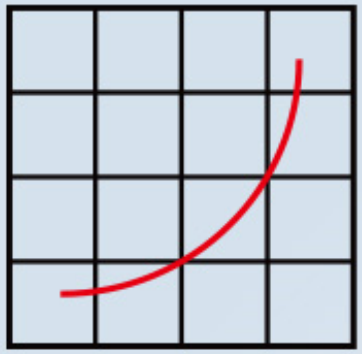

spec®

SPEC 2016亚洲峰会
SPEC 2016 ASIA SUMMIT

SPEC Measures and Needs of Chinese Market

Panelists:
Mathew Colgrove, NVIDIA/PGI
David Reiner, AMD
Arthur Kang, Inspur
Chunyu Jiang, CAICT

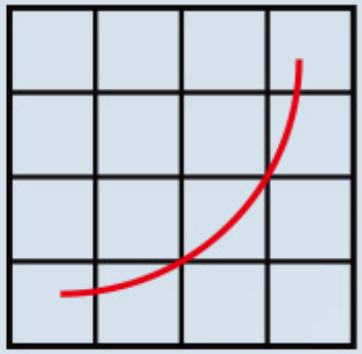

SPEC 2016亚洲峰会
SPEC 2016 ASIA SUMMIT

SPEC High Performance Group (HPG)

Mathew Colgrove, NVIDIA/PGI

SPEC HPG

Develops benchmarks to represent high-performance computing applications for standardized, cross-platform performance evaluation.

Benchmarks

- SPEC OMP2012
- SPEC MPI2007
- SPEC ACCEL

SPEC MPI2007

The **SPEC MPI® 2007** benchmark suite is for evaluating MPI-parallel, floating point, compute intensive performance across a wide range of cluster and SMP hardware**.**

SPEC MPI® 2007 focuses on performance of compute intensive applications using the Message-Passing Interface (MPI), which means this benchmark emphasizes the performance of:

- the type of computer processor (CPU),
- the number of computer processors,
- the MPI Library,
- the communication interconnect,
- the memory architecture,
- the compilers, and
- the shared file system.

SPEC OMP2012

The **SPEC OMP®2012** benchmark is designed for measuring performance using applications based on the OpenMP 3.1 standard for shared-memory parallel processing. The benchmark also includes an optional metric which includes power measurement.

The benchmark includes 14 scientific and engineering application codes, covering everything from computational fluid dynamics (CFD) to molecular modeling to image manipulation.

SPEC ACCEL

SPEC ACCEL provides a comparative performance measure

- Hardware Accelerator devices (GPU, Co-processors, etc.)
- Supporting software tool chains (Compilers, Drivers, etc.)
- Host systems and accelerator interface (CPU, PCIe, etc.)

Computationally-intensive parallel High Performance Computing (HPC) applications, benchmarks, and mini-apps

Portable across multiple accelerators

Two distinct suites

- OpenACC (API v1.0)
- OpenCL (API v1.1)

SPEC HPG Future Directions

SPEC ACCEL near future

- OpenMP 4.5 versions

Next generation SPEC ACCEL

- Hybrid MPI+Accelerator to measure large scale Accelerated systems

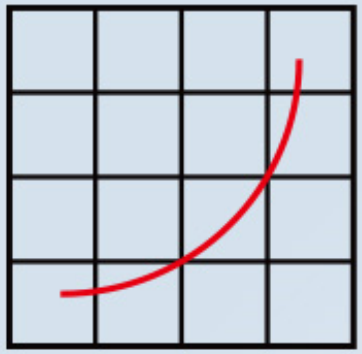

SPEC 2016亚洲峰会
SPEC 2016 ASIA SUMMIT

SPEC Graphics and Workstation Performance Group (SPEC/GWPG)

David Reiner

SPEC/GWPG

Worldwide leader in standardized performance evaluation for professional graphics and workstation systems

Benchmarks used by leading vendors, companies, user and R&D organizations, RFP issuers, analysts and publications to measure and compare workstation performance.

Benchmarks available for free download*

Approximately 10,000 benchmark downloads annually

More than 650 media mentions annually with potential audience of 2.63 billion

Three subcommittees: SPECapc, SPECgpc, SPECwpc

* to everyone except hardware and service providers using for financial gain

SPEC/GWPG

SPECapc (SPEC Application Performance Characterization)

Performance evaluation using popular workstation applications for CAD/Engineering and Media & Entertainment markets.

Measures total system performance, including graphics, CPU and I/O.

Current SPECapc benchmarks for PTC Creo, Siemens NX, Dassault Systemes SolidWorks, Autodesk 3ds Max and Maya.

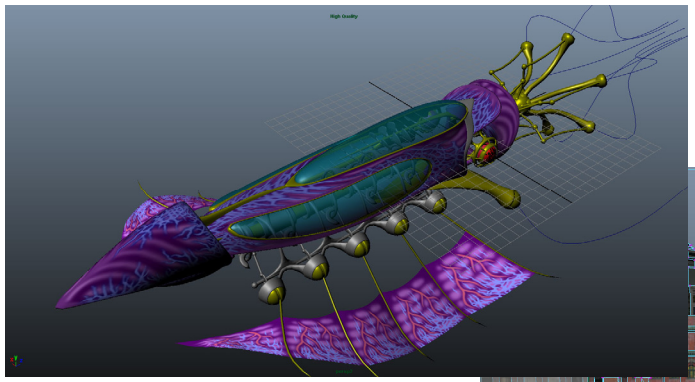

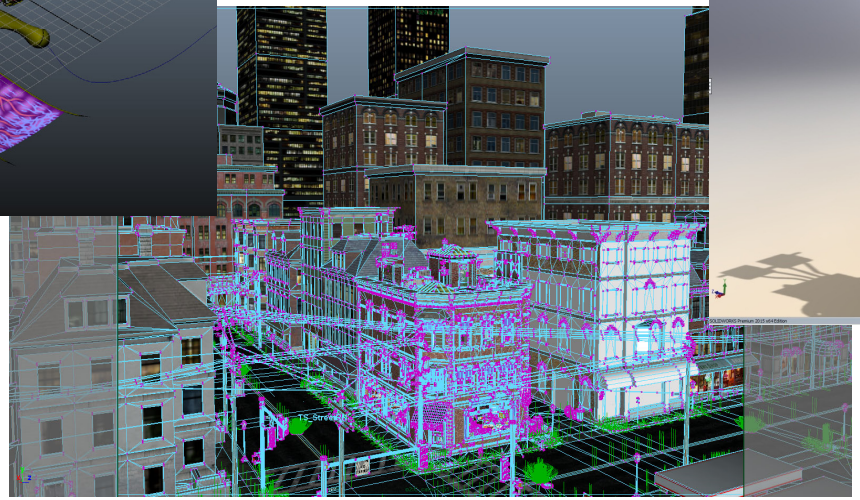

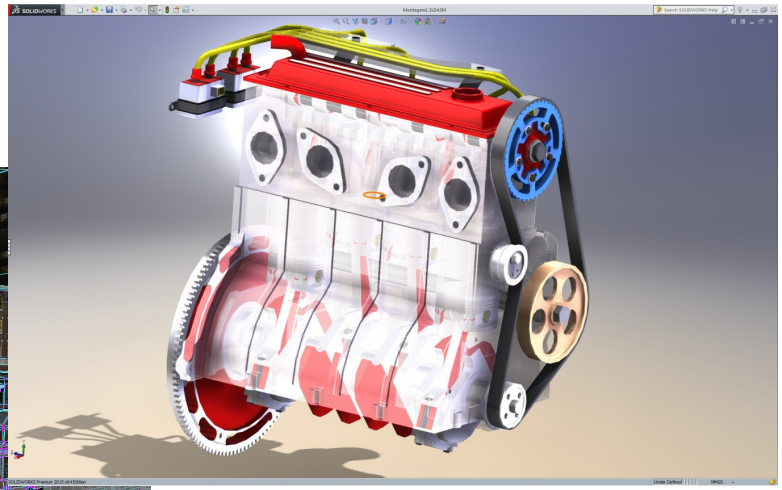

SPEC/GWPG

SPECgpc (SPEC Graphics Performance Characterization)

Provides standardized method for comparing graphics performance based on professional applications across vendor platforms

Develops and provides SPECviewperf 12.1, which measures the 3D graphics performance of viewsets based on popular professional applications that use various APIs, including OpenGL and DX11

No application installation required, easy to run, and available for free download

SPEC/GWPG

SPECwpc (SPEC Workstation Performance Characterization)

Measures key aspects of workstation performance based on diverse professional applications.

More than 30 workloads measuring CPU, graphics, I/O, and memory performance.

No application installations required, easy to run, and available for free download

Application categories include:

- Media and entertainment (3D animation, rendering)
- Product development (CAD/CAM/CAE)
- Life sciences (medical, molecular)
- Financial services
- Energy
- General operations

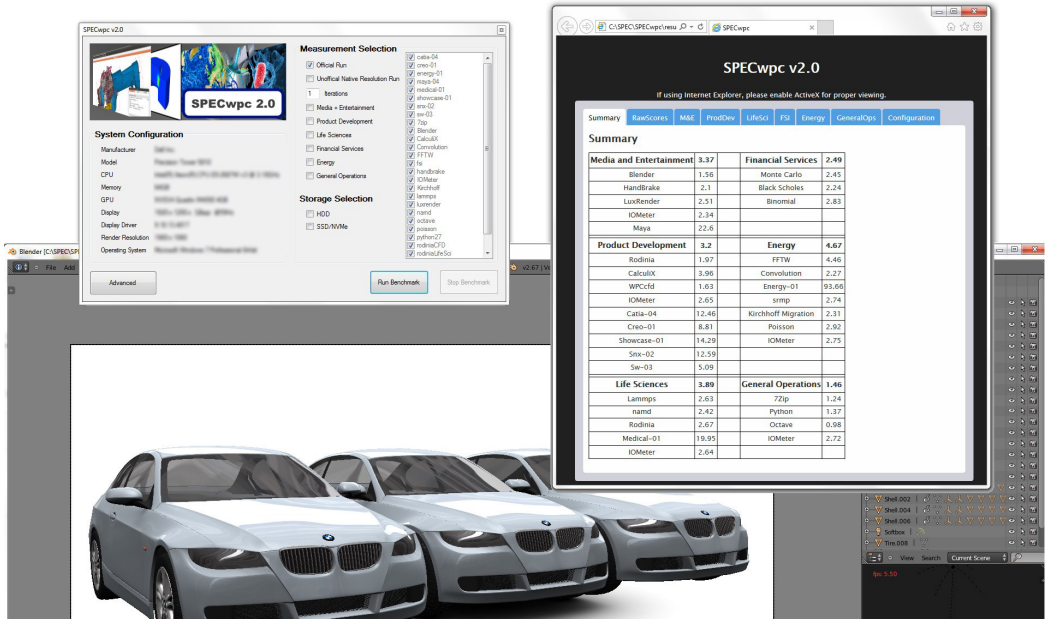

SPEC/GWPG

Looking ahead…

Continue to update SPECapc benchmarks for 3ds Max, Maya, SolidWorks, PTC Creo with the latest application versions and new features.

Continue to update SPECviewperf benchmark with updated and all-new application-based viewsets.

New SPECwpc benchmark with new and improved workloads and storage performance measurement.

More graphics and workstation benchmarking articles and blog posts -- Join the Graphics and Workstation Benchmarking group on LinkedIn for the latest updates: https://www.linkedin.com/groups/8534330

SPEC 2016亚洲峰会
SPEC 2016 ASIA SUMMIT

What Could SPEC Benchmark Do for Your Products?

Arthur Kang, Inspur

SPEC CPU Subcommittee Member

When should we use the Benchmarks?

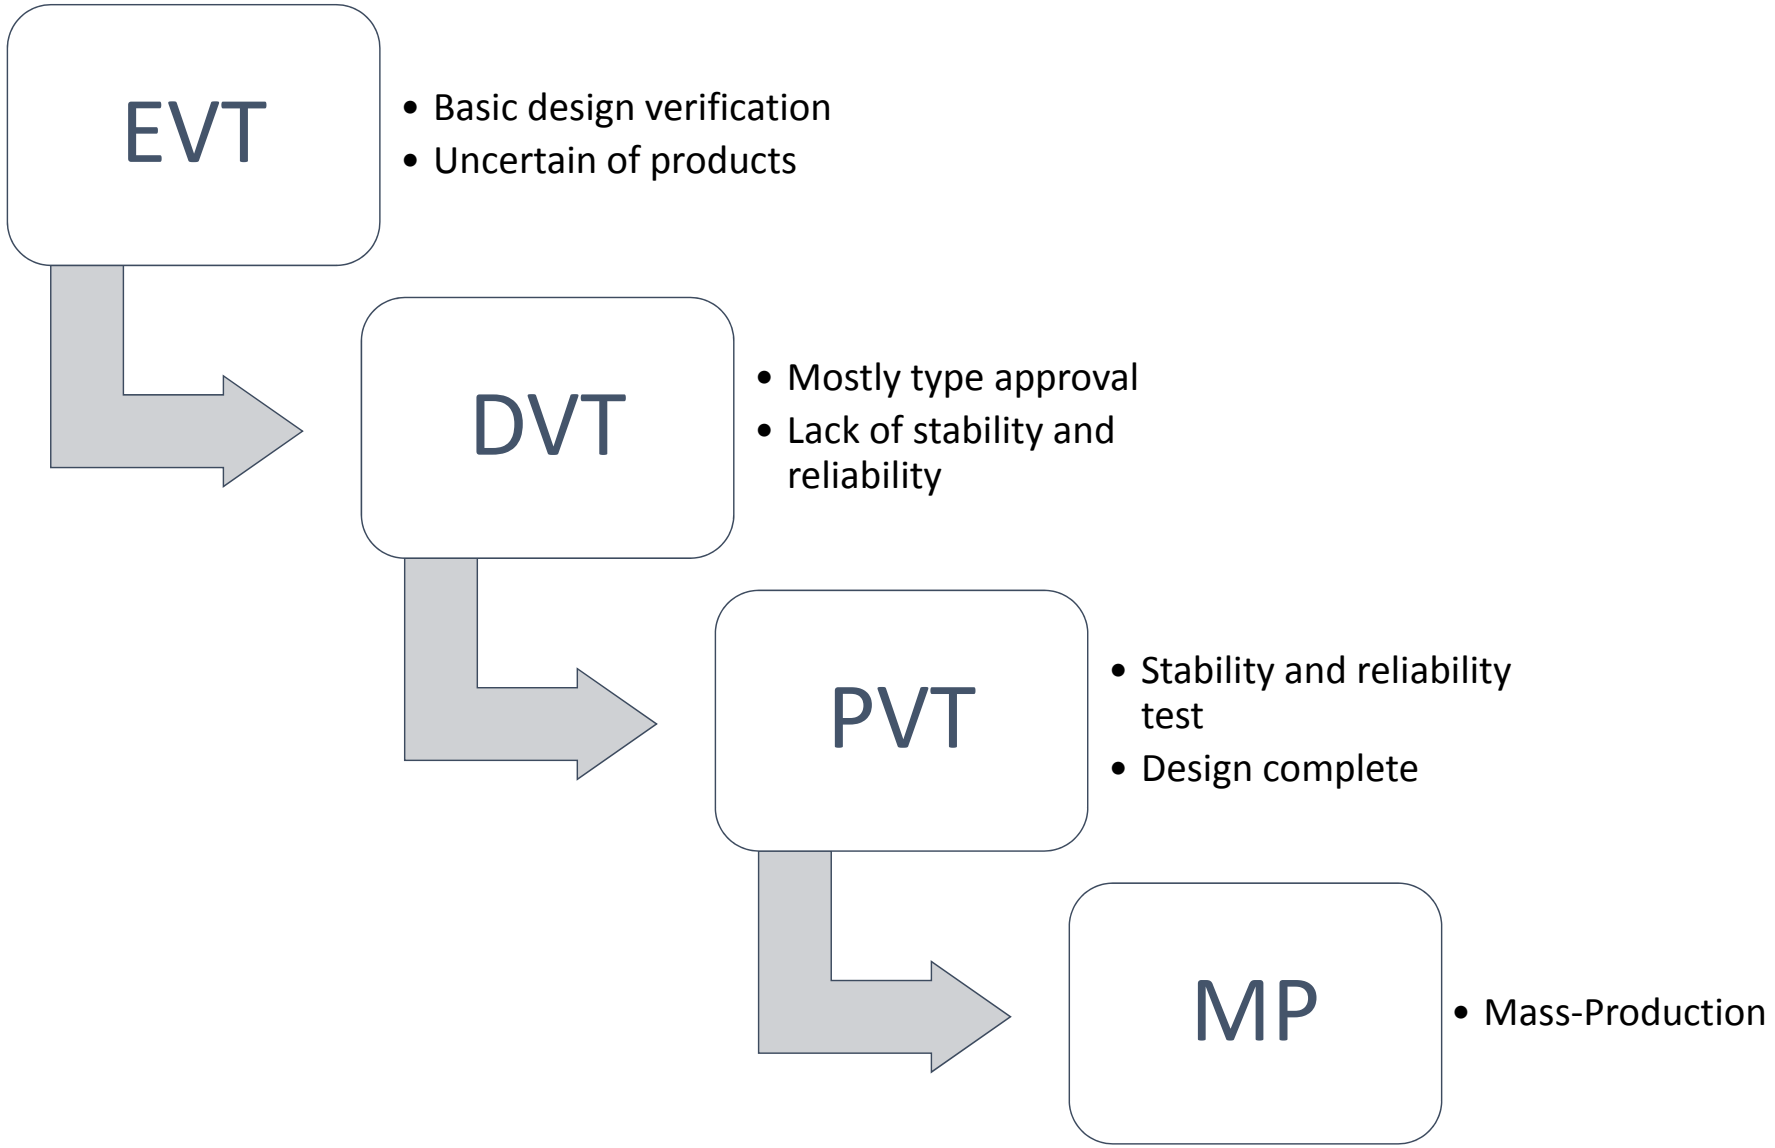

Where do we need for measurement?

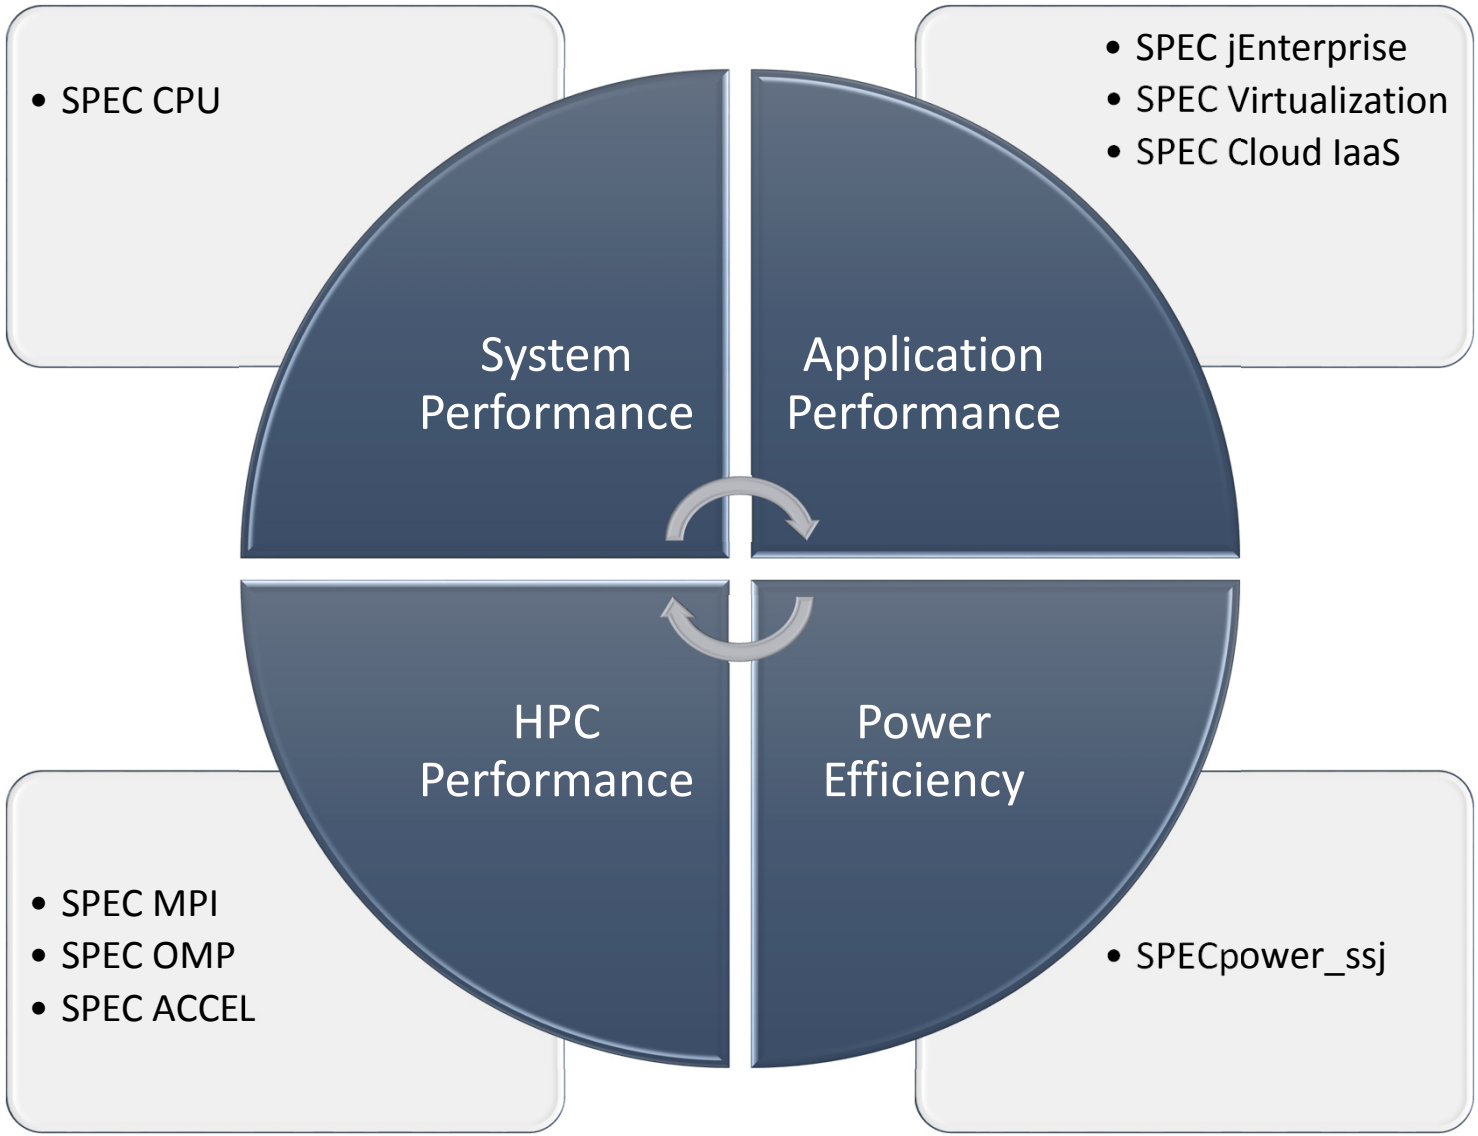

What Could Benchmarks Effect?

	DVT	PVT	MP
SPEC CPU2006 int rate Base	1678	1702	1702
Idle Power (Watt)	92	83	84
Max Power (Watt)	495	482	441

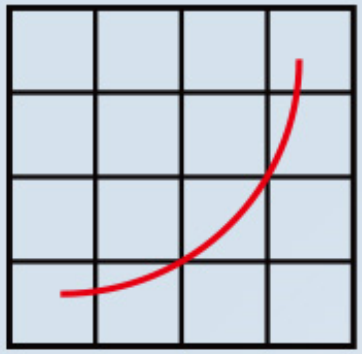

spec®

SPEC 2016亚洲峰会
SPEC 2016 ASIA SUMMIT

Big Data Benchmark and Standard Practice in China:
What does Chinese customers want?

Chunyu Jiang
Engineer in CAICT
10/28

Open source is a different beast

- Big data infrastructure technology is driven by open source community.
- Open Source ≠ commercialize products
  - The open-source software is not being straightforward to use.
  - Not stable
  - Security problem
  - Not user friendly
  - Lack of Support service

Open source is a different beast

Vender

- Lack a big data benchmarking standard to regulate the market.
- Need some domain-oriented benchmarking tools.
- Get bored of duplicated POC tests.

Evaluation Metrics

User needs oriented
Turn complicated products into easy understandable criteria

User

- Requirement is vague.
- Technology selection is difficult.
- Productions evaluation is costly.
- Daily operations is tedious.

There is a huge information gap between users and venders.

Motivation

- Evaluate the performance of different big data products and help users to choose them in a cost-effective way.
- Help venders to better understand the users' requirements.
- Help to mature big data products.
- Reduce the cost of POC test.

Methodology

- Define a big data benchmarking standard involving all related stakeholders.
- Set up a lightweight big data benchmarking test.
- Develop or utilize a big data benchmarking tool which covers many of application scenarios.

Big Data Products Certification

Certification

Basic Ability

Metric-oriented

- Function
- Operation and maintenance
- High Availability
- Security
- Compatibility
- Scalability
- Multi-tenant

Performance

Scenario-oriented

- Bach processing workloads
- SQL workloads
- Nosql workloads
- Machine learning workloads
- Graph processing workloads (future work)
- Streaming workload (future work)

Basic Ability

Operation / Maintenance	High Availability	Function	Compatibility	Security	Multi-tenant	Scalability
Auto deployment	Primary Namenode Failure Recovery	Data Import	ODBC compatible	Certification	Tenant Management	Cluster Extension
Resource monitoring	Standby Namenode Failure Recovery	SQL workload	JDB compatible	Authorization	Resource Management	Cluster Contraction
Job moinitoring	Datenode Failure Recovery	NoSQL workload	SQL supportive	Encryption	Resource isolation	
Cluster operations	Hmaster Failure Recovery	Machine Learning	Database synchronization	Audit	Permission Management	
Failure management	RegionServer Failure Recovery	Stream workload	Across different database tables associated actions			
Log management	HDFS Backup					
Configuration management	HBase Backup					
Authorization management	Management node failure recovery					
User Management						
Rolling upgrade						

Basic Ability is consist of seven metrics with 38 user cases in total.

Performance Workload

SQL Workload	NoSQL Workload	Machine Learning	Bach processing
I/O intensive	Load	Kmeans	Terasort
CPU intensive	95% Read，5% Write	Bayes	
Reporting	50% Read, 50% Write		
Data Mining	Read-modify-write		
Interactive			

Performance certification has 12 user cases which cover four kinds of scenarios : SQL workloads, NoSQL workloads, machine learning workloads and batch processing workloads.

Audit

Before Test

- Version check
- Test tools check
- Clear cache

Input Check

- Data size
- Table row and column check
- Data content
- Create Table query check
- Replication check
- Shell Script check

Processing Check

- Workload is running
- Cluster resource utilization monitoring

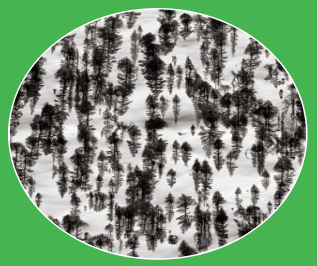

Output Check

- Record running time
- Output size
- Output check

File Retain

- Key jar file
- Job-history
- Tuning parameter

Roadmap

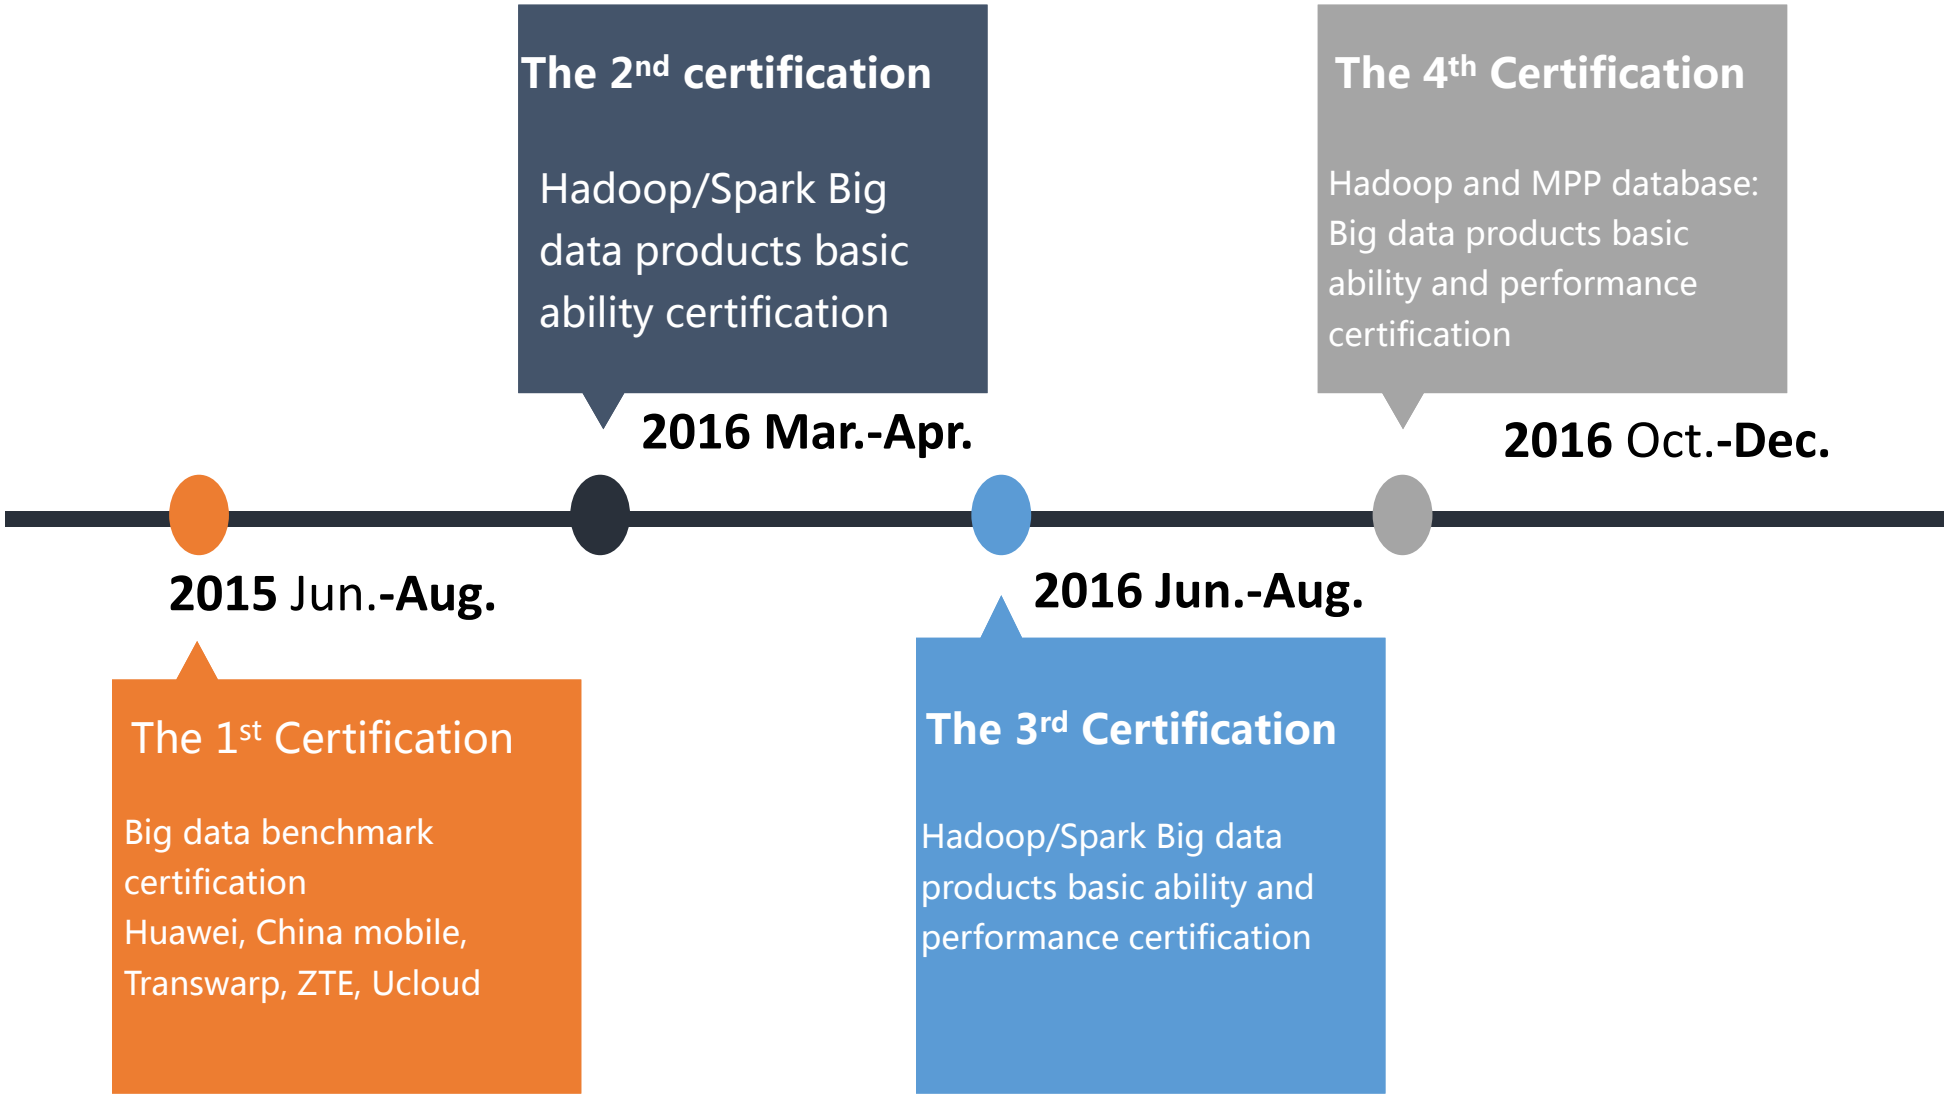

Products pass the certifications

The 1st Certification	The 2nd Certification	The 3rd Certification	
Big data Benchmark	**Big Data Products Basic Ability**	**Big Data Products Basic Ability**	**Big Data Products Performance**
HuaWei FusionInsight	SeaBox Data SeaBox Big Data platform	Baifendian BD-OS	H3C H3C DataEngine
China Mobile BC-Hadoop	Minglamp Minglamp Data Platform	Gridsum Gridsum Big Data Platform	Tencent Cloud Data Intelligence
Transwarp Transwarp Data Hub	Byitgroup Super Center Big Data Platform		SeaBox SeaBox Big Data Platform
ZTE Golden Data	H3C H3C DataEngine		Transwarp Transwarp Data Hub
Ucloud	Transwarp Transwarp Data Hub		Baifendian BD-OS
	Tencent Cloud Data Intelligence		

Companies participate the certification

The 1st Certification

The 2nd Certification

The 3rd Certification

Works we have done

- Proposed a technical specification on big data benchmarking in CCSA
- Set up a 32-nodes big data test environment
- Integrated a big data benchmarking tool
- Have completed the user guideline and use cases
- Finish three batch of Big data products certification

Future Plan

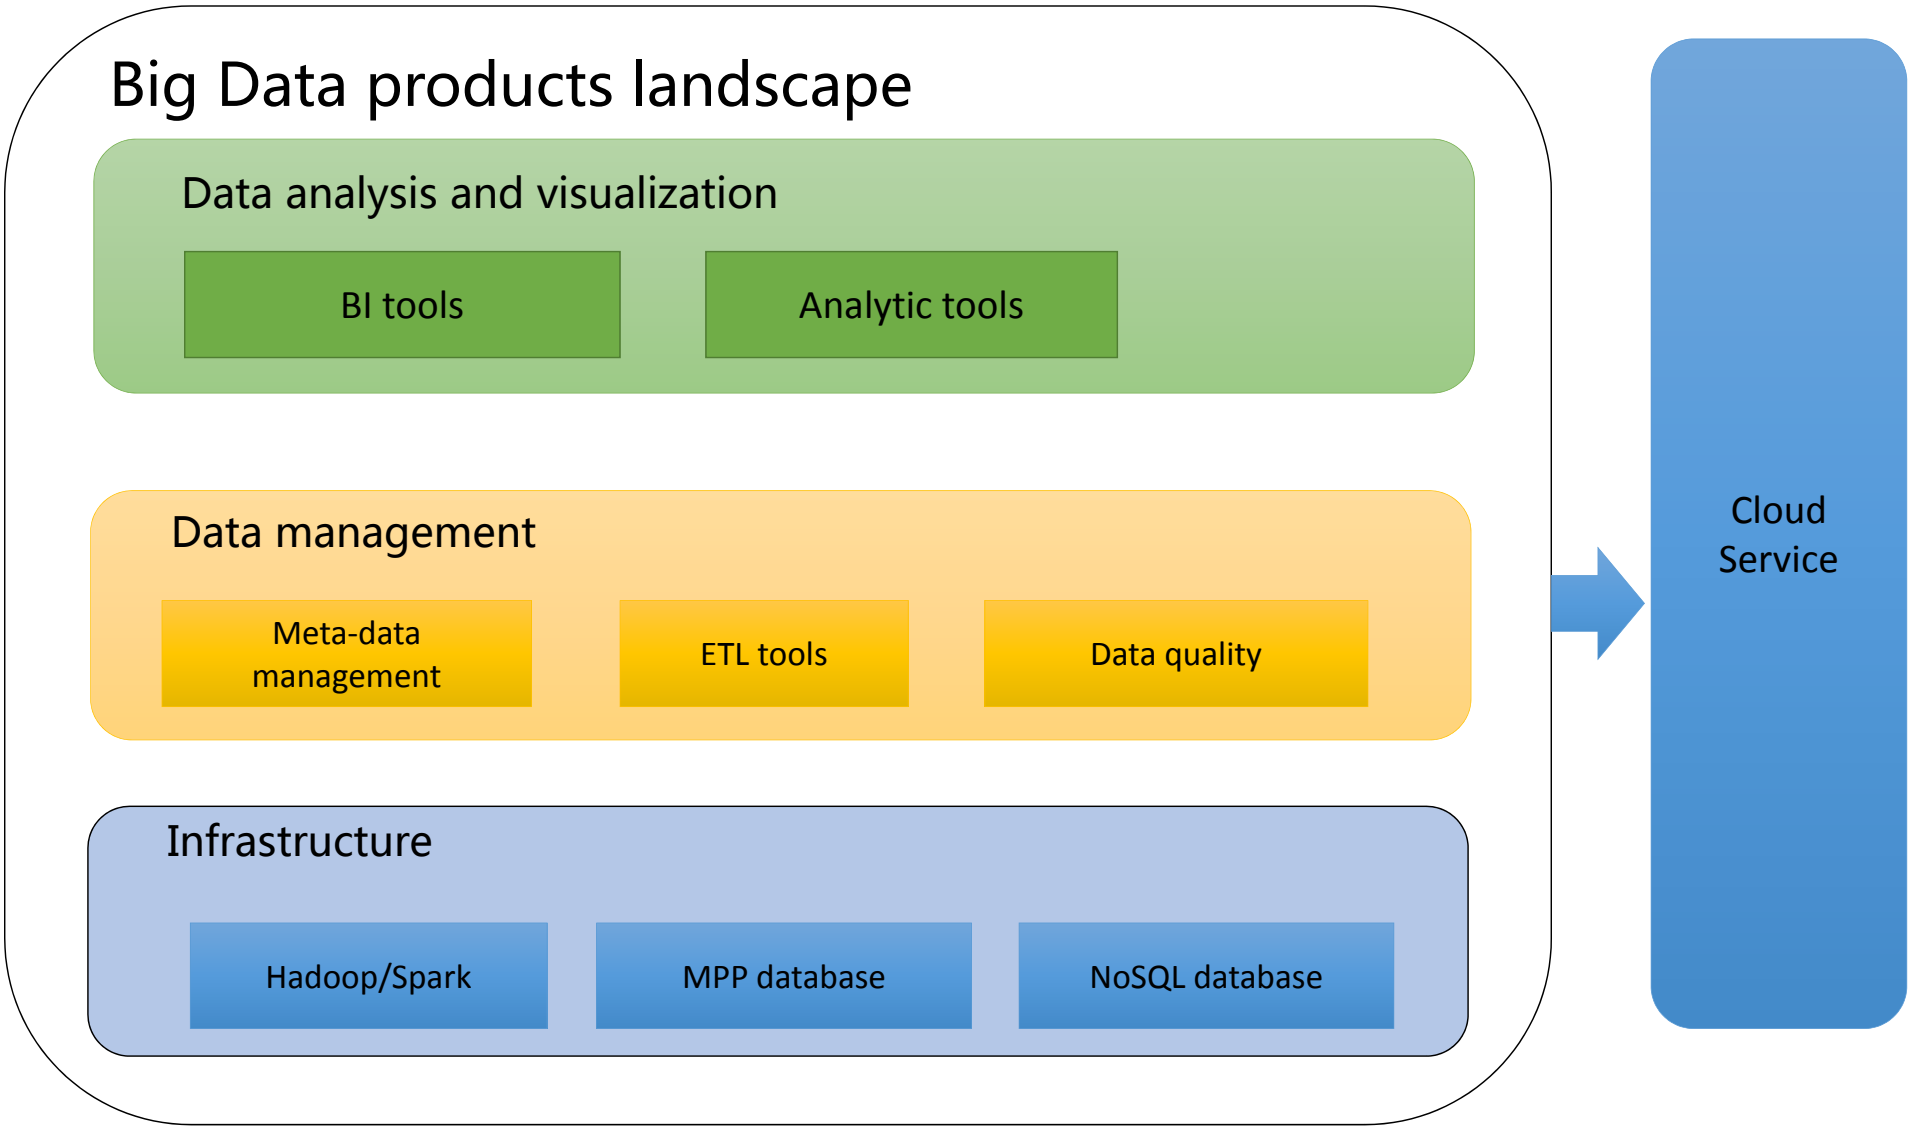

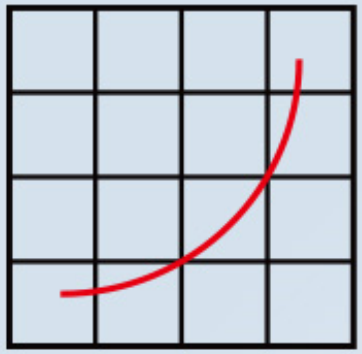

spec®

SPEC 2016亚洲峰会
SPEC 2016 ASIA SUMMIT

Q&A

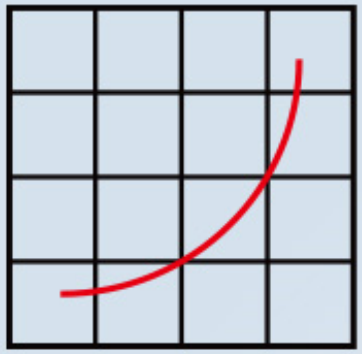

spec®

SPEC 2016亚洲峰会
SPEC 2016 ASIA SUMMIT

Thank you!

info@spec.org
www.spec.org