

Emerging Technologies Shaping the Future of Data Warehouses & Business Intelligence

Appliances and DW Architecture

OBJECTIVES

- To define an “appliance”
- Understand critical components of a DW appliance
- Learn how DW appliances fit within your existing DW architecture

WHAT IS AN “APPLIANCE?”

General Purpose

Multi-purpose

Purpose specific



AN APPLIANCE IS BORN...

Compatibility	Applications (Operational, BI, DSS)	Purpose
	Database	Purpose Specific Appliance
	Operating System (Windows, Unix, Linux, etc)	
	Computing Resources (CPU, memory)	
	Networking Resources	
	Storage Resources (Persistence)	

Horizontal versus Vertical Layers
When does it flip?

PROVEN APPLIANCES



Google search appliance



Symantec VPN appliance



Linksys Wireless router



TiVo

CHARACTERISTICS OF A GOOD APPLIANCE

- Clear, specific purpose
- Compatible with existing infrastructure
- Ease of installation & use
- Highly cost effective

MAKING A DW SPECIFIC APPLIANCE?

- Optimize for mixed BI workloads
 - Store and access lots of data
 - Perform analytical queries
 - Support operational report requirements
 - Expand data access to the mainstream DW user
- Combine Server + Optimization Software + Storage
- Ensure ease of use
- Offer low maintenance and TCO
- Integrate with existing infrastructure
 - BI tools with ANSI SQL
 - Ability to load data easily
 - Compatible with standard operating procedures
 - Build in flexibility to support changing requirements

PRE-CONFIGURED vs. APPLIANCE

- Appliances are not a “pre-configuration” or assembly of tuned products and technologies
 - An appliance is built for a specific purpose and can leverage commodity parts
 - Appliances require no user serviceable parts
-
- Test: Does it matter what’s inside the appliance?

DW APPLIANCE: HIGHLY COST EFFECTIVE

- Reduce DBA expertise and administration
 - No tablespaces, extents to manage
 - No archive, redo logs to manage
 - No partitioning management
- Reduce need for storage expertise
 - SAN storage architects
 - LUNs, meta volumes
- Relatively low acquisition cost vs. traditional model
 - Server manufacturer annual maintenance
 - Operating System license
 - Storage hardware
- Efficient power consumption

FUTURE DATA WAREHOUSE APPLIANCES?

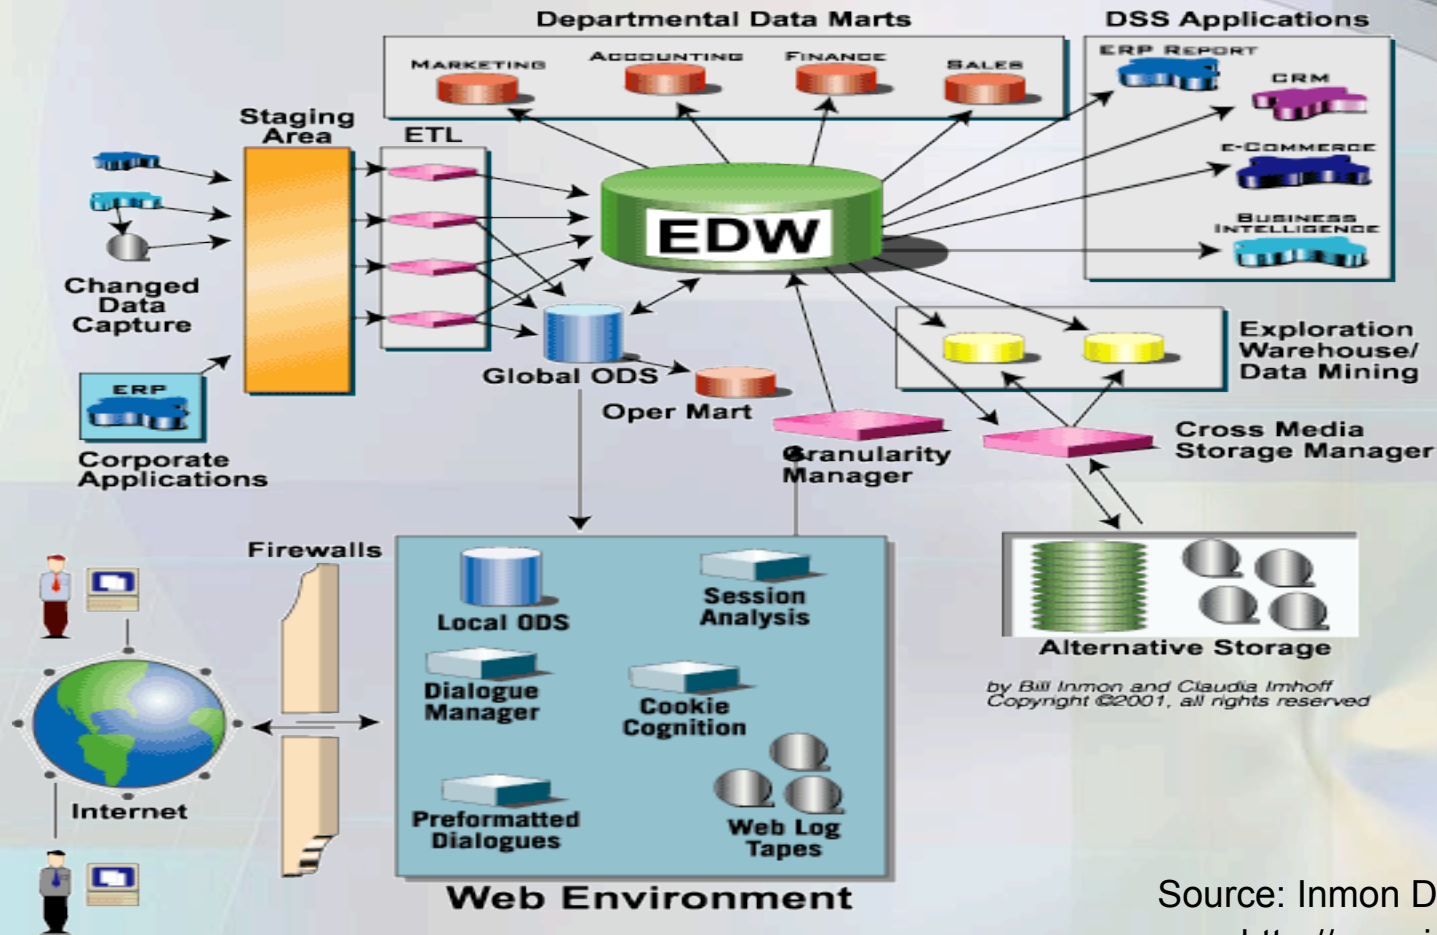
Will we see these specific appliances in the future?

Do they meet the criteria for success?

- ETL or ELT appliance ?
- BI appliance ?
- EII appliance ?
- OLAP appliance ?
- Audit appliance ?

WHERE DO DW APPLIANCES MAKE SENSE?

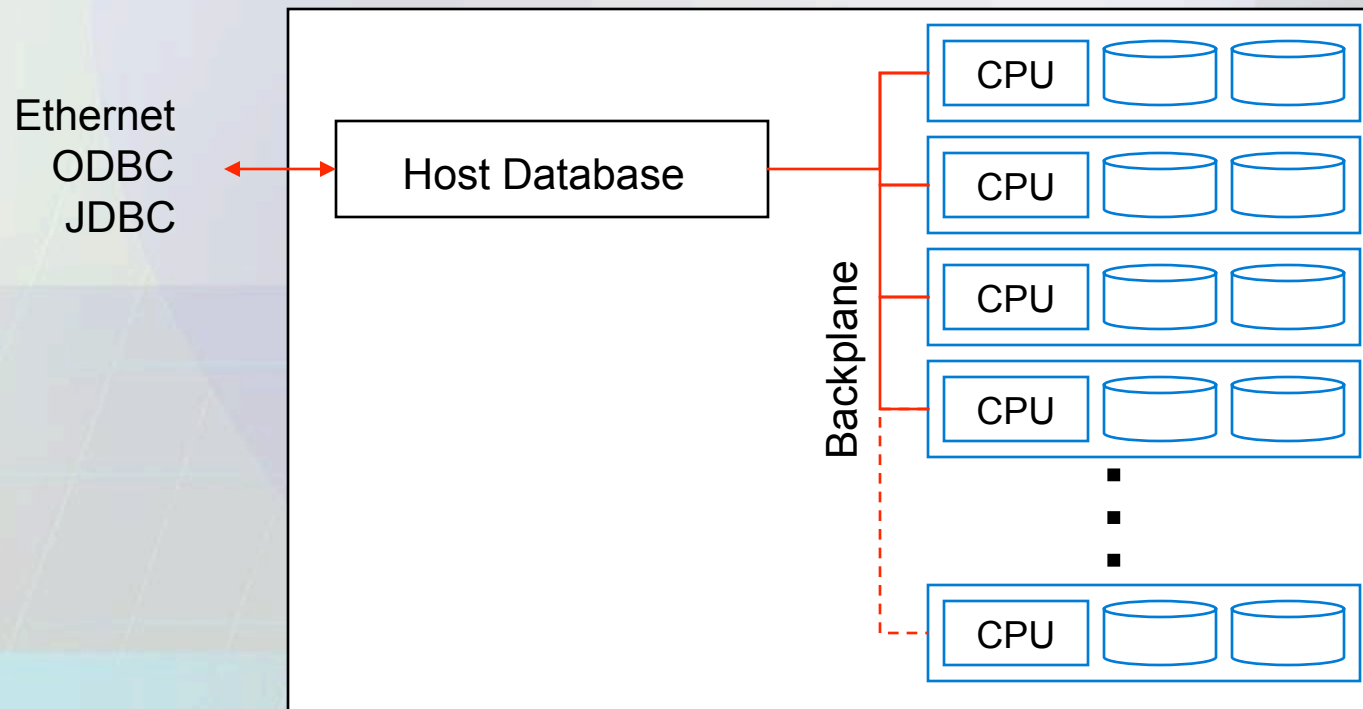
The Corporate Information Factory and the Web Environment



Source: Inmon Data Systems
<http://www.inmoncif.com>

DW APPLIANCE: HIGH PERFORMANCE

- Massive Parallelism inherent to the architecture
- Improved I/O throughput rates with effective I/O



DW APPLIANCE: LARGE SCALE DATABASES

- How do you effortlessly scale beyond 2TB as your data warehouse matures?
- Federated architecture
- Rapid expansion and data redistribution

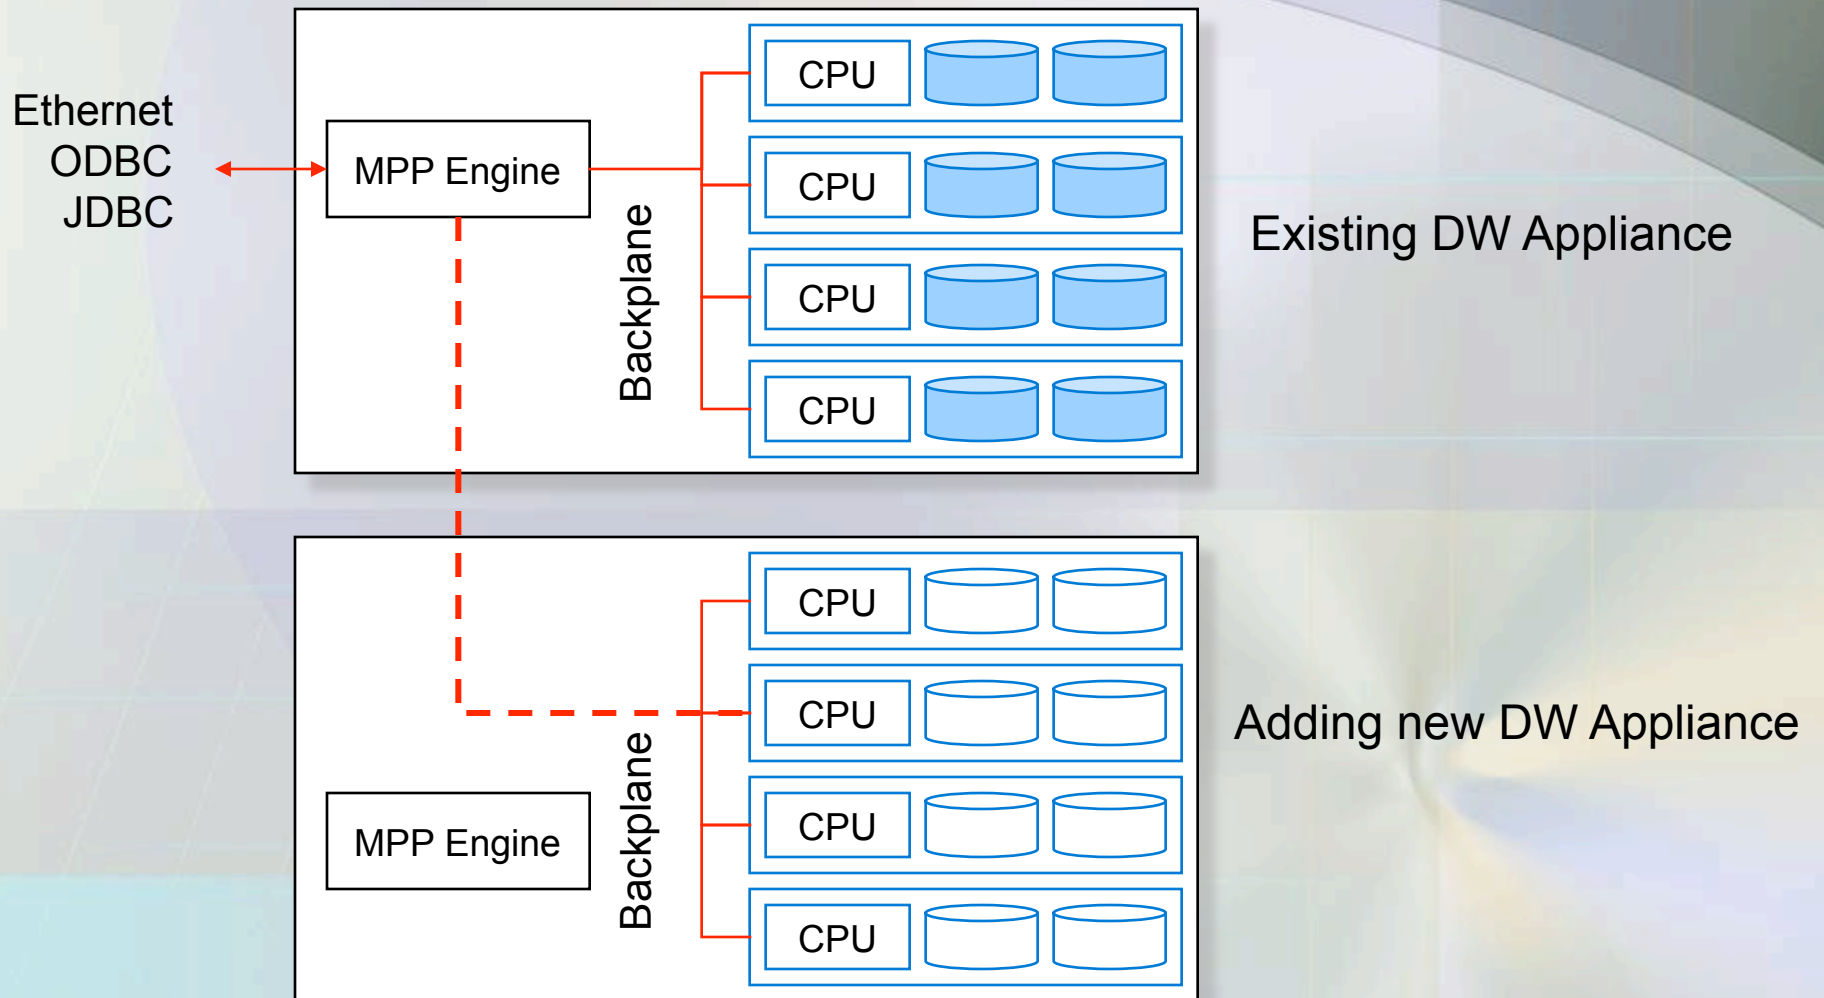
SCALABILITY: FEDERATED ARCHITECTURE

- Logical areas of the data warehouse architecture move to new appliances
 - Data marts, ODS, staging
- Subject areas or conformed dimensions stay together unless the capability to perform database joins across platforms exist
 - EII tool, remote tables
- Logical groups such as North America DW, Euro DW or corporate entities

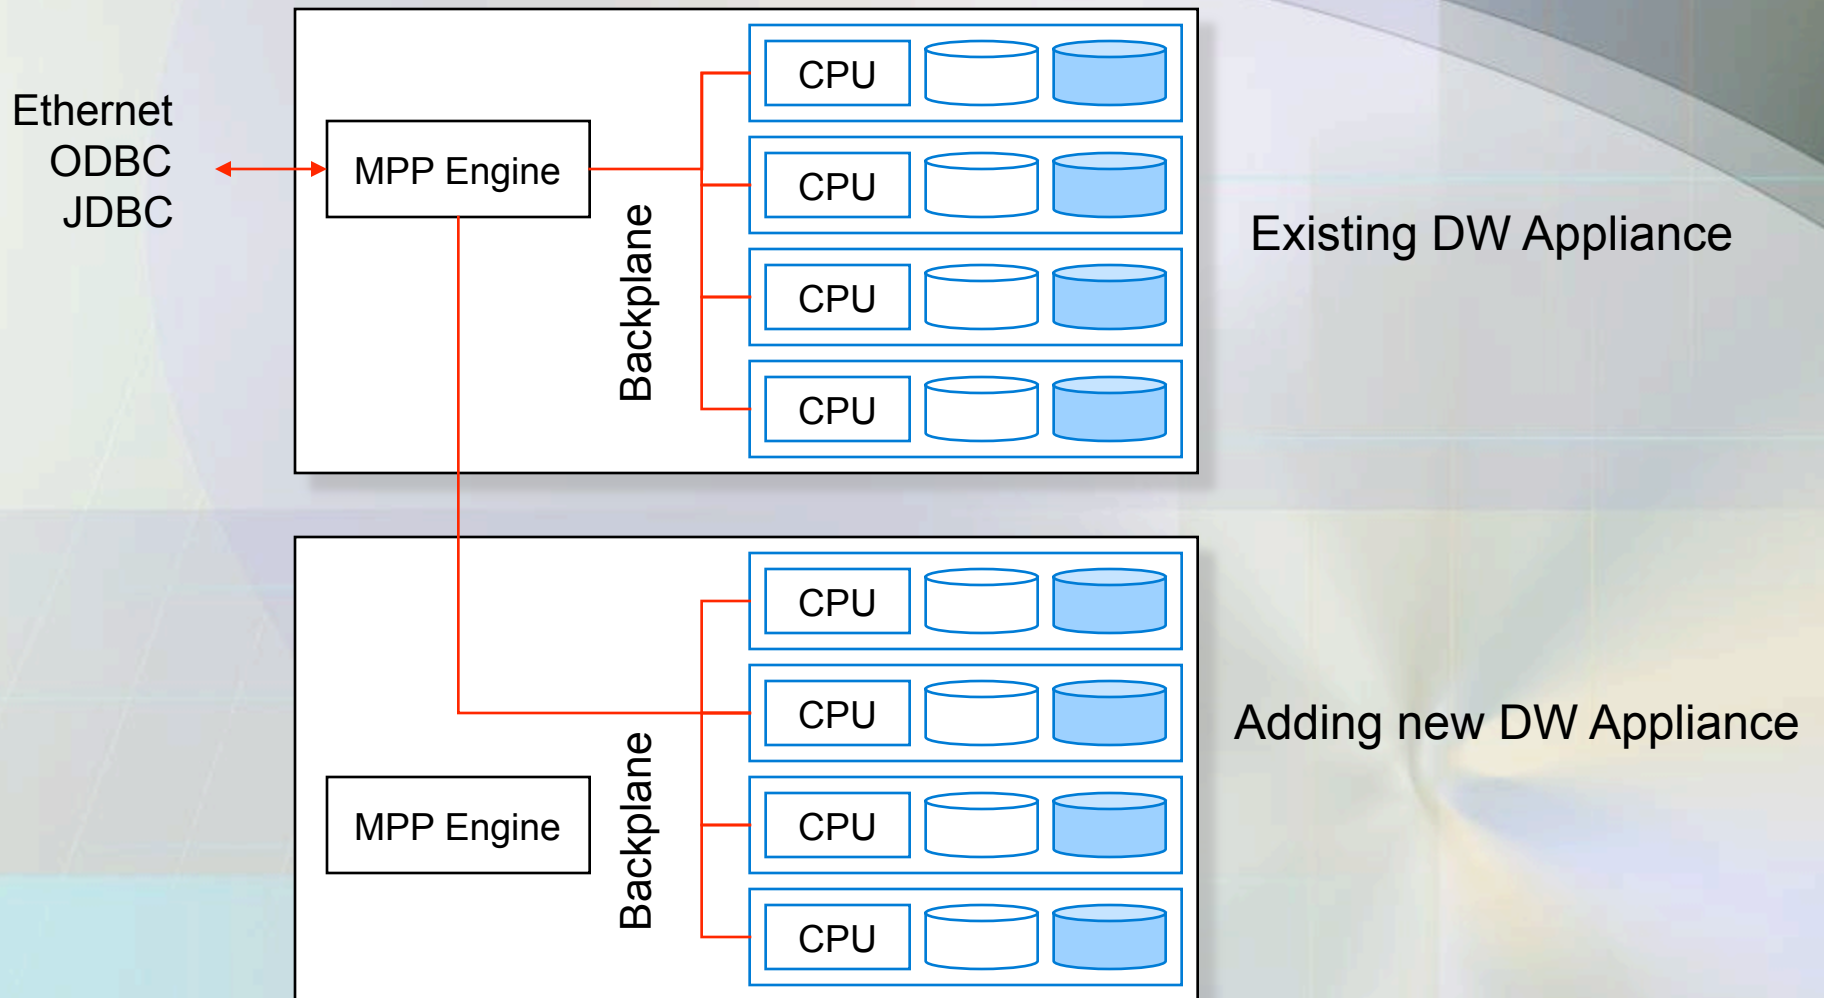
SCALABILITY: DATA REDISTRIBUTION

- Adding capacity may cause data to be redistributed
 - Buckets
 - SLA Driven
 - Redistribution

SCALABILITY: DATA REDISTRIBUTION



MAINTAINING PERFORMANCE WITH SCALABILITY



DW APPLIANCE: BACKUP AND RECOVERY

- Built-in RAID mirroring and striping
- The challenge is that large systems take massive amount of computing and network resources to backup changes in a high volume environment
- Recommended to compress and save load files since loads are faster than recoveries
- In a fully mirrored environment, chances of needing to restore from a backup are low
- Veritas or other standard APIs are available backups

DW APPLIANCE: FAST LOAD/UNLOAD

- Minimum 500GB per hour load rates
- Near physical limitations disk I/O and Ethernet connection to appliance
- Utilizes high performance ODBC drivers and special loader utilities

IMPACT OF DATA WAREHOUSE APPLIANCES

WHAT A FAST DATABASE CAN DO FOR YOU

- Before ETL tools
 - ETL was hand coded programs
 - ETL was code in database procedural languages
- ETL tools offered
 - Faster development with 3GL
 - Better performance than database code & SQL
- High Performance data warehouse appliance
 - Faster queries on large data sets
 - → High Performance SQL

ELT – A PARADIGM SHIFT

- Transformations typically are performed after extraction and before loading limiting the database workload which is thought to be query intense for data warehouses.
- High performance data warehouse appliances are being used to transform the data inside the high performance database and the results stored in the appliance or in other databases
 - Some companies point to ETL license savings as part of the ROI for appliances
 - Sunopsis is a tool being used in these cases

ROLAP vs. MOLAP

- Multidimensional database engines were created to make up for the performance deficiencies of relational OLTP databases at that time.
- MOLAP cubes are known for their high performance interactive capabilities, complex calculations and dimensional and hierarchical analysis but are challenged in:
 - Real-time environments
 - Large datasets
 - High degree of dimensional and hierarchical branches
- MOLAP cubes are also a duplication a data in another database format, the multidimensional database.
- ROLAP leverages on the underlying database for performance characteristics
- ROLAP metrics can be atomic level, pre-calculated, pre-rolled up depending on what works best

FOCUS ON BUSINESS VALUE OF DATA

- Data Architecture and Modeling
 - High performance databases shouldn't make up for poor data analysis and modeling efforts
- Good Requirements, Analysis and Reports
 - Make sure that the wrong answer doesn't just come back faster...

WHAT THE BUSINESS WANTS FROM DW

- On-Demand
 - New products and projects
- Lower Cost
 - Overall TCO and long term maintainability
- Flexibility
 - Simpler infrastructure to build and go
 - Able to update as the business needs change
- New capabilities
 - Scalability
 - Ability to store more data with less cost

Questions and Discussion