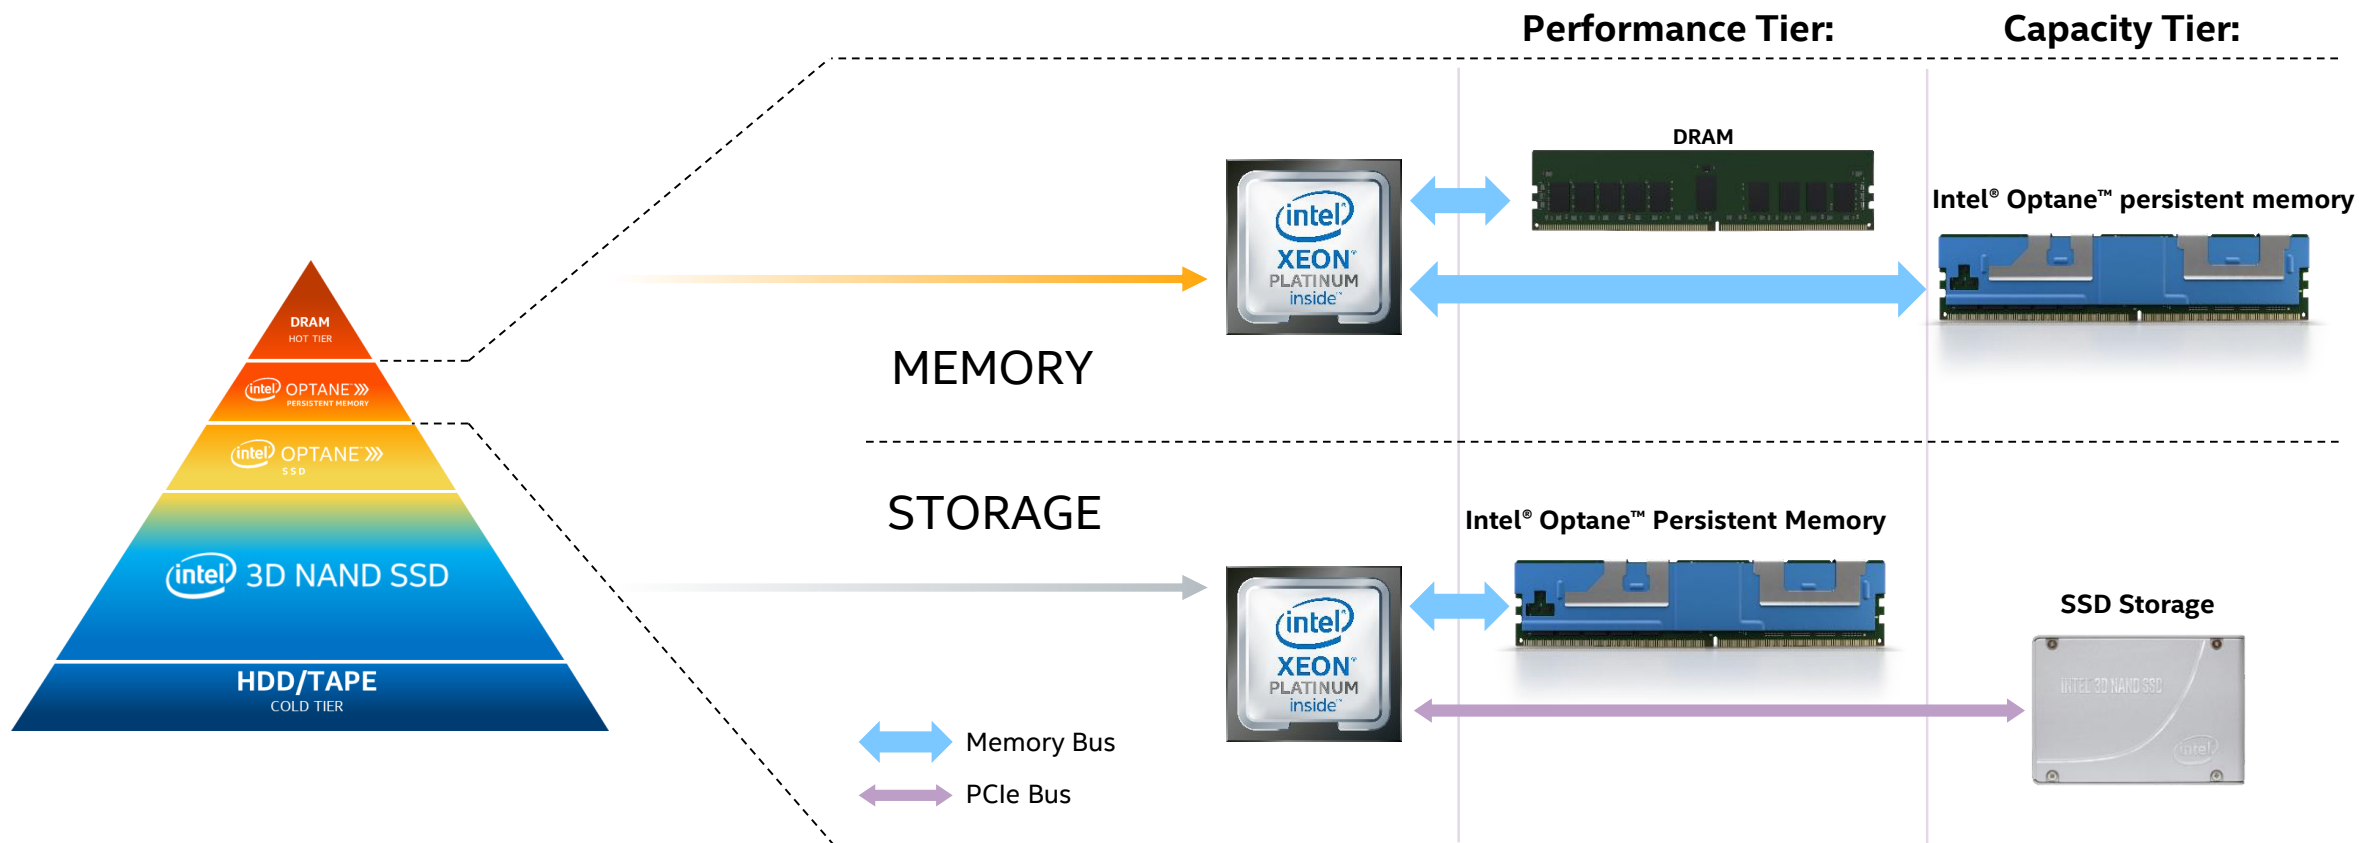
intel®

# 适应变化的数据环境



Source: Intel estimates; "3D NAND Technology – Implications for Enterprise Storage Applications" by J.Yoon (IBM), 2015 Flash Memory Summit; "DRAM Market Monitor Q1-20" by Yole Development

# 持久内存作为内存或存储的愿景



# 解决方案可以...

A deployable hardware/software stack that solves customer problems by...

## → COMMON PAIN POINTS ←

DRAM is too expensive

Scale up is expensive

Not enough capacity

Too many servers

Poor workload performance

Storage is too slow

## USE INTEL® OPTANE™ PERSISTENT MEMORY TO...

### SAVE MORE

DISPLACE DRAM

Systems >512GB

IMPROVE TCO

Workloads that need large memory

### DO MORE

INCREASE MEMORY SIZE

Large memory or SW license fees per core

CONSOLIDATE WORKLOADS

High VMs, with low CPU utilization

### GO FASTER

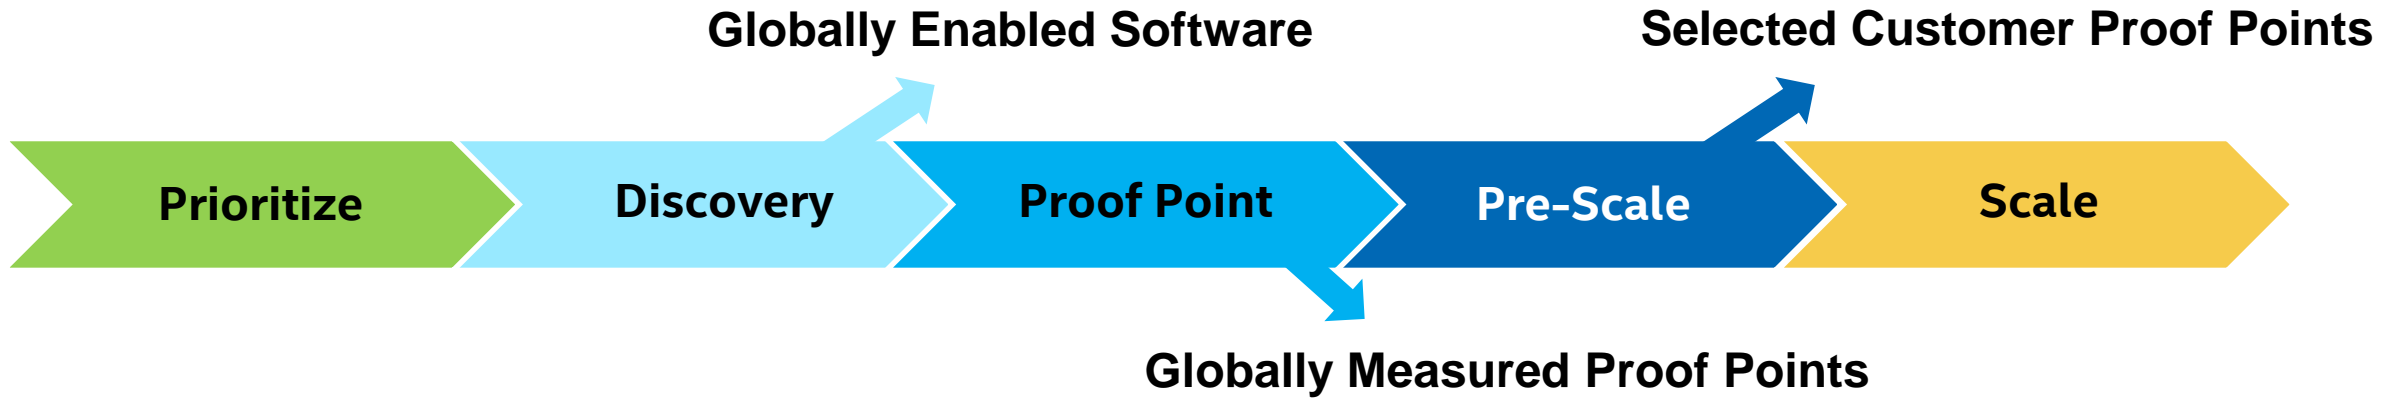
BREAK THE IO BOTTLENECKS

High Disk I/O Traffic

ADD HIGH SPEED STORAGE

Tiered storage subsystem

# Solution readiness examples



Collateral	Data	Format	Frequency	Visibility	How to access?
	Globally Enabled SW	Document – SW Ecosystem Support	~quarterly	Public	Contact your intel representative
	Globally Measured Proof Points	Document – Value Prop guide	~quarterly	NDA & public	
	Selected Customer References	Document – End user references	~monthly	NDA & public	
	+PRC Software Priorities	Slide	As needed	NDA	

# (1/4) SW Ecosystem Support Document

Describes application version, type of software and usage mode for over 40 apps

<b>Application Types</b>	<b>Operating Usage</b>	<b>Number of Applications</b> <i>Not including custom/internal apps</i>
AI / Analytics	App Direct	11 ISVs (e.g. Aerospike, SAP HANA, MemVerge) 5 Open source (e.g. Redis, MemKeyDB, Spark, Cassandra)
Database	Memory Mode	13 ISVs (e.g. HiSign, OracleDB, Gbase, Nari) 2 Open source (Redis, Memcached)
Infra & Storage	Storage	7 ISVs (e.g. Alluxio, AsiaInfo, Hazelcast, Formulus Black) 3 Open Source (RocksDB, Kafka, Percona+RocksDB))

# (2/4) Value Proposition Guide Document

Describes pain point, solution, measurements, HW/SW configuration for over 20 usages

Segments	Categories	Measurements taken	Example Usages	Key Benefits
Cloud Telco	AI / Analytics	<b>10+</b> Inc. Data Robot, Kudu, Spark, AsialInfo etc	Fast analytics Extension cache	<b>Save more:</b> 30-51% lower costs <b>Do more:</b> 1.21-1.66x more ops / second <b>Go faster:</b> 1.27-2X lower latencies
	Databases	<b>11+</b> Inc. SAP HANA, Redis, Memcached, Gbsae etc.	Scale out to scale up Node consolidation Virtualization	<b>Save more:</b> 31-60% lower costs <b>Do more:</b> 1.3-4x more sessions per node <b>Go faster:</b> 80% less DB restart, 2.5x perf
Enterprise & Government	Infrastructure & Storage	<b>4+</b> Inc. Vmware, Oracle Fusion, VDI & Azure Stack	Memory expansion Server consolitation	<b>Save more:</b> 12-28% lower cost per VM <b>Do more:</b> 1.25-2.2x more VMs per node
	Telco Services	<b>2+</b> Inc. CDN	Video streaming	<b>Save more:</b> 37-51% lower memory cost <b>Do more:</b> 2.5x more channels per server

# (3/4) Customer Proof Points

**CUSTOMER TRACTION SINCE LAUNCH**

**OVER 200**  
FORTUNE 500

**OVER 85%**  
POC TO SALE CONVERSION

**OVER 270**  
PRODUCTION WINS

## TCO Savings



**1.3X improvement**  
IN TCO (REDIS)



**30% reduction**  
IN RECOMMENDATION SYSTEM  
& REDIS SERVICE



**15X reduction**  
IN MEMORY FOR IMAGE  
PROCESSING



**22.5%-48% improvement**  
IN TCO (REDIS)



**41% reduction**  
ON INFRASTRUCTURE COST

## Increased Throughput



**~3X improvement**  
IN JOBS PER PHYSICAL HOST  
RATIO



**2.78X increase**  
IN GAMES HOSTED ON A  
SINGLE SERVER



**1.1X increase**  
IN CPU UTILIZATION  
VS. DRAM-ONLY



**~2X**  
VM INSTANTIATION FOR 5G  
MULTI-ACCESS EDGE (REDIS)



**~40% more**  
VMS & CONTAINERS WITHIN  
SAME BUDGET (REDIS)

## Faster Time To Insights



**8X faster**  
SOLVER RUN COMPARED TO  
LUSTRE FILESYSTEM



**80% LATENCY REDUCTION & 3X ACCELERATED INDEXING**  
(ELASTICSEARCH)



**15X faster**  
DATABASE DATA LOAD STARTUP  
(SAP HANA)



**13.7X accelerated**  
DATABASE STARTUP (SAP HANA)



**up to 17X faster**  
STORAGE APPLICATIONS  
(ROCKSDB, MONGODB, MYSQL)



# Solution Focus for ICX/BPS

	<b>Solution Tiers</b>	<b>Mode</b>	<b>Comments</b>
<b>1 – Scalable Solutions</b>	SAP HANA (bare metal, on VMware, on HCI)	AD, MM	Scale out HANA for analytics; exploring Non-Prod HANA in MM
	VDI (Citrix/Horizon on ESXi/vSAN/Nutanix)	MM	Target 2S with 1TB memory capacity
	Oracle Database (MM: 19C, AD: 19C, 21C)	MM & AD	19C: For OLAP WLs, virt DB; New AD-enabled features in 19C, 21C
	Oracle TimesTen	MM & AD	AD-enabled in future release
	MSFT SQL '19 (virtualized)	MM	Virt DB among apps on virt infrastructure
	VMware ESXi/vSAN (MM: VDI, virt apps, AD: SAP HANA)	MM & AD	Also Select Solutions
	Nutanix HCI (VDI, DB, virt apps)	MM & AD	AD-enabling in progress
	Microsoft Azure Stack HCI (DB, VDI, virt apps)	MM & SToAD	Via Select Solutions
	CDN/Visual Cloud Delivery Network	MM & AD	Via Select Solutions
<b>2 – Potential Scalable Solutions</b>	Open Source Redis	MM	Exploring scale beyond PRC to NW CSPs
	MemKeyDB (open source Redis)	AD	PRC community supported open source version of Redis
	MemVerge	AD	Middleware layer virtualizes DRAM & PMem as big memory to apps
	Kx kdb 4.0	AD	
	Aerospike 4.5+ (index) /4.8+ (data+index)	AD	Also enabled for Columbiaville
<b>3 – Launch Proofpoints</b>	Redis Enterprise 5.4	MM	
	MemSQL	MM	
	PRC-specific ISVs (AD: AsialInfo, Kingbase; MM: Neusoft)	MM & AD	
	Amdocs	MM	
	RPMEM + Columbiaville	MM	
	DAOS	AD	
<b>SW Enabling in Progress</b>	SAS Viya	AD	Estimating ~'21 release date
	MySQL, MariaDB	AD	Estimating ~'21 release date
	PostgreSQL	AD	Estimating ~'21 release date

# (4/4) Solution Deployments in PRC

## Open Source & Custom Apps for Cloud Native Environments

Businesses	Services	Enabling Technologies	Customers deploying	Key Benefits
(Video) Recommendation Systems	In-memory caching	Redis, MemKeyDB, Memcached, Aerospike	10+	Lower TCO Higher instance density
	Training/inference	Parameter server	3+	Lower TCO
Search	Indexing	RocksDB & variants	1+	Improved perf/TCO
eCommerce	Storage Caching	RocksDB & variants	3+	Higher read performance/throughput
Payment Transactions	Big data analytics	Spark OAP, Graph Compute	2+	Lower query latency Reduced network dependency

## Enterprise Applications

China maturity & deployment volume to date



- VMware
- Openstack
- VDI
- DeepGlint
- HiSign
- 4Paradigm
- Redis
- BONC
- AsiaInfo
- SAP HANA
- OEM BigData solns
- H3C UIS
- MS SQL Server
- SmartX HCI
- Oracle 19c/20x/X8M
- MemVerge
- Formulus Black



The Intel logo is centered on a solid blue background. It consists of the word "intel" in a white, lowercase, sans-serif font. A small blue square is positioned above the letter "i". To the right of the word "intel" is a registered trademark symbol (®) enclosed in a white circle.

intel®

# Case Studies/Whitepapers

Case Study/Solution Brief/Whitepaper	Related workload(s)	Link
Baidu Feed Stream:	Recommendation Engine Indexing	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/baidu-feed-memory-database-restruction-with-optane-technology.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/baidu-feed-memory-database-restruction-with-optane-technology.html</a>
ZTO Express:	Redis; Virtual Infra	<a href="https://www.intel.com/content/www/us/en/customer-spotlight/stories/zto-express-improves-infrastructure-video.html">https://www.intel.com/content/www/us/en/customer-spotlight/stories/zto-express-improves-infrastructure-video.html</a>
PingAn Cloud:	Redis	<a href="https://www.intel.com/content/www/us/en/customer-spotlight/stories/ping-an-cloud-customer-story.html">https://www.intel.com/content/www/us/en/customer-spotlight/stories/ping-an-cloud-customer-story.html</a>
DeepGlint:	Image search engine/feature retrieval	<a href="https://dpgresources.intel.com/asset-library/very-large-scale-data-governance-case-study/">https://dpgresources.intel.com/asset-library/very-large-scale-data-governance-case-study/</a>
Kuaishou	Redis, short video recommendation system	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/kuaishou-recommendation-system-and-redis-services-storage-upgrade.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/kuaishou-recommendation-system-and-redis-services-storage-upgrade.html</a>
AsialInfo	Telco Billing system	<a href="https://www.intel.cn/content/www/cn/zh/communications/aisware-5g-billing-system-and-network-transformation-for-5g.html">https://www.intel.cn/content/www/cn/zh/communications/aisware-5g-billing-system-and-network-transformation-for-5g.html</a>
Hisign	Fingerprint/Palmprint Identification & Comparison	<a href="https://www.intel.cn/content/www/cn/zh/analytics/accelerate-haixin-finger-and-palm-print-recognition-brief.html">https://www.intel.cn/content/www/cn/zh/analytics/accelerate-haixin-finger-and-palm-print-recognition-brief.html</a>
Baidu BigSQL	SparkSQL	<a href="https://www.intel.cn/content/www/cn/zh/analytics/baidu-bigsql-delivers-faster-spark-interactive-queries.html">https://www.intel.cn/content/www/cn/zh/analytics/baidu-bigsql-delivers-faster-spark-interactive-queries.html</a>
DeepGlint:	Graphic search	<a href="https://www.intel.cn/content/www/cn/zh/analytics/artificial-intelligence/deepglint-search-by-image-engine-solution.html">https://www.intel.cn/content/www/cn/zh/analytics/artificial-intelligence/deepglint-search-by-image-engine-solution.html</a>
China Unicom Wo Cloud	Redis; MySQL	<a href="https://www.intel.cn/content/www/cn/zh/communications/china-unicom-wo-cloud-optimizes-its-data-center.html">https://www.intel.cn/content/www/cn/zh/communications/china-unicom-wo-cloud-optimizes-its-data-center.html</a>
T-systems   HPE	SAP HANA	<a href="https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/t-systems-in-memory-database-customer-story.html">https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/t-systems-in-memory-database-customer-story.html</a>
VK社交网络存储架构	social network storage arch	<a href="https://www.intel.cn/content/www/cn/zh/communications/reimagining-storage-for-social-networks.html">https://www.intel.cn/content/www/cn/zh/communications/reimagining-storage-for-social-networks.html</a>
SmartX 超融合平台	Finance IT system	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/smartx-hci-platform-meet-challenges-in-the-financial-industry.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/smartx-hci-platform-meet-challenges-in-the-financial-industry.html</a>
4th Paradigm AI	AI Integration	<a href="https://www.intel.cn/content/www/cn/zh/analytics/artificial-intelligence/intel-helps-enterprises-meet-the-threshold-of-ai-adoption.html">https://www.intel.cn/content/www/cn/zh/analytics/artificial-intelligence/intel-helps-enterprises-meet-the-threshold-of-ai-adoption.html</a>
Nitrado提升游戏实例密度, 降低玩家成本	Gaming Hosting Server	<a href="https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/nitrado-online-gaming-customer-story.html">https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/nitrado-online-gaming-customer-story.html</a>
提升可扩展性并加速元数据搜索	MetaData Searching	<a href="https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/phoenixnap-panzura-customer-story.html">https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/phoenixnap-panzura-customer-story.html</a>
强化实时数据处理, 助力智能制造	Intelligent Manufacturing	<a href="https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/catl-customer-story.html">https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/catl-customer-story.html</a>
新华三 ONEStor 分布式存储系统	Distributed Storage	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/h3c-onestor-distributed-storage-system.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/h3c-onestor-distributed-storage-system.html</a>
微软Azure Stack HCI方案	HCI	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/transform-data-center-efficiency-with-innovation-solution-brief.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/transform-data-center-efficiency-with-innovation-solution-brief.html</a>
VMware虚拟数据中心方案	Vmware	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/boost-memory-capacity-lower-tco-in-data-centers-solution-brief.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/boost-memory-capacity-lower-tco-in-data-centers-solution-brief.html</a>
ORACLE数据库强化应用优势	Oracle	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/oracle-database-and-analytics-on-intel-architecture-brief.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/oracle-database-and-analytics-on-intel-architecture-brief.html</a>
SAP HANA搭配英特尔傲腾持久内存	SAP HANA	<a href="https://www.intel.cn/content/www/cn/zh/big-data/partners/sap/persistent-memory-and-sap-hana-brief.html">https://www.intel.cn/content/www/cn/zh/big-data/partners/sap/persistent-memory-and-sap-hana-brief.html</a>
利用傲腾持久内存及FPGA优化数据库性能	FPGA	<a href="https://www.intel.cn/content/www/cn/zh/cloud-computing/csp-improve-customer-experience-eguide.html">https://www.intel.cn/content/www/cn/zh/cloud-computing/csp-improve-customer-experience-eguide.html</a>
Aerospike实现高性价比PB级实时计算	Real-time computing	<a href="https://www.intel.cn/content/www/cn/zh/processors/xeon/scalable/software-solutions/aerospike-enterprise-database-brief.html">https://www.intel.cn/content/www/cn/zh/processors/xeon/scalable/software-solutions/aerospike-enterprise-database-brief.html</a>
傲腾持久内存助力AI训练	AI model training	<a href="https://www.intel.cn/content/www/cn/zh/artificial-intelligence/posts/memory-bottleneck-healthcare.html">https://www.intel.cn/content/www/cn/zh/artificial-intelligence/posts/memory-bottleneck-healthcare.html</a>
基于 VMware vSAN 的英特尔® 精选解决方案	Vmware vSAN	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/intel-select-solution-vmware-vsan.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/intel-select-solution-vmware-vsan.html</a>
软银在VMware虚拟机环境中采用傲腾持久内存	telco infrastructure	<a href="https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/softbank-vm-customer-story.html">https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/softbank-vm-customer-story.html</a>
金融AI解决方案落地 - 4PD	financial	<a href="https://www.intel.cn/content/www/cn/zh/analytics/artificial-intelligence/ai-guidebook-fsi-2020.html (P60-70)">https://www.intel.cn/content/www/cn/zh/analytics/artificial-intelligence/ai-guidebook-fsi-2020.html (P60-70)</a>

# Product and Solution Contents

PMem Soluton Contents	Link
Optimized Storage with PMem for Efficient Cloud Services	<a href="https://www.intel.cn/content/www/cn/zh/service-providers/optimized-storage-for-efficient-cloud-services.html">https://www.intel.cn/content/www/cn/zh/service-providers/optimized-storage-for-efficient-cloud-services.html</a>
How Intel Optane DC Persistent Memory Tackles Data Analysis Challenges in Practice	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/optane-dc-persistent-memory-case-gallery.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/optane-dc-persistent-memory-case-gallery.html</a>
Reduce Redis Deployment Cost, Complexity with Intel Optane DC Persistent Memory	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/reduce-redis-deployment-cost-and-complexity-infographic.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/reduce-redis-deployment-cost-and-complexity-infographic.html</a>
Cut Infrastructure Complexity and Cost for Redis	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/cut-infrastructure-complexity-and-cost-for-redis.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/cut-infrastructure-complexity-and-cost-for-redis.html</a>
Improve customer experience in B2B SaaS with optimized database performance	<a href="https://www.intel.cn/content/www/cn/zh/service-providers/improve-b2b-saas-with-optimized-database-performance.html">https://www.intel.cn/content/www/cn/zh/service-providers/improve-b2b-saas-with-optimized-database-performance.html</a>
How to Save Big on Redis Deployments	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/how-to-save-big-on-redis-deployments.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/how-to-save-big-on-redis-deployments.html</a>
Application Performance Verification Highlights	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/application-performance-verification-highlights-of-intel-optane-dc-persistent-memory.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/application-performance-verification-highlights-of-intel-optane-dc-persistent-memory.html</a>
The Challenge of Keeping Up With Data	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/the-challenge-of-keeping-up-with-data.html?cache=1111">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/the-challenge-of-keeping-up-with-data.html?cache=1111</a>
英特尔® 傲腾™ 技术助力 SDS 释放人工智能无限潜力 video	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/accelerate-ai-in-sds-with-intel-optane-memory-and-storage-hierarchy.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/accelerate-ai-in-sds-with-intel-optane-memory-and-storage-hierarchy.html</a>
如何利用英特尔傲腾持久内存扩展虚拟化环境的内存容量并降低成本	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/expanding-memory-capacity-in-virtualized-environments-business-brief.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/expanding-memory-capacity-in-virtualized-environments-business-brief.html</a>
利用英特尔傲腾持久内存提升可扩展性并加速元数据搜索	<a href="https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/phoenixnap-panzura-customer-story.html">https://www.intel.cn/content/www/cn/zh/customer-spotlight/stories/phoenixnap-panzura-customer-story.html</a>
基于英特尔至强平台和傲腾持久内存定义内存数据库的未来	<a href="https://www.intel.cn/content/www/cn/zh/products/docs/memory-storage/optane-persistent-memory/accenture-partner-whitepaper.html">https://www.intel.cn/content/www/cn/zh/products/docs/memory-storage/optane-persistent-memory/accenture-partner-whitepaper.html</a>
Communications services providers (CoSPs) faq sheet	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/optane-faq-sheet.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/optane-faq-sheet.html</a>
PMem Marketing video	<a href="https://videoportal.intel.com/media/0_to5z24cm">https://videoportal.intel.com/media/0_to5z24cm</a>
PMem Technical video	<a href="https://videoportal.intel.com/media/0_nq36vwgm">https://videoportal.intel.com/media/0_nq36vwgm</a>
PMem 200 series Product Brief	<a href="https://www.intel.cn/content/www/cn/zh/architecture-and-technology/pmem-product-brief-better-data-work.html">https://www.intel.cn/content/www/cn/zh/architecture-and-technology/pmem-product-brief-better-data-work.html</a>