AWS Summit

AWS 技术峰会 · 北京 2014

DynamoDB在高性能广告平台大数据需求下的实践

赵堃亮 木瓜移动架构师

2014年12月12日

提纲

- 木瓜移动简介
- DynamoDB介绍
- DynamoDB设计模式
- DynamoDB 在AppFlood中的应用

PAPAYA

papayamobile.com

木瓜移动创立于2008年，
总部位于北京，现有近200名员工

木瓜产品

AppFlood

Papaya Games

木瓜移动于2012年推出了AppFlood，AppFlood现在已经发展成为中国最大的程序化移动广告平台。目前AppFlood已经接入了82,000款应用，并为广告主提供了每天超过8亿次的展示和曝光的机会。

amazon web services

Dynamo: Amazon's Highly Available Key-value Store

Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall and Werner Vogels

Amazon.com

ABSTRACT

Reliability at massive scale is one of the biggest challenges we face at Amazon.com, one of the largest e-commerce operations in the world; even the slightest outage has significant financial consequences and impacts customer trust. The Amazon.com platform, which provides services for many web sites worldwide, is implemented on top of an infrastructure of tens of thousands of servers and network components located in many datacenters around the world. At this scale, small and large components fail continuously and the way persistent state is managed in the face of these failures drives the reliability and scalability of the software systems.

This paper presents the design and implementation of Dynamo, a highly available key-value storage system that some of Amazon's core services use to provide an "always-on" experience. To achieve this level of availability, Dynamo sacrifices consistency under certain failure scenarios. It makes extensive use of object versioning and application-assisted conflict resolution in a manner that provides a novel interface for developers to use.

Categories and Subject Descriptors

D.4.2 **[Operating Systems]**: Storage Management; D.4.5 **[Operating Systems]**: Reliability; D.4.2 **[Operating Systems]**: Performance;

General Terms

Algorithms, Management, Measurement, Performance, Design,

One of the lessons our organization has learned from operating Amazon's platform is that the reliability and scalability of a system is dependent on how its application state is managed. Amazon uses a highly decentralized, loosely coupled, service oriented architecture consisting of hundreds of services. In this environment there is a particular need for storage technologies that are always available. For example, customers should be able to view and add items to their shopping cart even if disks are failing, network routes are flapping, or data centers are being destroyed by tornados. Therefore, the service responsible for managing shopping carts requires that it can always write to and read from its data store, and that its data needs to be available across multiple data centers.

Dealing with failures in an infrastructure comprised of millions of components is our standard mode of operation; there are always a small but significant number of server and network components that are failing at any given time. As such Amazon's software systems need to be constructed in a manner that treats failure handling as the normal case without impacting availability or performance.

To meet the reliability and scaling needs, Amazon has developed a number of storage technologies, of which the Amazon Simple Storage Service (also available outside of Amazon and known as Amazon S3), is probably the best known. This paper presents the design and implementation of Dynamo, another highly available and scalable distributed data store built for Amazon's platform. Dynamo is used to manage the state of services that have very

DynamoDB 是一种全托管的 NoSQL 数据库服务

amazon web services

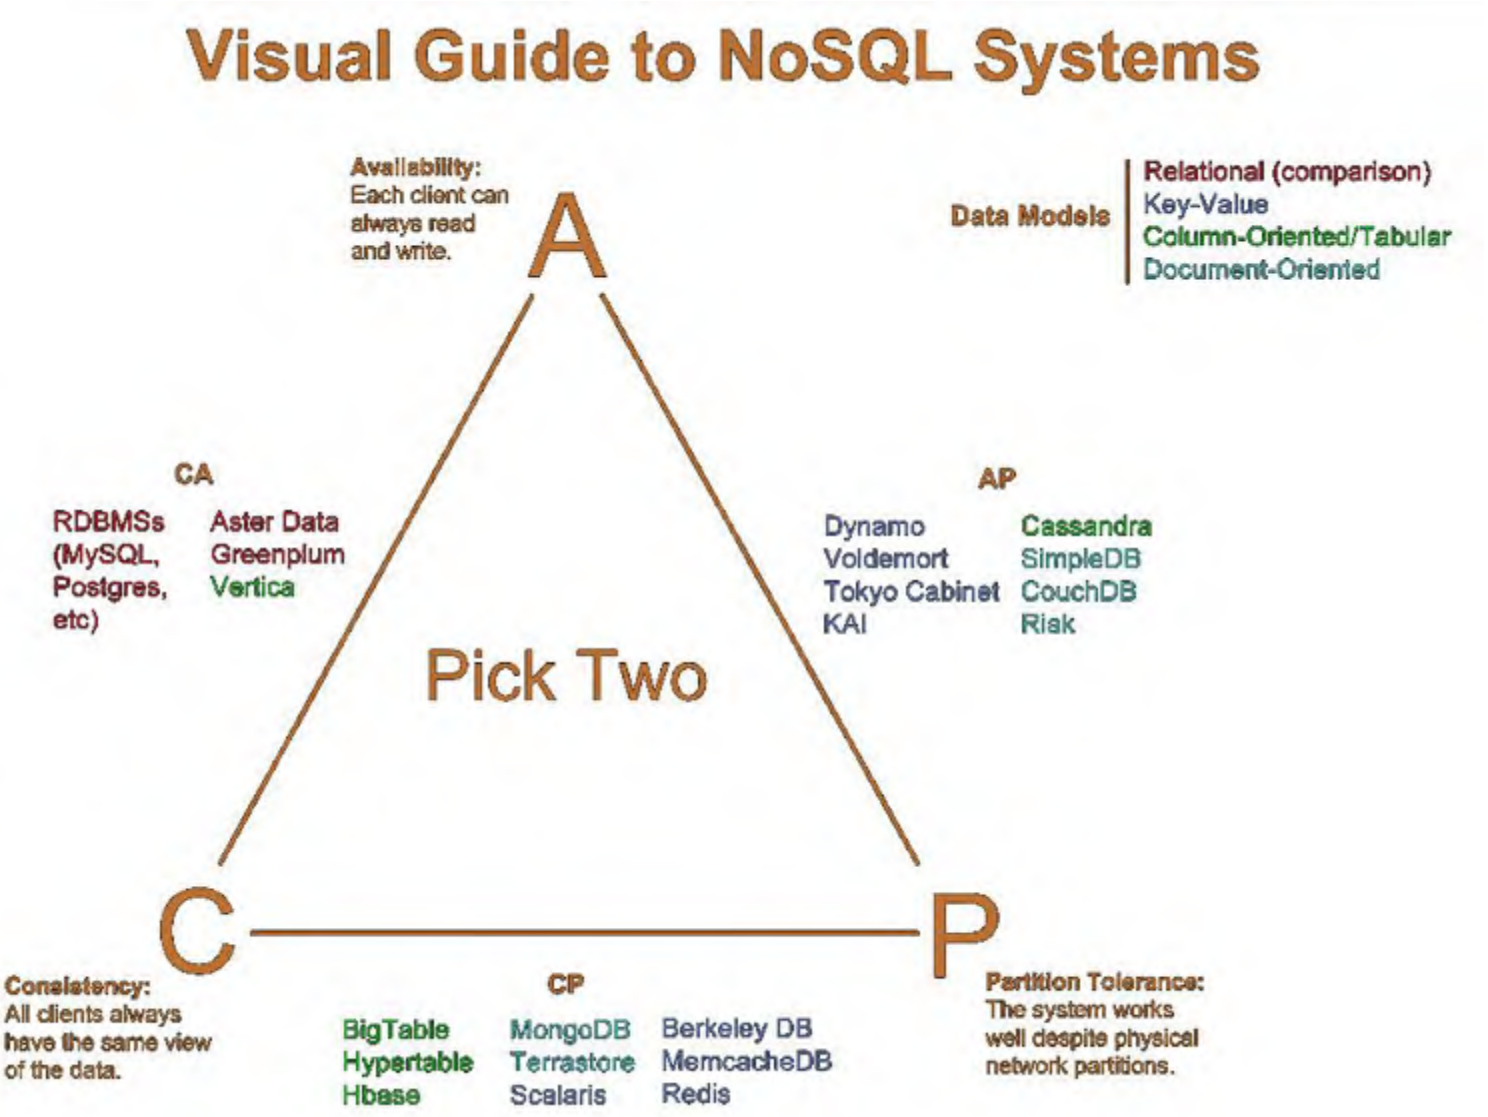

NoSQL数据库的CAP理论

DynamoDB的优势

- 没有运营和维护的负担
- 快速而可预测的性能
- 实现无缝扩展
- 关注自己的产品

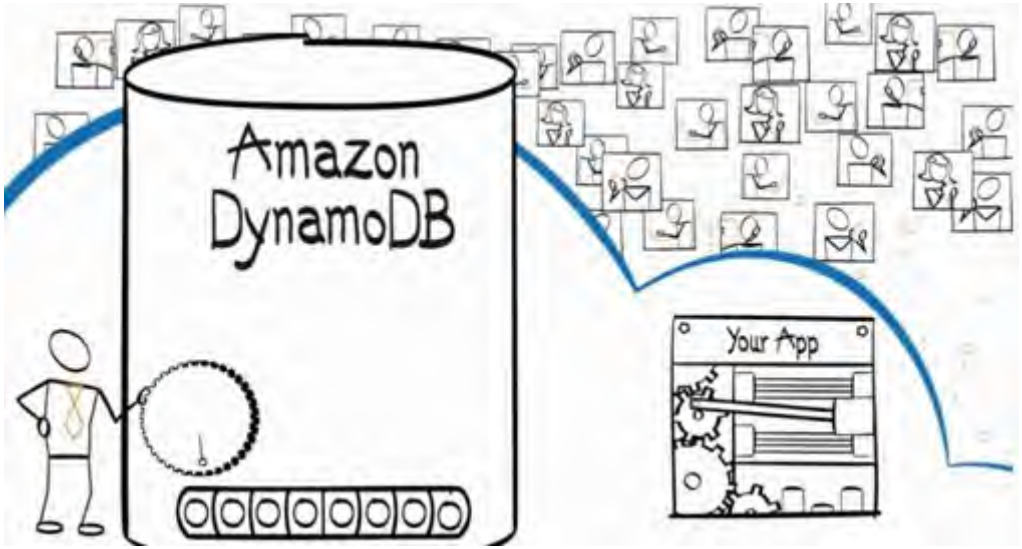

DymamoDB基本概念

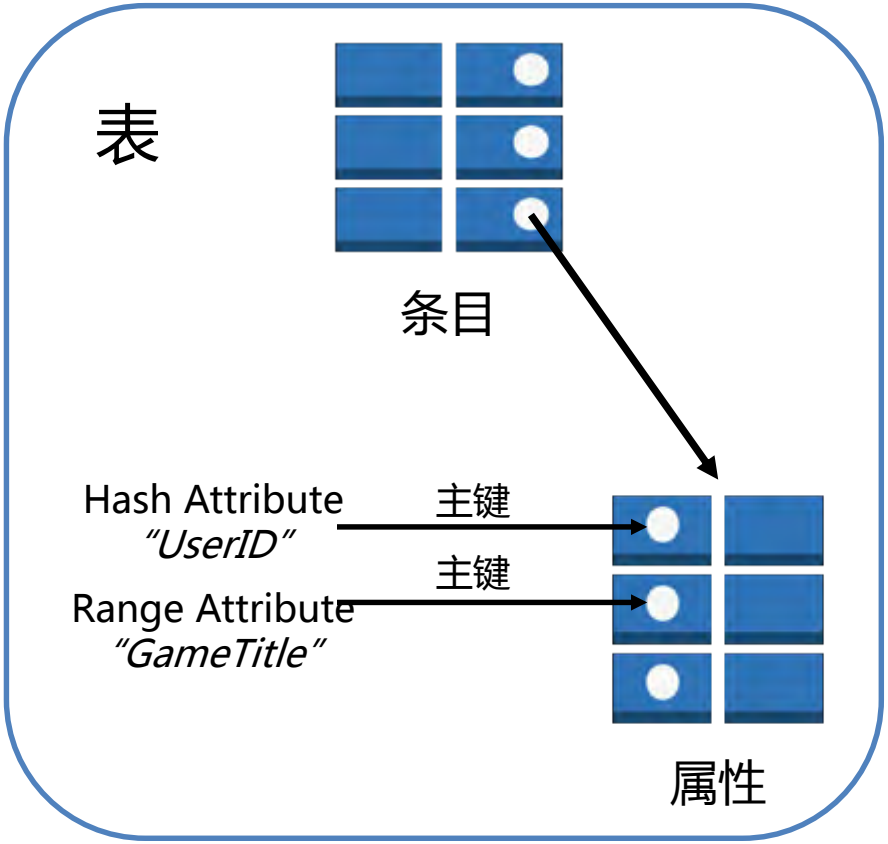

- 不需要严格的表结构定义
- 提供主索引和二级索引
- 可以用使用主键读取条目
- 可以查询和扫描表数据
- 所有接口都提供HTTPS访问

表设计

- 选择合适的主键

User Profile Table				
Hash Key	**Range Key**	**attr. 1**	**…**	**attr. n**
user_id	game_id			

Request Log Table				
Hash Key	**Range Key**	**attr. 1**	**…**	**attr. n**
http_code	timestamp			

预配置吞吐量

- 写吞吐量
  - 1个写单元 = 1 KB数据大小 × 次数 / 秒
  - 并发操作，保证条目级别的ACID
- 读吞吐量
  - 1个读单元 = 4 KB数据大小 × 次数 / 秒 (强一致性)
  - 强一致性 vs 最终一致性

存储大条目

- 将一个大条目查分存储在多个小条目中

Post Message Table		
Message_id	**Part**	**attr. 1**
100	0	first section
100	1	second section
100	2	third section

- 利用S3服务

Post Message Table	
Message_id	**S3 location**
100	s3://post_msgs/id_100

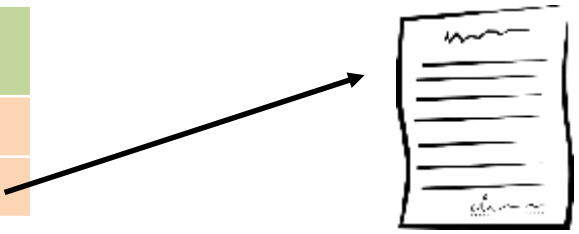

amazon web services

存储时序型数据

- 使用DynamoDB记录日志，点击行为，浏览行为，操作记录等数据

Ad View Table				
User id	**timestamp**	**attr. 1**	**…**	**attr. n**
100	2014-12-05 14:32:30			
100	2014-12-05 16:03:27			
100	2014-12-12 15:15:13			
101	2014-12-02 01:19:30			
101	2014-12-08 13:56:09			

DynamoDB 在AppFlood中的应用

- 所有生产数据都存在DynamoDB中
- 仅美西，有100亿条数据，索引存储空间超过2T
- 我们使用DynamoDB本地版本用于开发和单元测试
- 使用DynamoDB自动扩展工具，调整吞吐量，保证性能的基础上节省成本

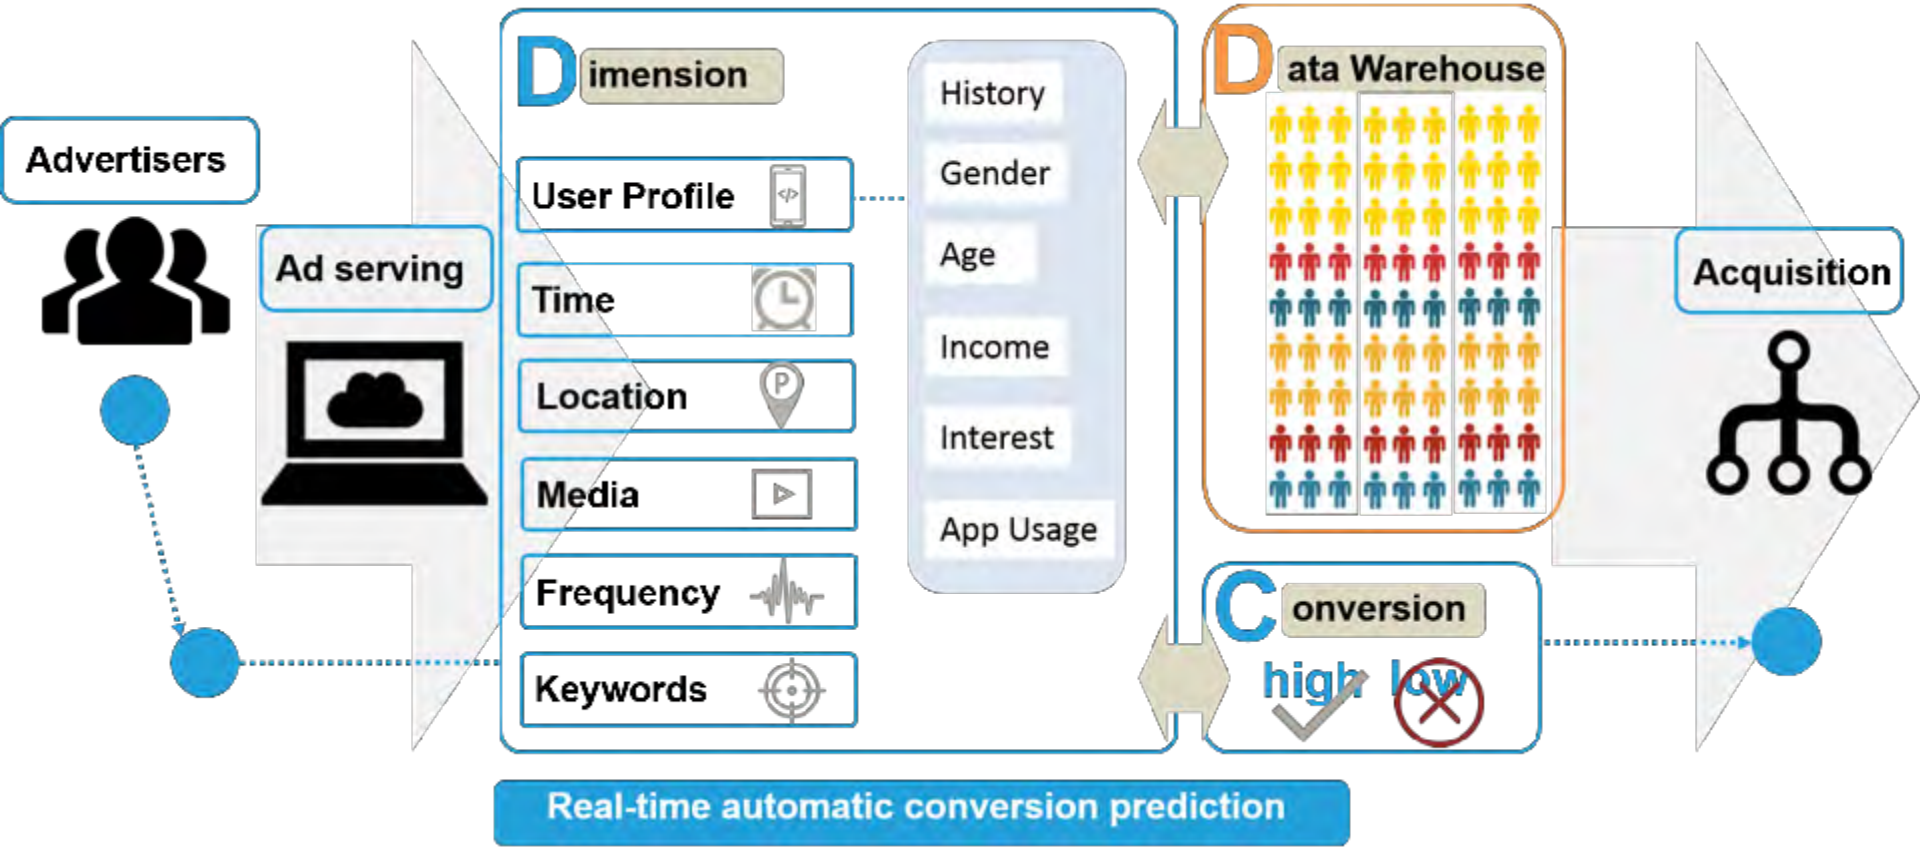

AppFlood 实时预测系统

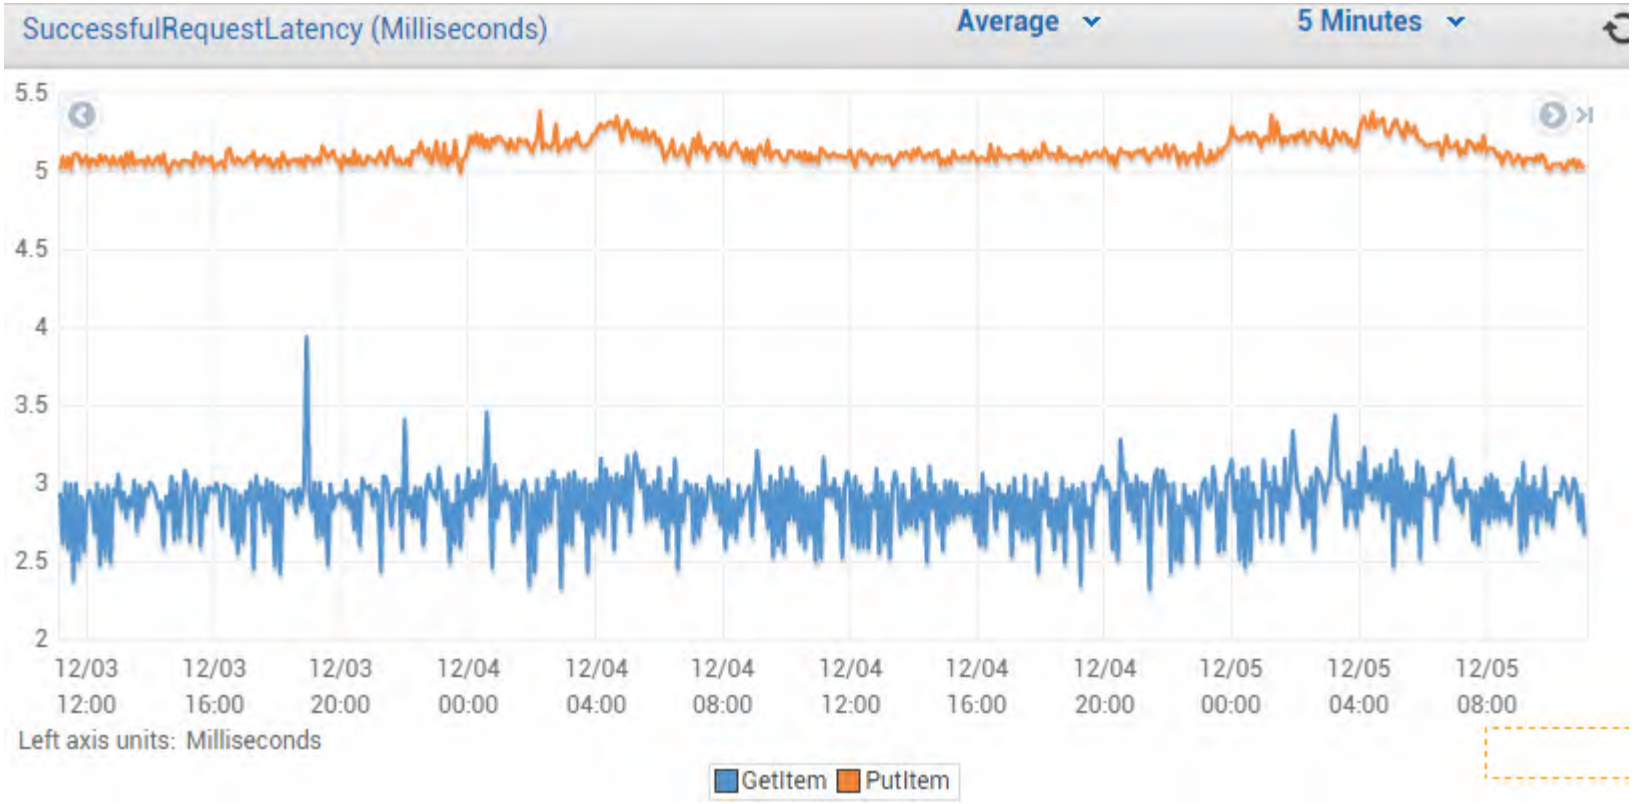

读、写延迟样例

THANKS

赵堃亮

zhaokunliang@papayamobile.com

Amazon Redshift

amazon web services

Amazon Redshift

Amazon Redshift:
一种快速，完全托管的PB级数据仓库服务

性能优越
价格低廉
更加简便

PB级数据仓库

大规模并行处理 (MPP)

关系型数据仓库 (SQL)

管理简便、大幅扩容

$1,000/TB/Year

从小到PB级数据仓库

Extra Large Node (DW1.XL)
2 TB HDD, 16 GB RAM, 双核

单节点 (2 TB)

集群 2-32 节点 (4 TB – 64 TB)

在线扩展集群以改善性能或增加容量

Eight Extra Large Node (DW1.8XL)
16 TB HDD, 128 GB RAM, 16 核, 10 GigE

集群 2-100节点 (32 TB – 1.6 PB)

amazon web services

Redshift的应用场景

✓	在线分析处理 (OLAP) 大量数据的复杂查询
✗	在线交易处理 (OLTP) 数据量小的简单查询 数据库如RDS和DynamoDB 是更好的选择
✗	非结构化数据 数据可能需要进行预处理加载到Redshift (可以使用EMR 运行 MapReduce)


与现有BI工具集成

Amazon Redshift

JDBC/ODBC

使用PostgreSQL.org的驱动程序实现连接

ACTUATE The BIRT Company
MicroStrategy
birst
tableau
JASPERSOFT the intelligence inside
infor
pentaho
Logi ANALYTICS
PERVASIVE
JReport JINFONET SOFTWARE
SiSense

及更多.....

Pinterest
airbnb

AWS数据库服务

Amazon RDS

- SQL 关系型数据库: MySQL、Oracle 、SQL Server和 PostgreSQL

Amazon ElastiCache

- 高性能的分布式内存对象缓存系统 Memcached & Redis

Amazon DynamoDB

- 大规模可扩展性，高耐久性，高吞吐量，低延迟 NoSQL 数据库

Amazon Redshift

- 性能优越的PB级别的数据仓库服务

The Next Big

AWS数据库服务演进

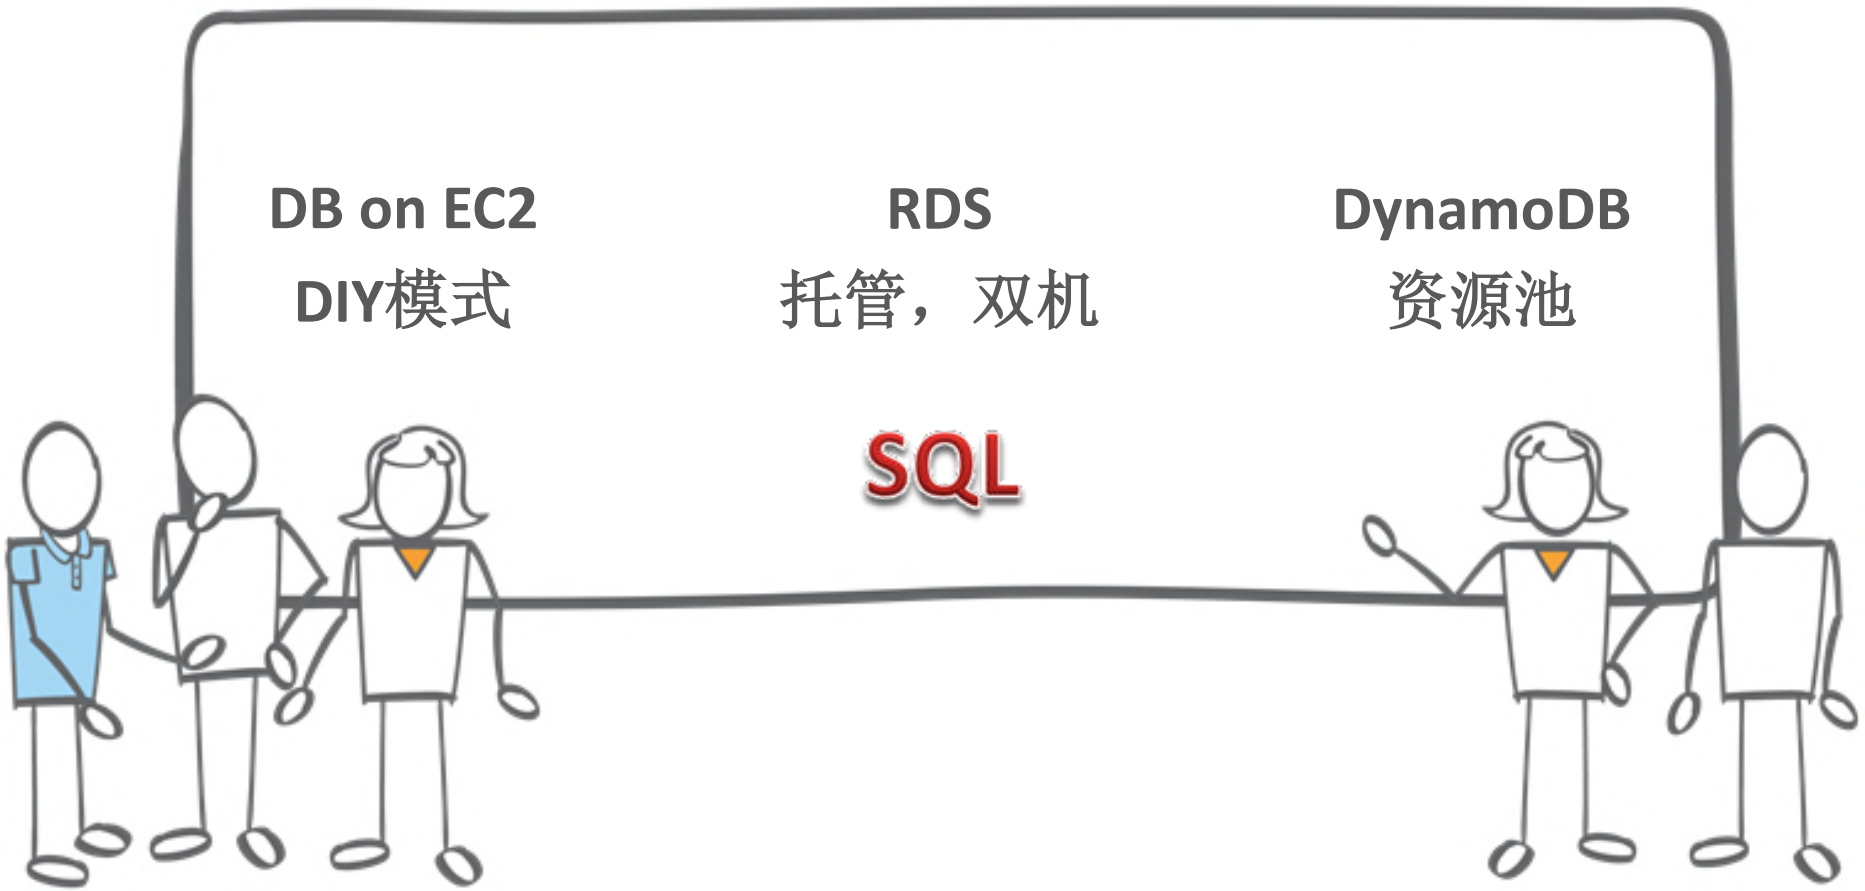

amazon web services

关系型数据发展多年，
客户的选择依然只有两种… …

Open Source 开源数据库

- 容易上手
- 成本可接受

但是… ….

- 扩展性差
- 可靠性有不确定性

Commercial 商业数据库

- 扩展性强
- 高级别的耐久性

但是……

- 使用复杂
- 昂贵
- 客户容易被锁定

客户应该有第三种选择！

Amazon Aurora

Amazon Aurora 综述

- RDS 的第五种关系型数据库引擎，全面兼容MySQL， 结合了商用数据库的高性能、高可用，以及开源数据库的简洁性与低成
- 只需几个点击，客户可以将MySQL数据库迁移至Aurora，并获得5倍的性能提升，成本却只是商用数据库的十分之一
- Aurora自动承担并解决了数据库中耗时耗力的工作：包括备份，监控，补丁升级，故障检测，备份以及灾难恢复
- 按需付费，无预付，无许可证费用，不被绑定

	vCPU	Mem	Included Storage	Included IOs	Hourly Price
db.large	2	15.25	40 GB	10 million	$0.29
db.xlarge	4	30.5	80 GB	20 million	$0.58
db.2xlarge	8	61	160 GB	40 million	$1.16
db.4xlarge	16	122	320 GB	80 million	$2.32
db.8xlarge	32	244	640 GB	160 million	$4.64

- Additional storage consumed, up to 64 TB, is billed at $0.10 a GB per month
- Additional IOs consumed is billed at $0.20 per million IO
- Prices are for Virginia

What's New

amazon web services

Thank You!
find AWS China online

中国官网
www.amazonaws.cn

官方博客
blog.csdn.net/awschina

官方微博
weibo.com/amazonaws

微信
AWS 中国

视频中心：
http://aws.amazon.bokecc.com/

amazon web services