

## Storage Reinvented: Smarter managed storage with XIV

Hugh Hulleman

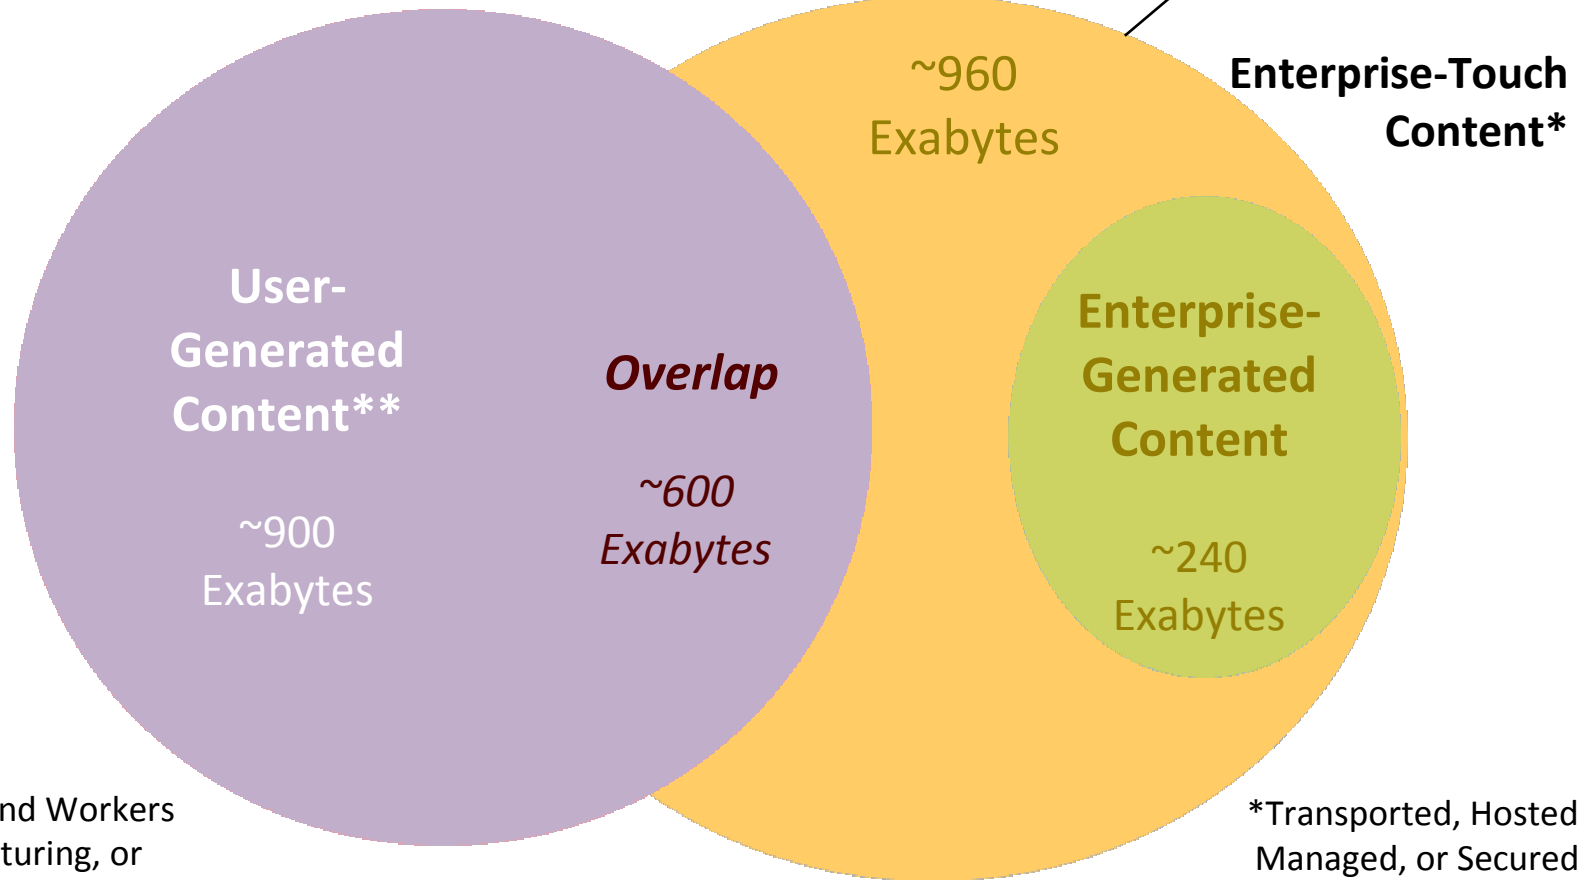
Senior Storage Architect - IBM XIV Storage Systems



# Structured and un-structured data, the growth



- Structured data: databases, stable growth
- Unstructured data: pictures, documents, multimedia, etc, growth is exploding



\*\*Consumers and Workers Creating, Capturing, or Replicating Personal Information

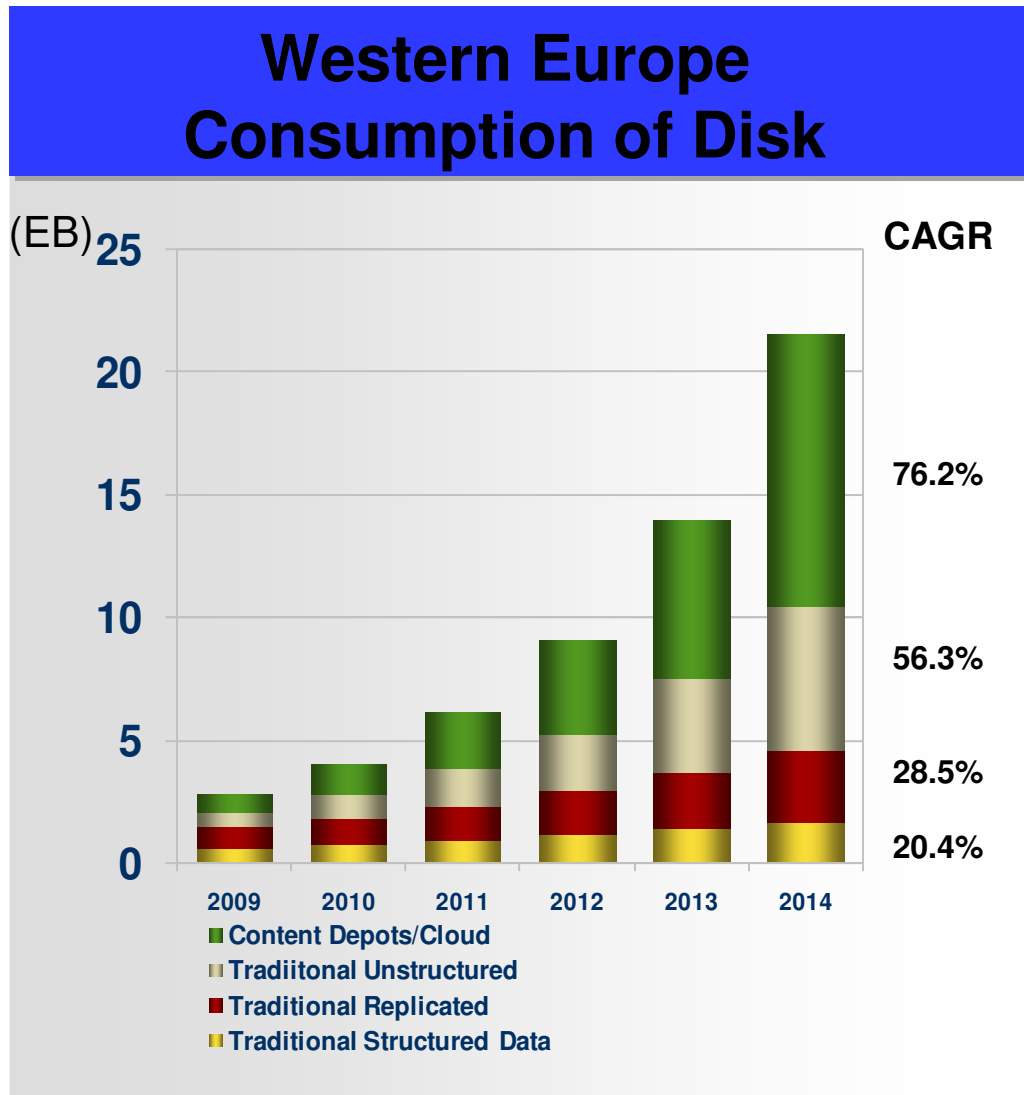
\*Transported, Hosted, Managed, or Secured

**2010**  
**1,200 Exabytes**

Source: IDC 2010



Gain actionable insights from data



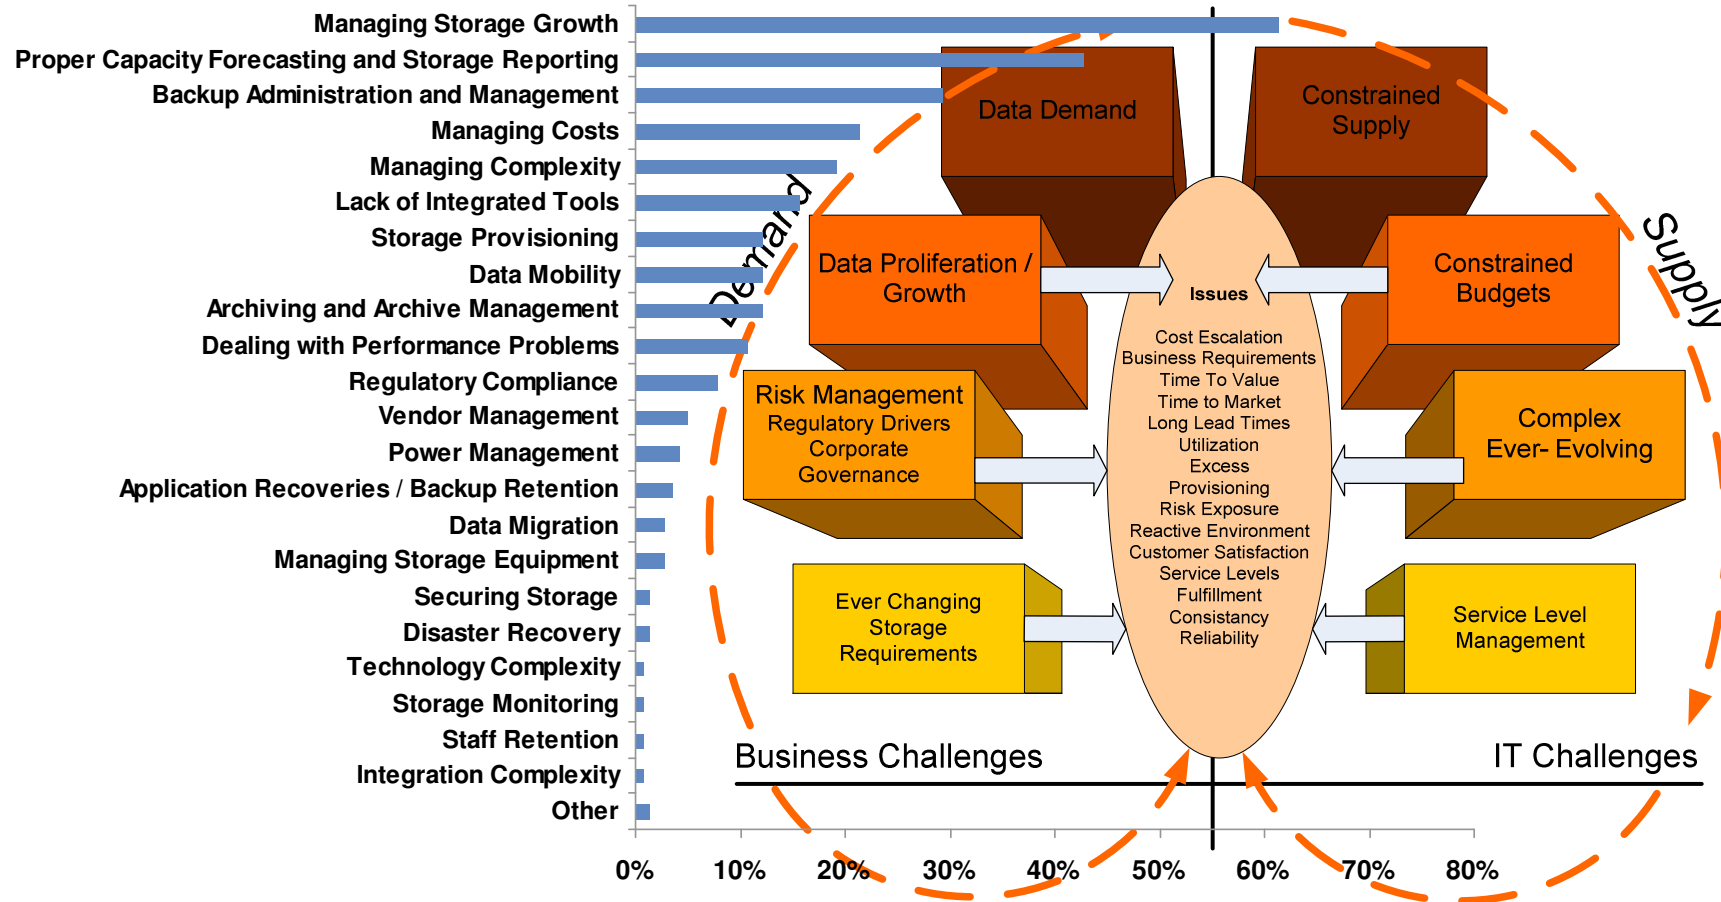
- Structured, transaction-oriented data accounts for declining portion of storage needs
- Compliance and analysis driving replication
  - Disk-based recovery
  - Business analytics
- Explosion in fixed content data
- Content Depots & cloud change the rules for data
  - Petabyte scale storage

Source: IDC 2010



Gain actionable insights from data

# The storage management problem



Gain actionable insights from data

# Escalating complexity of traditional storage



- Disk Tiers
  - **1-3 Types and sizes of disk**
  - **Flash/SSD?**
  - **Replica pool**
  - **Spares**
  - **# RAID groups / Different RAID Levels**
- Data center efficiency
  - **Power**
  - **Space**
  - **Utilization**
- Responsiveness / cost to business units

This results in THOUSANDS of components to manage in a traditional array!

Complexity =  
COST =  
*Decelerated ROI*

***With XIV, manage storage capacity - NOT technology!***



# What is XIV storage?



- IBM Enterprise Grid Storage Solution: 27-161TB **USABLE** capacity
  - **30-50% Higher effective utilization related to traditional storage**
- Cost effective | Unmatched TCO
  - **Built from industry standard hardware components**
- Ground-breaking innovations in ease of use, flexibility and manageability
- Innovative, all inclusive design
  - **No additional software license fees**
  - **Functionally integrated software design**

## AVAILABILITY

5 9s+ (99,999%)

Multiple levels of redundancy

Autonomous self-healing

Ultra-fast fault recovery

## PERFORMANCE

Self-tuning

Automatic load balancing  
eliminates bottlenecks

Performance scales with  
capacity

## EASE OF USE

Provision new storage in <1  
minute

Dynamically resize/adapt

Easily managed by junior  
staff

**Drive out  
cost while  
exceeding  
existing  
SLAs.**




Gain actionable insights  
from data

# Improved capacity utilization



- XIV sold as USABLE capacity
  - NO lost capacity due to : spares, special system areas, volume set asides for replication, etc.
- Capacity usage easy to monitor
  - Complete system, storage pool, or volume
- XIV all virtual
  - Single disk type and no RAID groups minimizes islands of capacity
  - No physical disk binding
  - THIN provisioning standard
- Designed to perform well at >90% capacity utilization




Gain actionable insights  
from data

# Ongoing administration/management



- Tier-less, hybrid design eliminates complexity
- Self healing & load balancing design
- Major reduction in tuning
- No reconfiguration charges
  - All logical config changes by your staff real-time
- New management paradigm:
  - Capacity vs. technology (disks & raid groups)
  - Train admins < 2 hours
  - Simple SNAPs can be delegated to DBA's



***Industry analysts estimate administration is 70% of TCO!***



Gain actionable insights  
from data



# Agility, responsiveness, flexibility



- Provision storage in < 1 minute
- Resize volumes dynamically
- High performance SNAPshots
  - Replicate data for backups or other apps
  - Empower DBAs to manage SNAPs
- Multi-tenant management
  - Delegate responsibility
- Add new capacity quickly, transparently
  - Added capacity available near real-time
  - Existing data automatically re-balanced across all storage resources

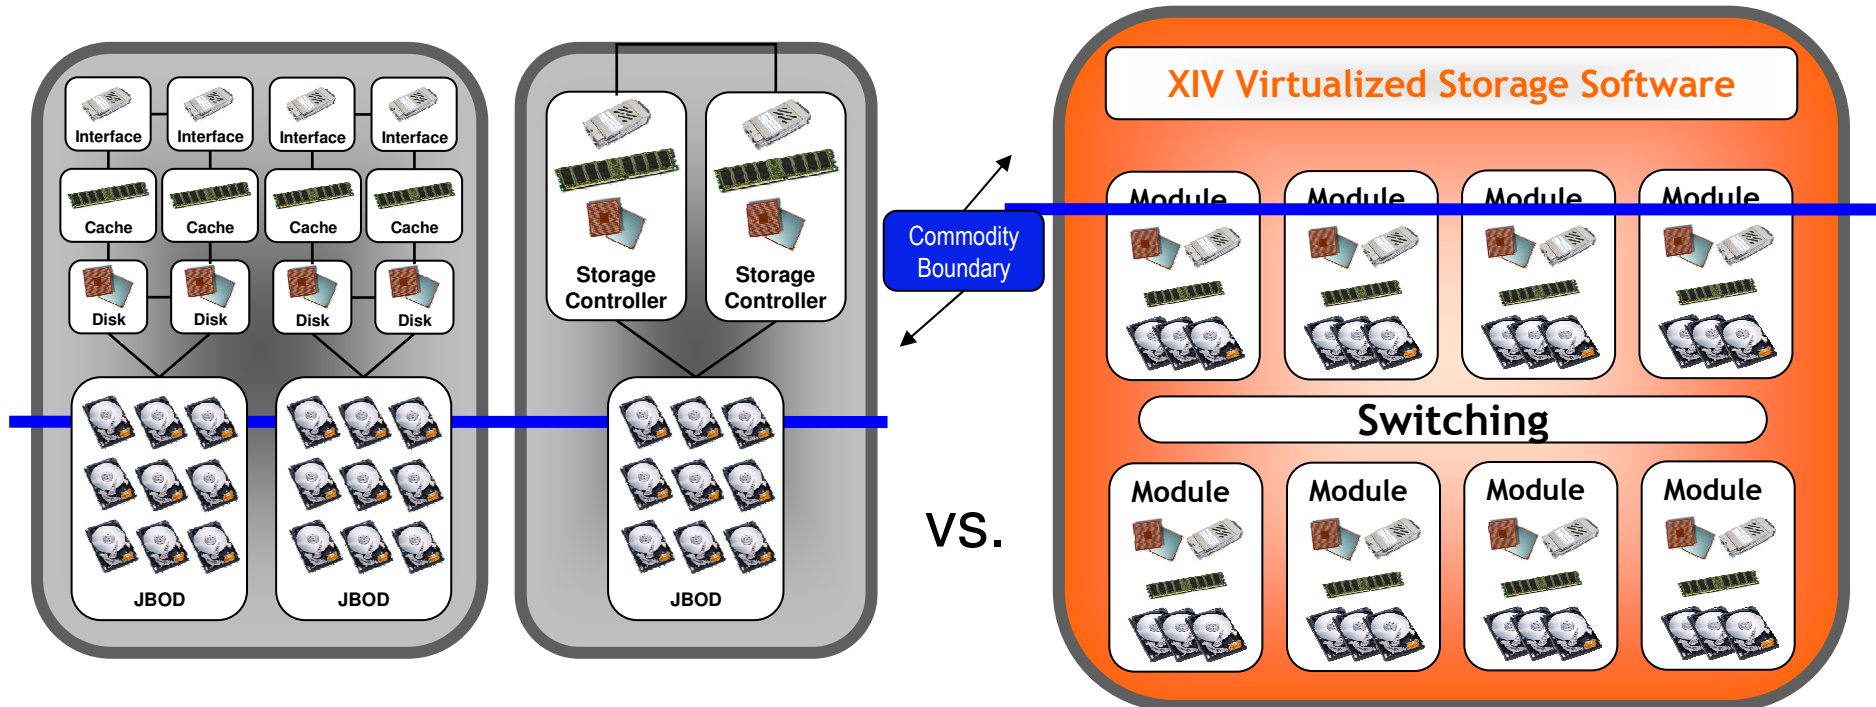


Gain actionable insights  
from data

- Comparable systems (equivalent usable space, performance):
  - 50% greater usable capacity per sq. meter
  - 25-30% less power & cooling per usable TB
- Commodity hardware components
  - Low power SATA II disk drives
  - Power-efficient Intel-based server modules
- Very efficient capacity utilization
  - Single level storage = no orphaned space
  - Virtual design + thin provisioned = efficient



# The XIV solution – core technology



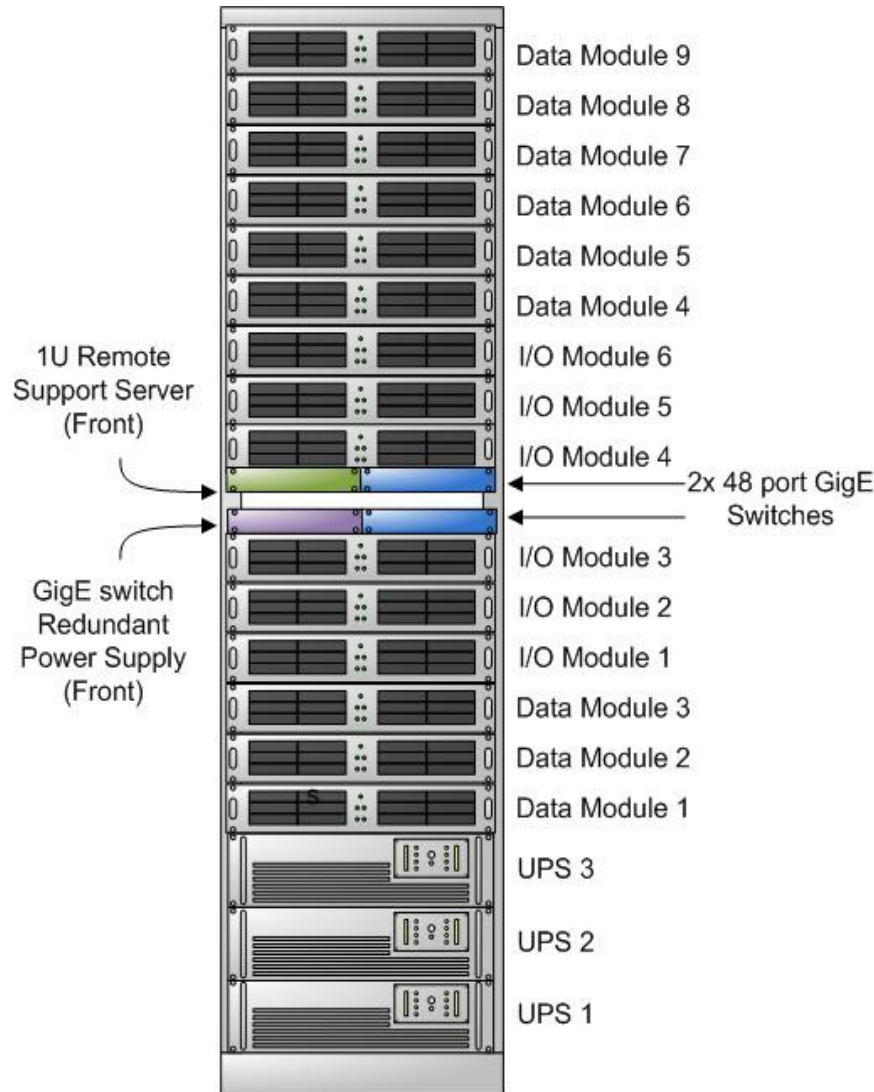
- Controller-centric
  - **Custom hardware & software**
- Centralized, shared cache
  - **Needs lots of shared bandwidth**
  - **Complex cache lock management**
- How do you scale beyond the controller?

- Distributed grid of commodity servers
  - **Software only**
- Distributed cache
  - **No shared bandwidth needed**
  - **No complex lock management**
- Scales in every dimension



Gain actionable insights  
from data

# XIV System Components



- System hardware
  - 15 modules with 12 drives per module
  - 120GB or 240GB memory
  - 24 x 4Gb FC ports
  - 6 x 1Gb iSCSI ports
  - 3 x redundant uninterruptible power supplies
  - 180 x 1TB or 2TB SATA Drives
  - 79TB or 161TB Useable
- SW features
  - Access control lists
  - NDU / CCL
  - LDAP
  - SMIS / TPC Compliancy
  - XIV Remote Support Center / VPN
  - Virtualized env / storage pools
  - Enterprise snapshot
    - 16K snaps / Unlimited
    - Snapshots of writeable snapshots
    - Redirect on write
    - Consistency groups
    - Full volume copies
  - Thin provisioning
  - Synchronous replication
  - Async 10.2
  - DM – Data Migration
  - Thick to Thin



Gain actionable insights  
from data

# XIV Storage System Distribution Algorithm



- Each volume is spread across all drives
- Data is “cut” into 1MB “partitions” and stored on the disks
- XIV distribution algorithm automatically distributes partitions randomly

XIV disks behave like connected vessels, as the distribution algorithm aims for constant disk equilibrium.

Thus, XIV’s overall disk usage approaches 100% in all usage scenarios.

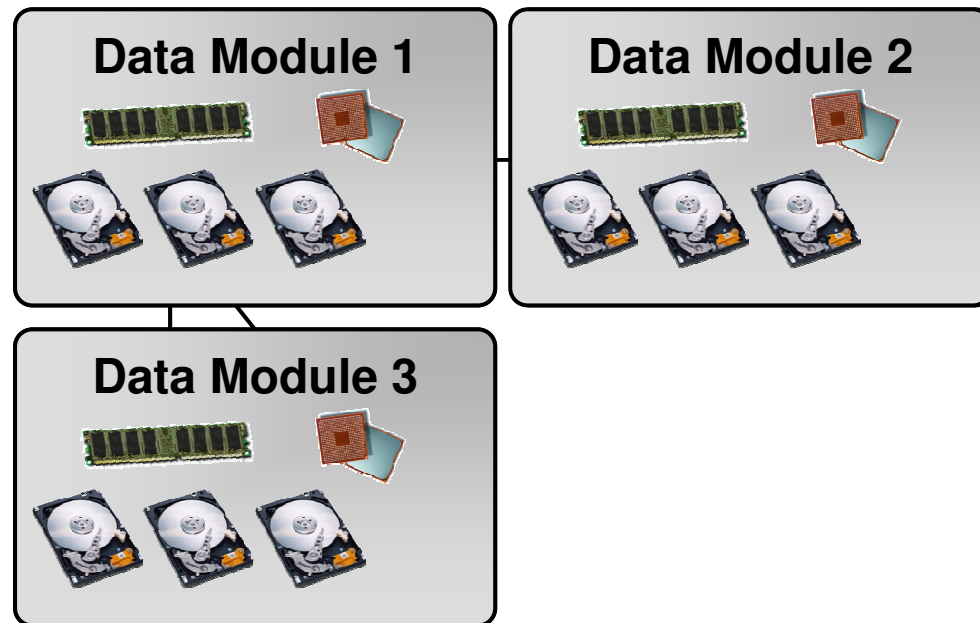



Gain actionable insights from data

# Distribution Algorithm on System Changes




- Data distribution only changes when the system changes
  - Equilibrium is kept when new hardware is added
  - Equilibrium is kept when old hardware is removed
  - Equilibrium is kept after a hardware failure



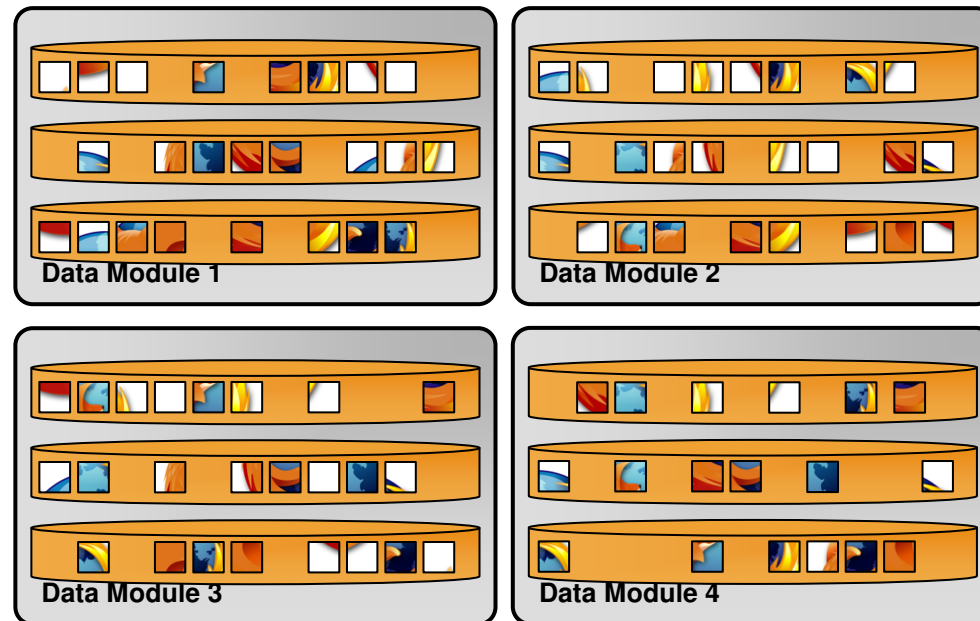
# Distribution Algorithm on System Changes



- Data distribution only changes when the system changes
  - **Equilibrium is kept when new hardware is added**
  - **Equilibrium is kept when old hardware is removed**
  - **Equilibrium is kept after a hardware failure**



[ hardware upgrade ]



Gain actionable insights  
from data

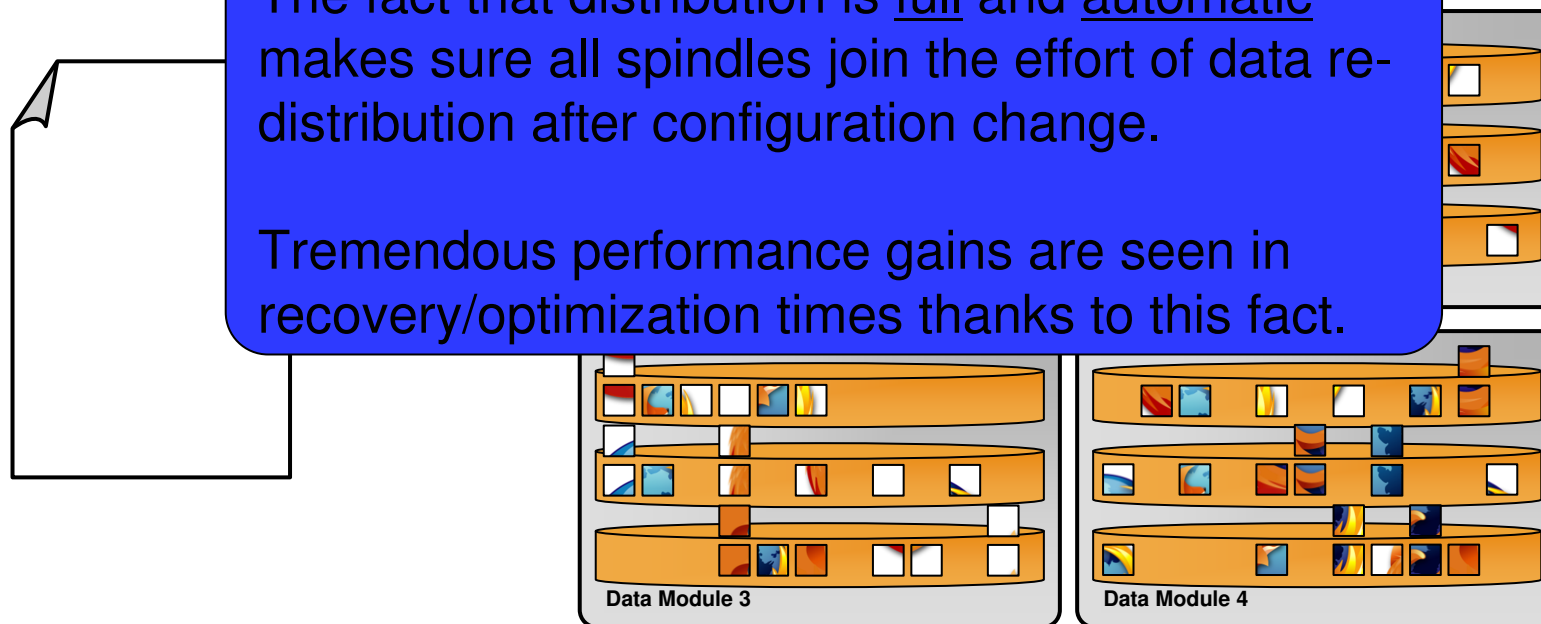
# Distribution Algorithm on System Changes



- Data distribution only changes when the system changes
  - **Equilibrium is kept when new hardware is added**
  - **Equilibrium is kept when old hardware is removed**
  - **Equilibrium is kept after a hardware failure**

The fact that distribution is full and automatic makes sure all spindles join the effort of data re-distribution after configuration change.

Tremendous performance gains are seen in recovery/optimization times thanks to this fact.

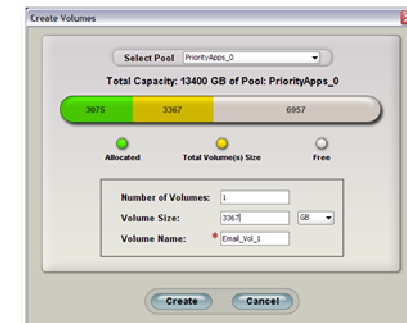




# Administration Made Easy




- Use your time to provision storage:
  - **Define volumes in seconds**
  - **Resize volumes painlessly**
  - **Create instant snapshots**
  - **Create test environments with writable snapshots**
- Don't spend time optimizing:
  - **No need for performance tuning**
  - **No need for complex layout patterns**
- Role-based management allows safe delegation of tasks
  - **Application administrators manage their own snapshots**
  - **Integrate with the organization's LDAP**



# Customer case

---



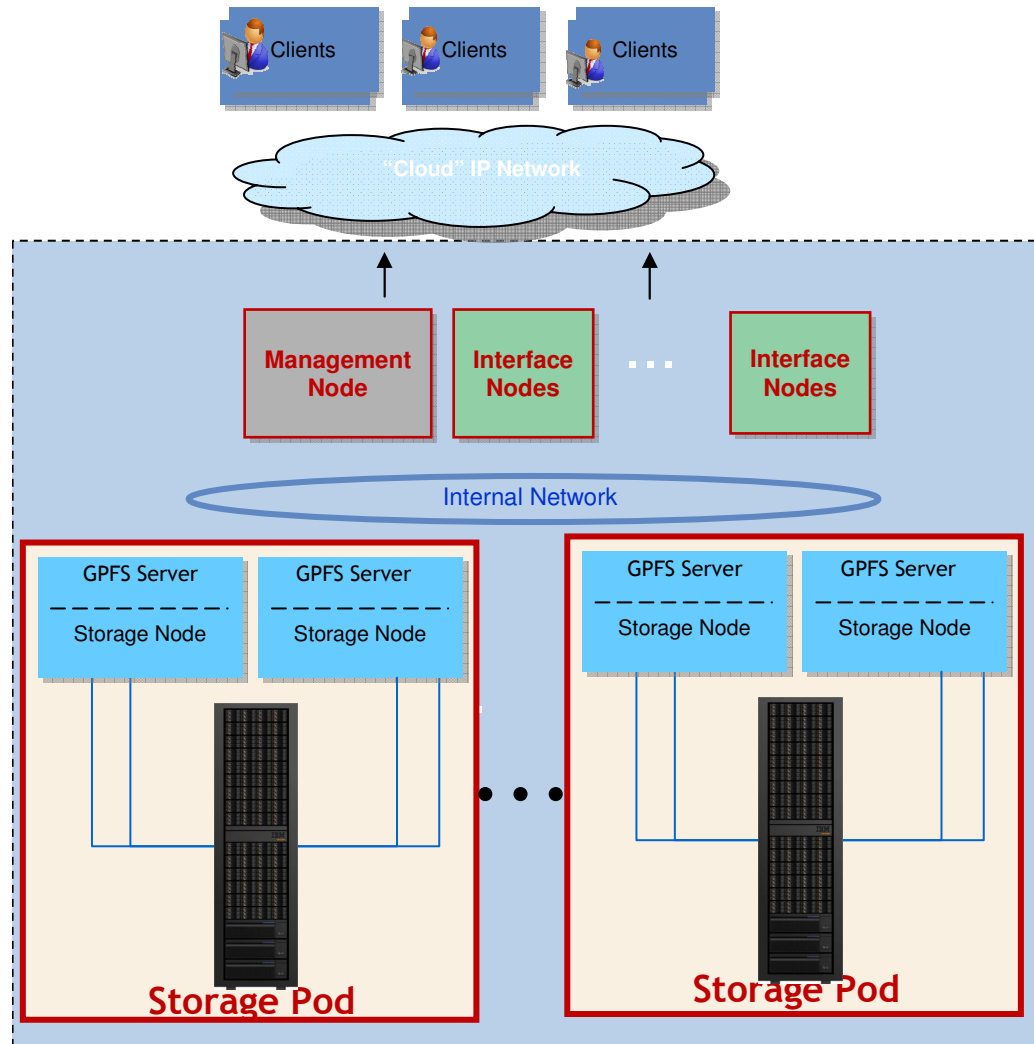
Gain actionable insights  
from data



IBM Scaled Out Network Attached Storage  
+  
IBM XIV  
=  
Manageability and scalability + performance



# SONAS and XIV – Architecture



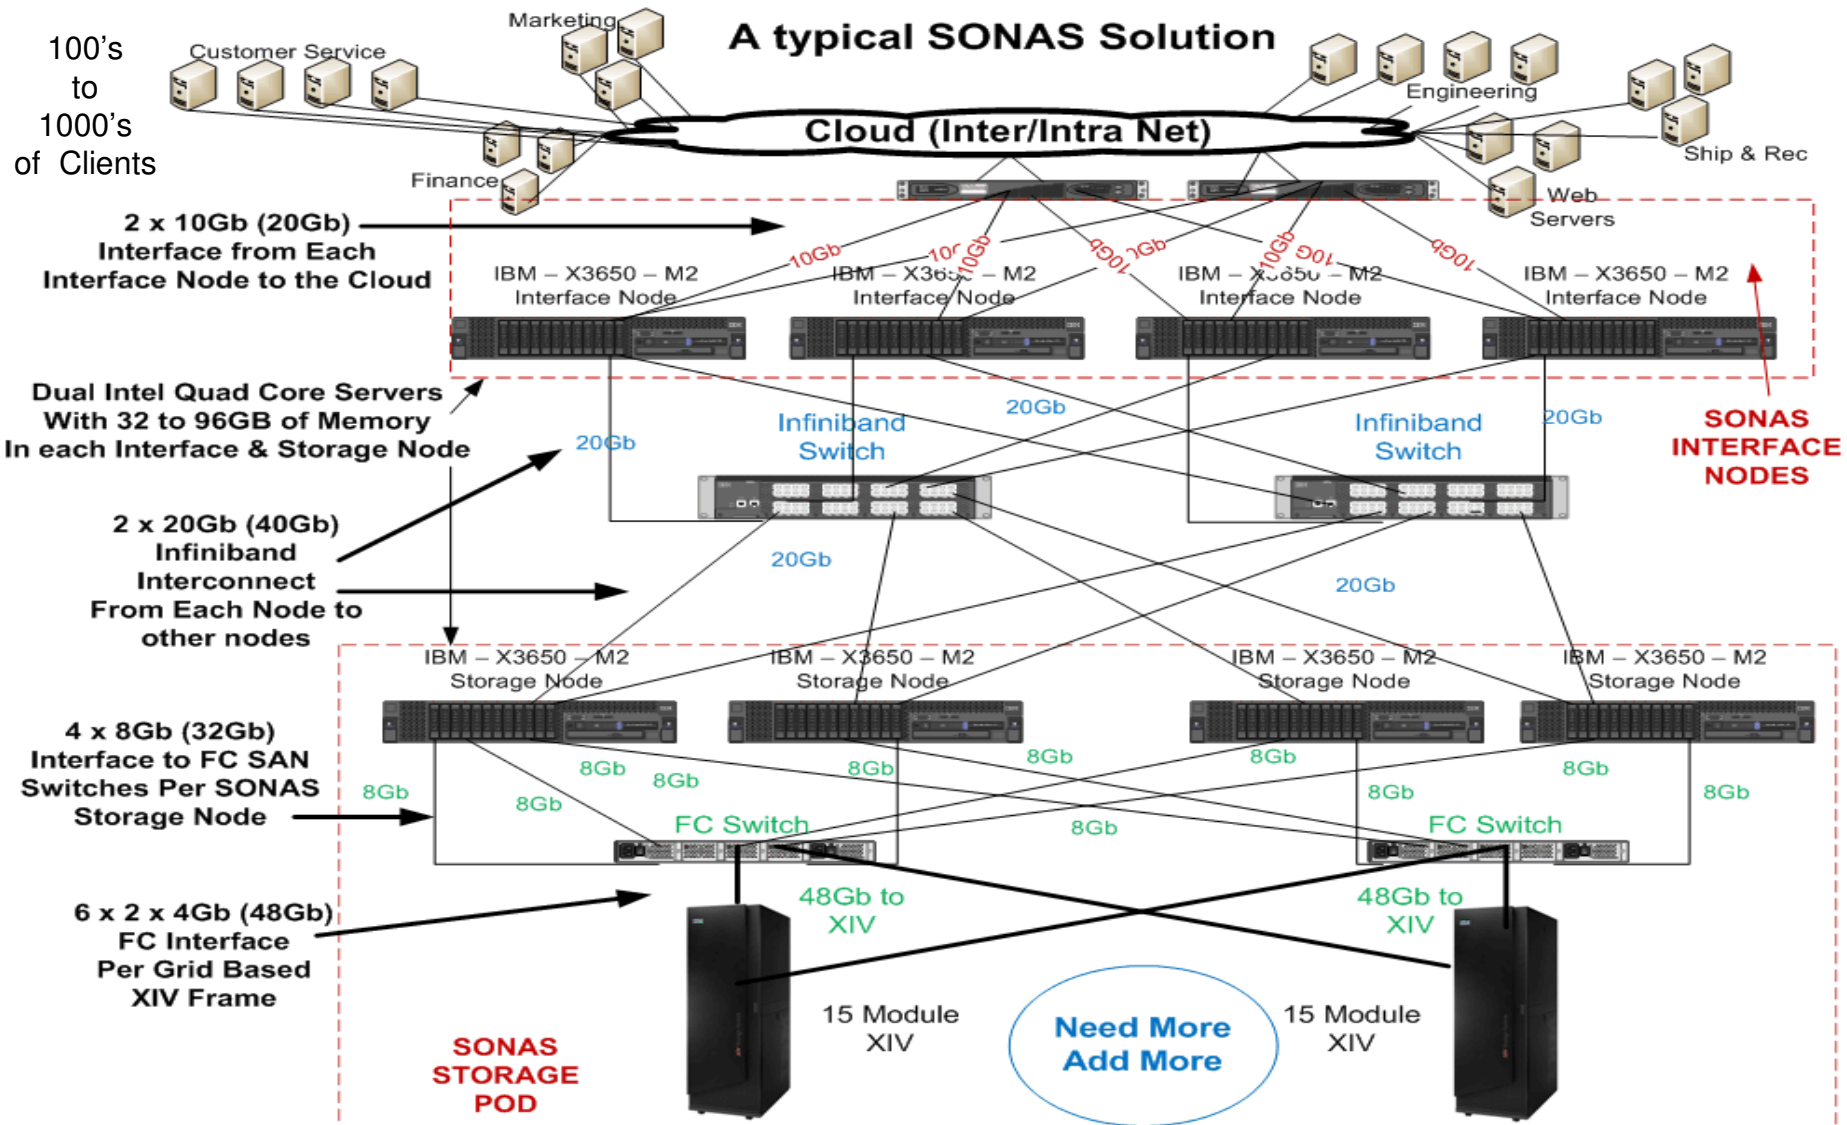
NFS, CIFS and other popular network file protocols

- 1GbE and 10GbE host interfaces
- Work load and data is evenly distributed across all nodes and disk pools, eliminating hot spots.
- All nodes continuously serve all files – every node is the same (not failover pairs).
- Near linear aggregate performance and capacity scaling
- Healthy nodes take over workload of failed node
- High speed, low latency (20 Gbps) private Infiniband cluster data network



Gain actionable insights from data

# SONAS and XIV Key Advantage - Scalability



Gain actionable insights from data

# Where?



## 3000 units

### Performance: SAP/Oracle:

- EMC DMX to XIV – 2X Capacity 60% Cap-Ex
- Reduction 5-6X Performance Sustained Through Failures – **Gerber Scientific**

### Scalability: VMware

- EMC DMX to XIV - 12,000 VMs on 10 XIV Frames Dramatically Faster Provisioning
- New Corporate Standard – **Airline Manufacturer**

### Migration – Exchange

- NetApp to XIV - Set Up/Use Same Day, No More
- SW Licensing, Predictable DR Costs OpEx Control/Savings – **Carillion Health**

### Ease of Management: SAS/Oracle

- HP EVA to XIV, 2X Performance 40% TCO Reduction, 30% Throughput Increase
- Saving Money, Exceeding SLAs Cost – Oracle & VMware - **Aetna Insurance**

### Cost – VMware

- EMC DMX to XIV VMware | Windows | Unix
- \$2.6M CapEx Savings Outperforms EMC w/Lower Human, Capital Cost – **Amylin Pharmaceutical**

### Manageless – Exchange

We are almost not managing the XIV system hosting our 200,000 exchange



Gain actionable insights from data

## 2 analysts



### Async – Snaps

throttle your RPO by app Slash Network Costs

### Then: Announcement

“For many workloads, this type of storage system appears to be the future of storage, offering lower acquisition cost, increased flexibility of data management, massive scalability and much easier management”

### Now: +1 year


“XIV could well be a piece of computer history in the making because its guiding light, when at EMC, once took on and beat IBM at its own game. Could EMC's former benefactor and acknowledged storage maven now become its biggest enterprise storage headache? Quite possibly.” -- John Webster, Principal IT Adviser, Illuminata



## 20 tier 1 customers tell all



Three-year Costs for Use of IBM XIV and EMC V-Max Systems for Tier 1 Applications



# 63%

Savings



- Innovative, enterprise class storage solution
- Key component in new IBM storage strategy
- Field proven since 2003 – aggressive, cross-Industry customer and market adoption
- Game-changing architecture – measurable, real world benefits – unrivaled TCO/ROI benefits
- SONAS + XIV bizarre and extreme NAS performance

***IBM XIV customer: “IBM XIV has reduced our total cost of operations while improving our service levels to our worldwide SAP user community. The XIV architecture has allowed Gerber Scientific to improve performance for operational tasks: one function that used to take eight hours is now completed in just 15 minutes, giving our IT staff time to focus on other projects.”***

Raf Cohen,  
CIO of Gerber Scientific



# Finally, the future?



Movie source: [www.dumpert.nl](http://www.dumpert.nl)



Gain actionable insights  
from data



---

# THANK YOU



[hugh.hulleman@nl.ibm.com](mailto:hugh.hulleman@nl.ibm.com)



Gain actionable insights  
from data