BMW11:

Dealing with the Massive Data Generated by Many-Core Systems

Dr Don Grice

Title: Dealing with the Massive Data Generated by Many Core Systems.

Abstract:
Multi-core and Many-core architectures are enabling computing systems that are more powerful than ever. The amount of data being generated by these systems is becoming an issue in several areas, including storage of results, movement of intermediate and final results, and the ability to consume the data and transform it into 'information'. As we move forward we need to be developing HW and SW methods to deal with this massive data explosion. Data reduction/simplification and real time analytics will involve more computation but may be one of the most promising methods for dealing with this flood of newly generated data.

TYPES OF DATA MANIPULATION

COMPUTE INTENSIVE
DATA INTENSIVE
NETWORK INTENSIVE

Styles of Massively Parallel Workloads

Data Intensive (Data At Rest)

Hadoop/MapReduce (BigInsights)

Data at Rest*:
- **High Volume**
- **Mixed Variety**
- **Low Velocity**

= compute node

(*pre-partitioned)

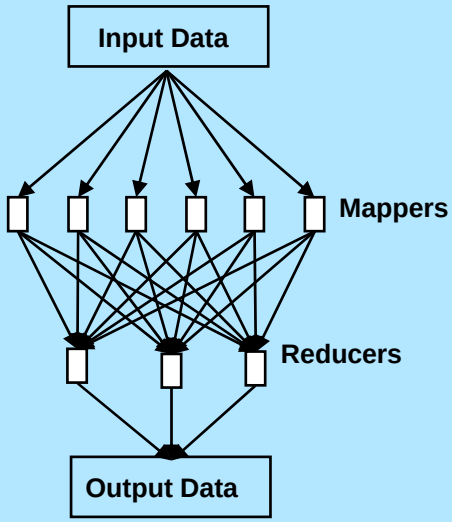

Data Intensive (Data Needs to Move)

Global Analytics:

View of All Data Required

Data 'Must be Moved'
Higher Velocity
Network is Critical

Data Intensive : Data in Motion (Streaming)

SPL, C, Java

Data in Motion:
- **High Velocity**
- **Mixed Variety**
- **High Volume***

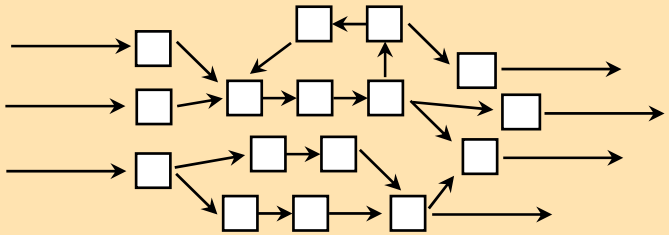

Reactive Analytics
Extreme Ingestion

(*over time)

Compute Intensive (Data Generators)

C/C++, Fortran, MPI, OpenMP

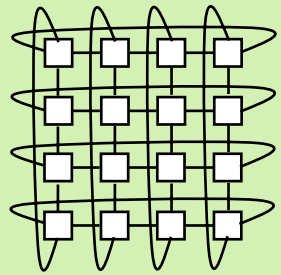

- **Long Running**
- **Small Input**
- **Massive Output**

Generative Modeling
Extreme Physics

Styles of Massively Parallel Workloads

IBM

Embarassingly Parallel

Data Intensive (Data At Rest)

Hadoop/MapReduce (BigInsights)

Data at Rest*:
- **High Volume**
- **Mixed Variety**
- **Low Velocity**

= compute node

(*pre-partitioned)

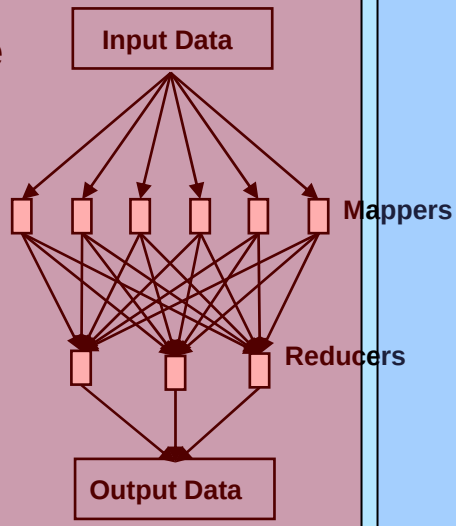

Network Dependent

Data Intensive (Data Needs to Move)

Global Analytics:

View of All Data Required

Data 'Must be Moved'
Higher Velocity
Network is Critical

Data Intensive : Data in Motion (Streaming)

SPL, C, Java

Data in Motion:
- **High Velocity**
- **Mixed Variety**
- **High Volume***

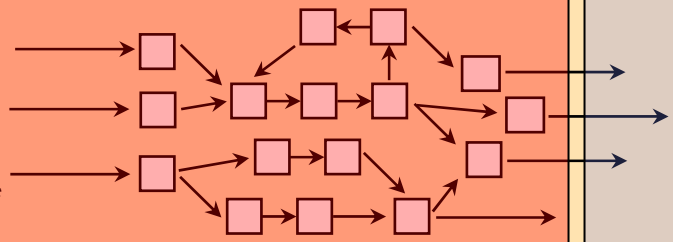

Reactive Analytics
Extreme Ingestion

(*over time)

Compute Intensive (Data Generators)

C/C++, Fortran, MPI, OpenMP

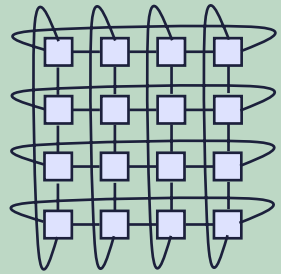

- **Long Running**
- **Small Input**
- **Massive Output**

Generative Modeling
Extreme Physics

Data Intensive Applications
(Large Data)

New “Big Data” Brings New Opportunities, Requires New Analytics

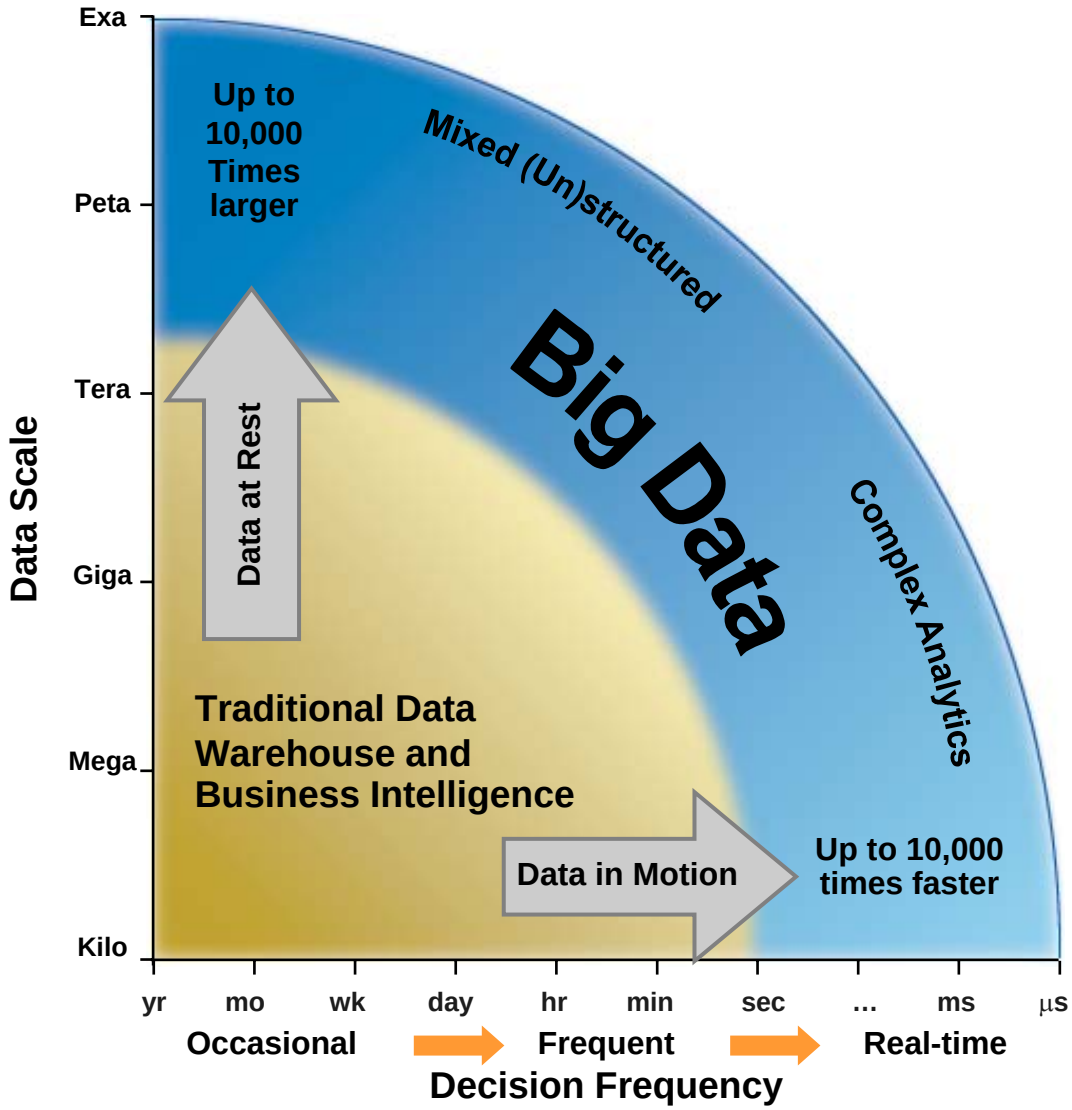

Telco Promotions

100,000 records/sec, 6B/day

10 ms/decision

270TB for Deep Analytics

DeepQA

100s GB for Deep Analytics

3 sec/decision

Smart Traffic

250K GPS probes/sec

630K segments/sec

2 ms/decision, 4K vehicles

Petascale Analytics, Appliances and Ecosystem

Big Data is the new resource. **The new opportunity is Big Analytics.** Every Smarter Planet solution will depend on it.

Market leadership in the Era of Analytics will be taken by **the first player to deliver high volumes** of easy-to-use Smarter Planet solutions.

Ultimate success will require a Petascale Analytics Appliance and a rich ecosystem of data, algorithms and skills.

Maximum Insight Requires Combining Deep and Reactive Analytics

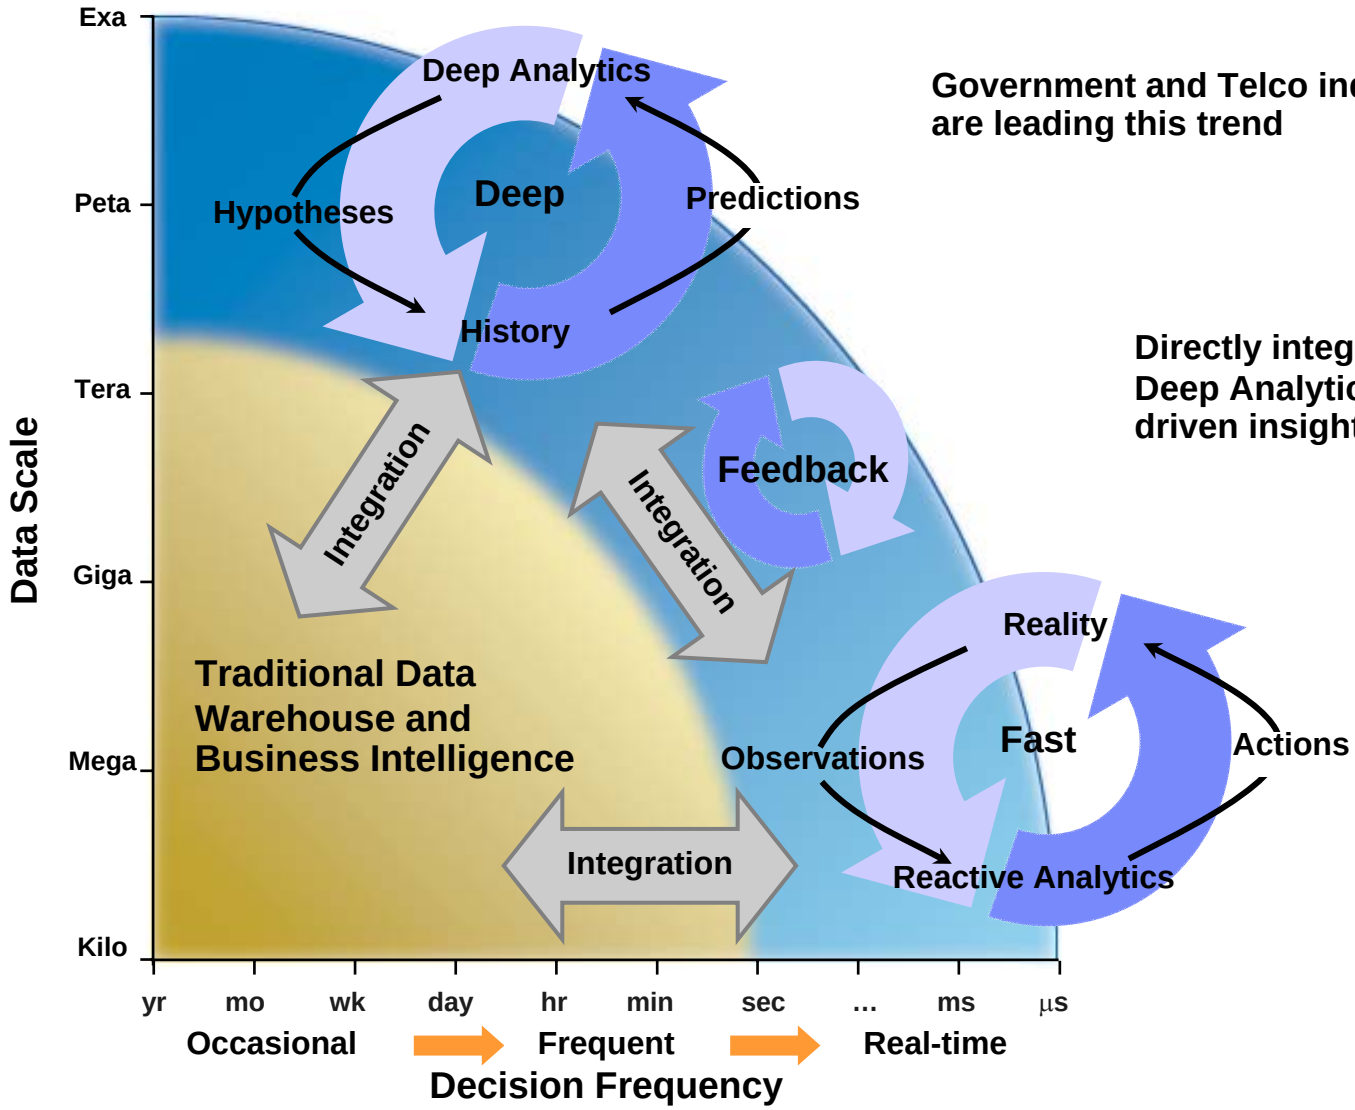

Government and Telco industries are leading this trend

Directly integrating Reactive and Deep Analytics enables feedback-driven insight optimization

Watson

- IBM Research built a computer system that is able to compete with humans at the game of Jeopardy: Human vs. Machine contest.
- Named "Watson," the computer is designed to rival the human mind
- Answering questions in natural language poses a grand challenge in computer science, and the Jeopardy! clue format is a great way to showcase:
  - Broad range of topics, such as history, literature, politics, popular culture and science
  - Nature of the clues, requires analyzing subtle meaning, irony, riddles and other complexities
- Based on the science of Question Answering (QA); differs from conventional search
- Natural Language / Human Interactions
- Critical for implementing useful business applications such as:
  - Medical diagnosis
  - Regulatory compliance
  - Customer relationship management
  - Help desk support

IBM

Compute Intensive Workloads (Traditional ‘HPC’)

Fundamental Issues with Large Scale HPC Compute Intensive Workloads

- Power Efficiency
  - TCO
- Programmability and Scaleout
  - Frequency is Plateaued
  - More Parallelism is needed
  - Balanced BWs are required for ‘sustained’ Perf
  - Shared Memory Model vs I/O ‘Accelerator’ Model
- Availability and Reliability
  - More Circuitry is required
  - Technology Scale makes it worse
  - Design for Availability is required
- Data Management and Cost of Storing/Moving Data
  - Time Steps & Checkpoints
  - Storage Cost, Energy Cost, BW, Latency
  - Life Cycle Management

Amount of Data Generated Growing Much Faster Than BW to Store or Retrieve it

- Example: 100x improvement in Machine Performance

- Core Frequency has Plateaued
- 100x Performance -> >100x more cores
- Memory per core ~ constant? -> >100x more memory

- Checkpoint Data Increase >100x
  Plus frequency may increase due to reliability changes
- Time Step Data Increase at least 100x? (with Performance)

- Disk and Tape BWs are basically Plateaued (~100MB/s)
  - Compression Methods are not improving much
    Only provides ~2x BW boost at most in any event
  - Capacity Growing at 20-30% CGR but not BW
- Amount of Disk/Tape needs to grow >100x to match BW
- Some relief possible with Write Duty Cycle Utilization
  - Cache locally and take full interval to write it out
  - Pre-stage Reads

Example of Data Volume Gap

- Example of Data Volume Gap Growing for Commercial Users
- BW Gap is even larger!

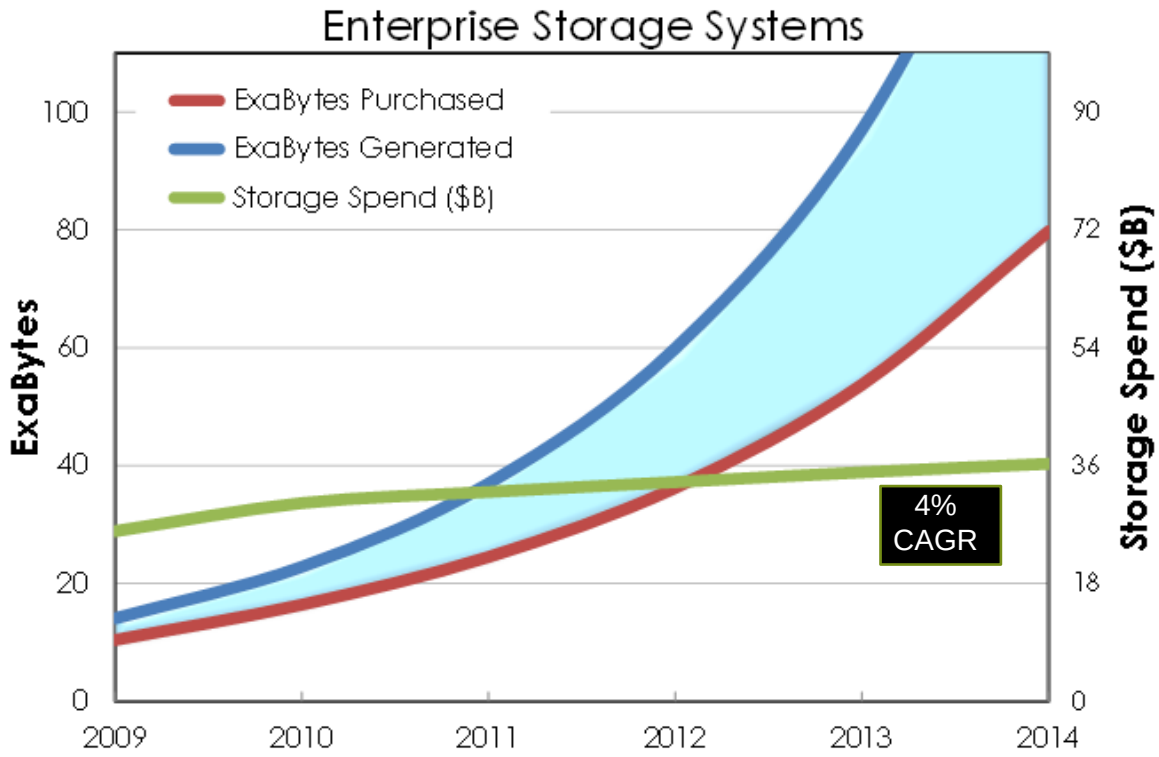

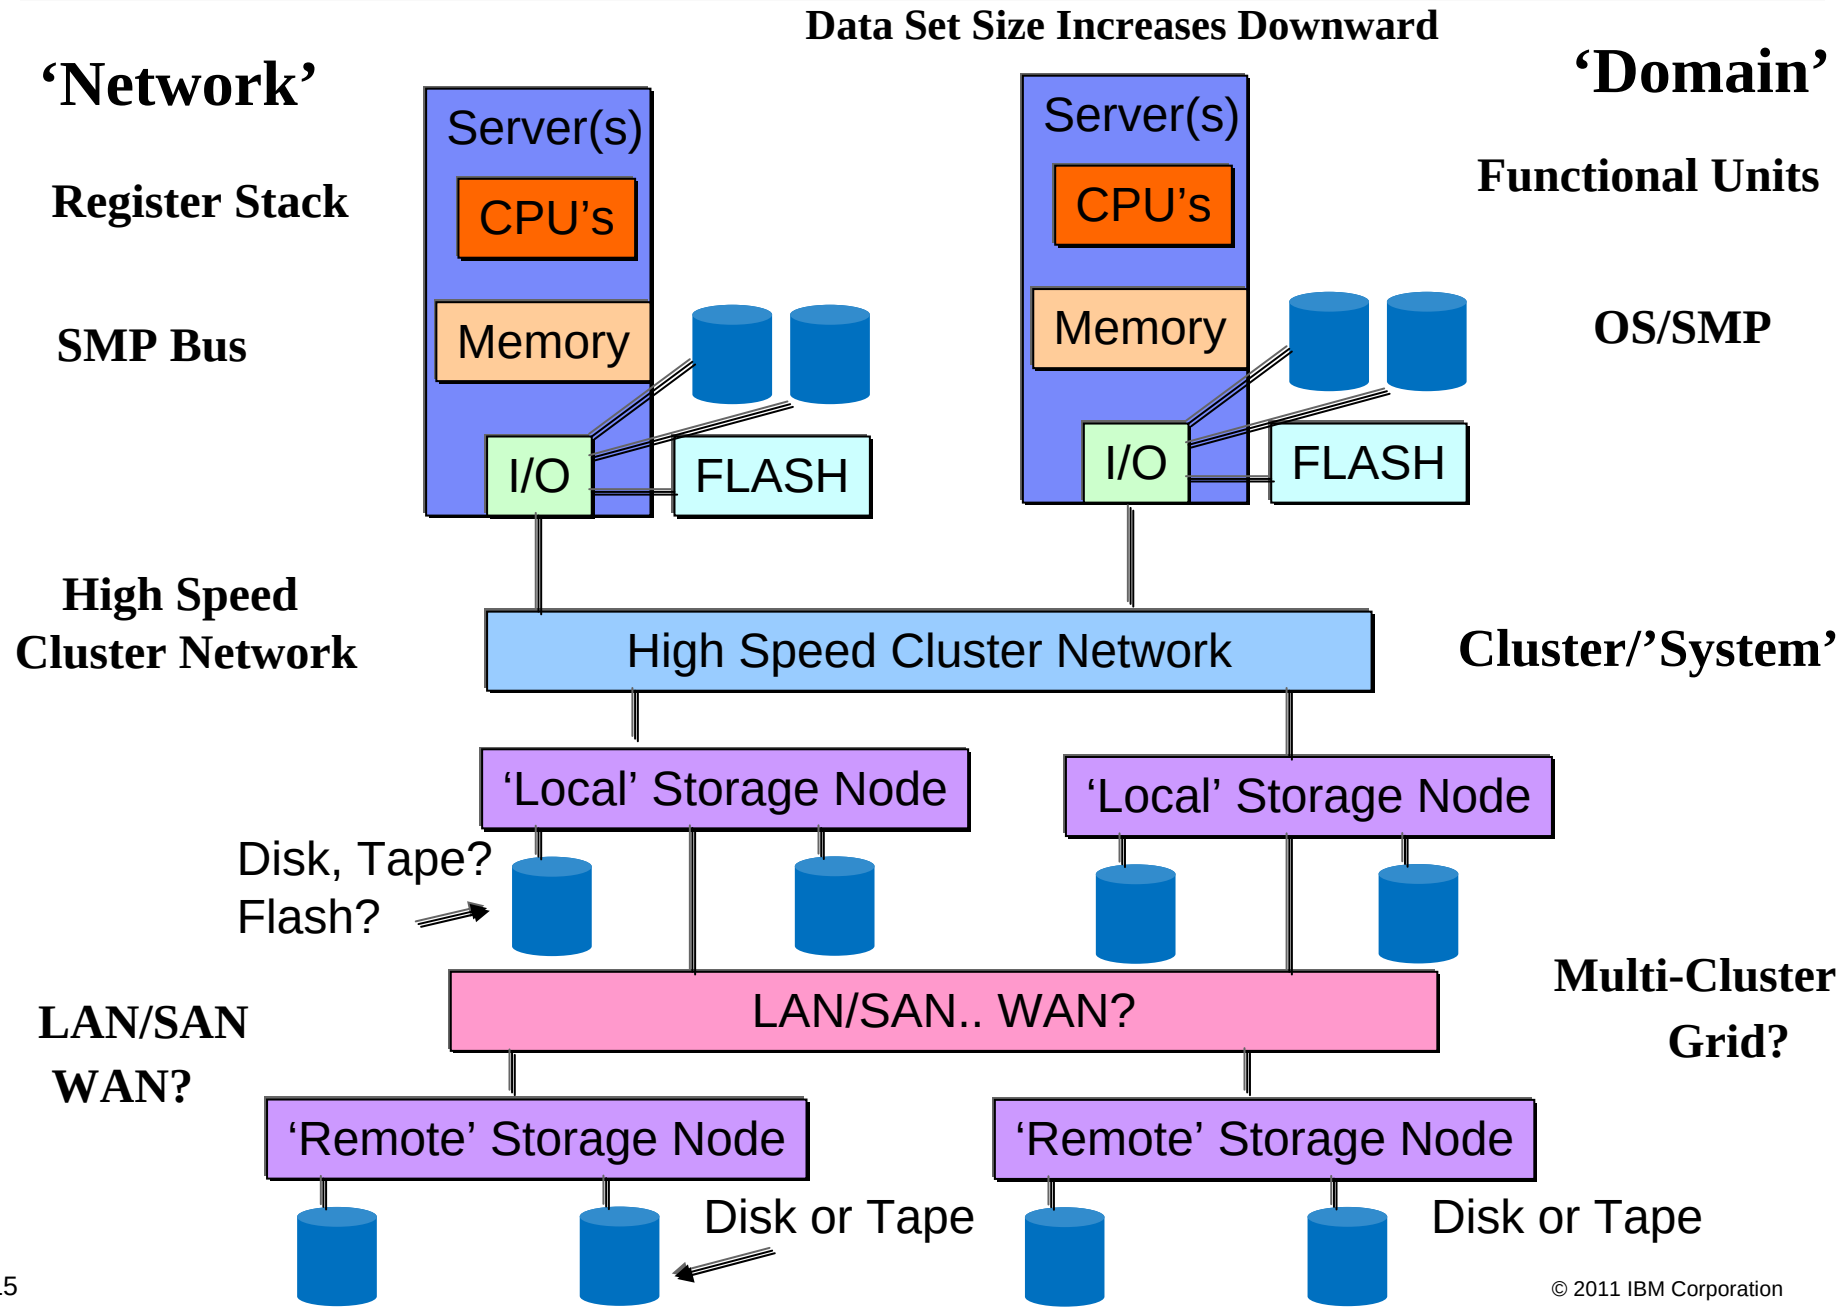

SUMMARY

- Data Volume and BW is Exploding in many Areas
- Multicore/Many Core Compute Intensive Systems
  - Are generating more data and faster than ever before
  - Also using more Memory due to Frequency Stabilization
- Data Storage BWs are not improving much
- Balance of Compute to I/O and Storage will need to shift
- Compute Intensive Workloads will also interact with Data Intensive Workloads in Workflow environments
- Data Life Cycle Management, Prestaging and Intelligent Writing will become increasingly more important as machines grow in capability

Thank you...

...any Questions?