

Addressing Data Growth, Threats and Compliance: Unified Storage Management

Kalyan Pola,
World Wide Product Marketing Manager
Tivoli Storage Software



Business without **LIMITS 2012**

[30/8/2012, Bangalore]



Agenda

- Introduction
- Information Explosion
- Choices for reducing your data storage footprint
 - Avoid duplicating data – Treat the cause not the symptom
 - Categorize the data for migration & deletion
 - Automate the migration, archival and deletion
 - Compress and deduplicate what's left
 - Maximize utilization of storage resources



Information Explosion

Volume

Every day, **15 petabytes** of new information are being generated. **By 2010 onwards**, the codified information base of the world is expected to double every **11** hours.



Variety

80% of new data growth is unstructured content, generated largely by email, with increasing contribution by documents, images, and video and audio.



Velocity

An average company with 1,000 employees spends **\$5.3 million** a year to find information stored on its servers. **42%** of managers say they use the wrong information at least once per week.



Information Week Survey:

70% say Regulations drive up storage demands, businesses realizing they **must** classify, manage, delete data - not just keep adding storage

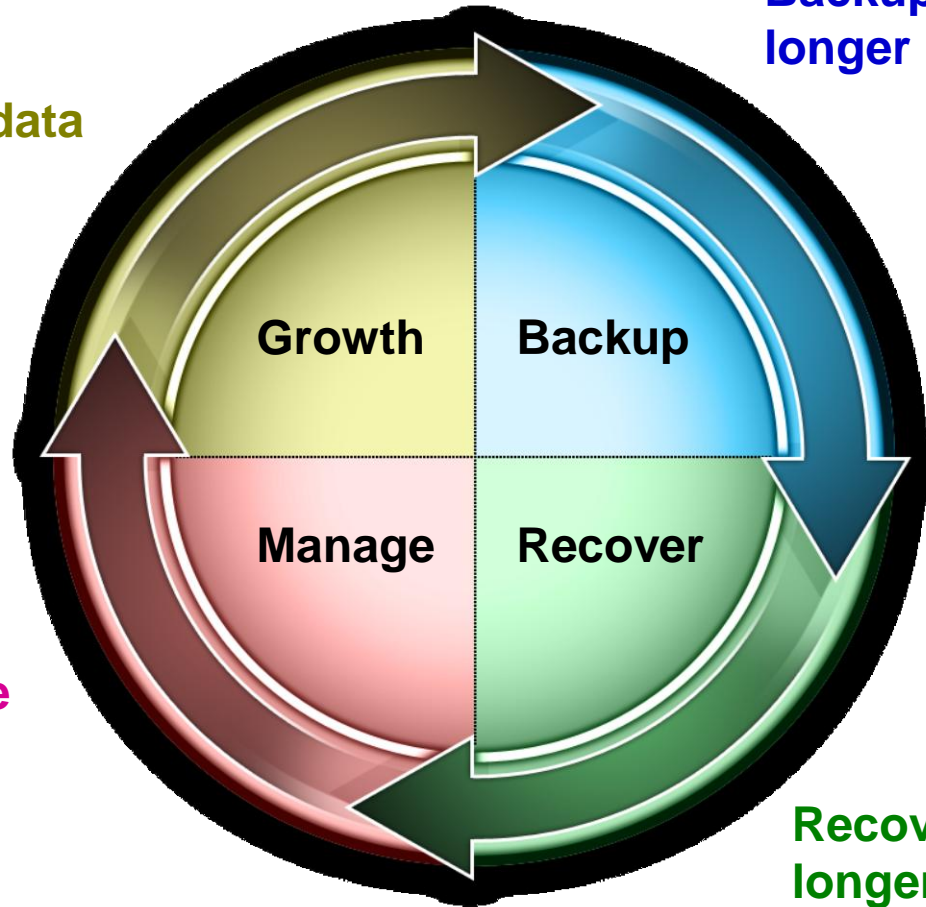


The pressures on administrators are growing

The consequences of data growth

More new data coming

Backup takes longer



Recovery Considerations



RPO



RTO



Labor, Systems Bandwidth Costs




Impact on Applications

Can't buy more storage

Recovery takes longer



Survey - what are your two biggest storage pain points?




* TheInfoPro Wave 14 Storage Study: F1000 Sample. n=149. Other n=14. *Multiple responses recorded



Surviving the Tidal Wave

Reducing your data storage footprint will:

- Reduce your costs
 - Less storage = less capital expenditures
 - Less data = simplified management and administration
- Improve service levels
 - Less downtime = higher application availability
 - Improved competitiveness and customer satisfaction
- Mitigate **Risks**
 - Eliminate consequences of data loss
 - Respond faster to events and legal/government inquiries – **Aid Compliance**




IBM can help you build a dynamic storage infrastructure that will intelligently improve service levels, reduce costs and manage risks



Choices for reducing your data storage footprint:

1. Avoid duplicating data – treat the cause, not the symptom
(Progressive Incremental BackUp)
 - **Periodic full backups** are the #1 cause of duplicate data
 - Performing progressive incremental backups eliminates duplication
2. Categorize the data for migration & deletion (TPC for Data)
 - Older data should be **moved off production systems**
 - This will shorten backup cycles and improve application performance
3. Automate the migration, archival and deletion (TSM HSM for Windows / TSM for Space Management)
 - Set policies based on business requirements
 - Move older, less-frequently accessed data to archive storage tiers
4. Compress and deduplicate what's left (Native De Dup & Appliance based De Dup)
 - Redundant copies may still exist on different source systems
 - **Deduplication** can reduce capacity requirements by another 40-95%
5. Maximize utilization of storage resources (VSC)
 - Gain better Visibility, Control and Automation - Management
 - **Virtualize heterogeneous storage** silos into a single pool of capacity



1. Avoid duplicating data – treat the cause, not the symptom



Business without **LIMITS 2012**



Comprehensive Data Reduction

Treat the cause, not the symptom

- Reduce Windows backup storage costs by 94%
- Reduce Oracle deduplication results – 75% reduction of backup storage.



Source (Client) -Side Data Deduplication

- Enables a 24:1 deduplication ratio



Target (server)-side deduplication

- Shortens backup times by 20 percent

Device side deduplication (ProtecTIER)

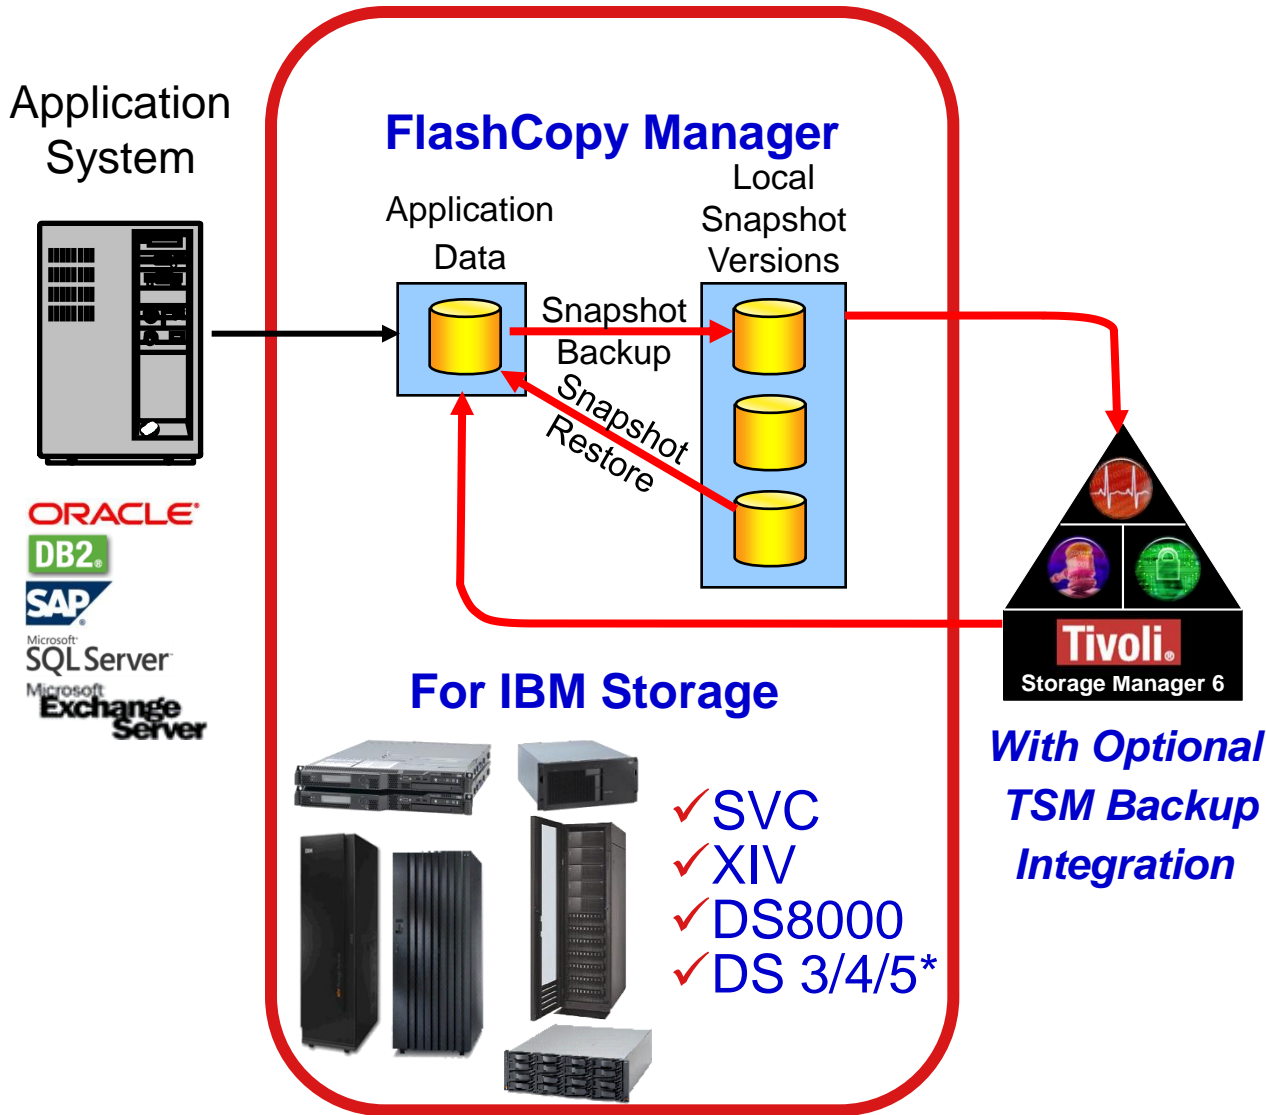
- Enables a 24:1 deduplication ratio
- Shortens backup times by 20 percent
- Reduces recovery times by more than 50 percent

Progressive Incremental backup for ever

- Backup storage capacity reduction by 70%



FlashCopy Manager Overview (for Data)



ORACLE®
DB2®
SAP
Microsoft
SQL Server
Microsoft
Exchange Server

- ✓ Online, near instant snapshot backups with minimal performance impact
- ✓ High performance, near instant restore capability
- ✓ Integrated with IBM Storage Hardware
- ✓ Simplified deployment

*VSS Integration



Where we've been ...

Leading the way for more than 19 years

**“Business runs on IT,
IT runs on Tivoli !”**

- Leading the way for more than 19 years
- IBM Continue to gain market share
- A major market share player
- Tivoli Storage Management Suite for Unified recovery will provide complete solution to meet RPO, RTO and SLAs.
- Comprehensive Data Reduction method
- *IBM has a highly scalable product*
- Realize measurable cost savings and superior ROI with a comprehensive storage management solution

Innovation from the start

- 1st disk-to-disk backup solution
- 1st to provide support for both disk and tape
- 1st to use hierarchical storage management for offline storage
- 1st to offer a fully incremental approach for backup window management
- 20,000+ loyal customers
- Industry's broadest range of OS and HW platform *from laptop to mainframe*
- *Emphasizing data reduction since 1993*



Where we've been ...

Tivoli Storage Manager – the INNOVATION continues

Tivoli Storage Manager has been designed to address our customers' most challenging issues: the continuing growth of data and the complexity of managing it.

New

in TSM v6.1 (03/2009):

- Increased scalability, availability & performance (embedded DB2 database)
- **Built-in target-side data deduplication**
- Simplified management & deployment
- Improved integration with VMware, NetApp / N series and Windows

New

in TSM v6.2 (03/2010):

- Tight integration with TSM FastBack
- **Built-in source-side data deduplication**
- Automatic deployment of Windows client updates
- Support for VMware vStorage API and VSS in Hyper-V

New

in TSM v6.2.2 (12/2010):

- Even better support for VMware, Windows and N series / NetApp environments
- High Availability of the TSM Server, using DB2 HADR Replication
- Introduction of **TSM for Virtual Environments** in March 2011 – optimized protection and flexible recovery for VMware vSphere
- Introduction of **TSM Suite for Unified Recovery** in June 2011 – capacity-priced bundle of 10 TSM and FastBack products




New enhancements in TSM 6.3

Tivoli Storage Manager v6.3

- Improve Return on Investment
 - Increase scalability to 4 Billion data objects per TSM Server
 - 800% improvement over TSM v5.5
- Simplify Administration
 - Automated client software updates across all platforms – simplify admin
 - Improved reporting and monitoring across the family; faster custom report creation
- Improve service levels and availability
 - New client data/metadata replication for warm standby disaster recovery
 - Faster internal database backup

TSM for Virtual Environments v6.3

- Simplify Administration
 - Support for VMware vSphere 5
 - Management plug-in for vCenter Client
 - Centrally manage multiple Linux and Windows vStorage Servers



Tivoli Storage Manager for z/OS Media

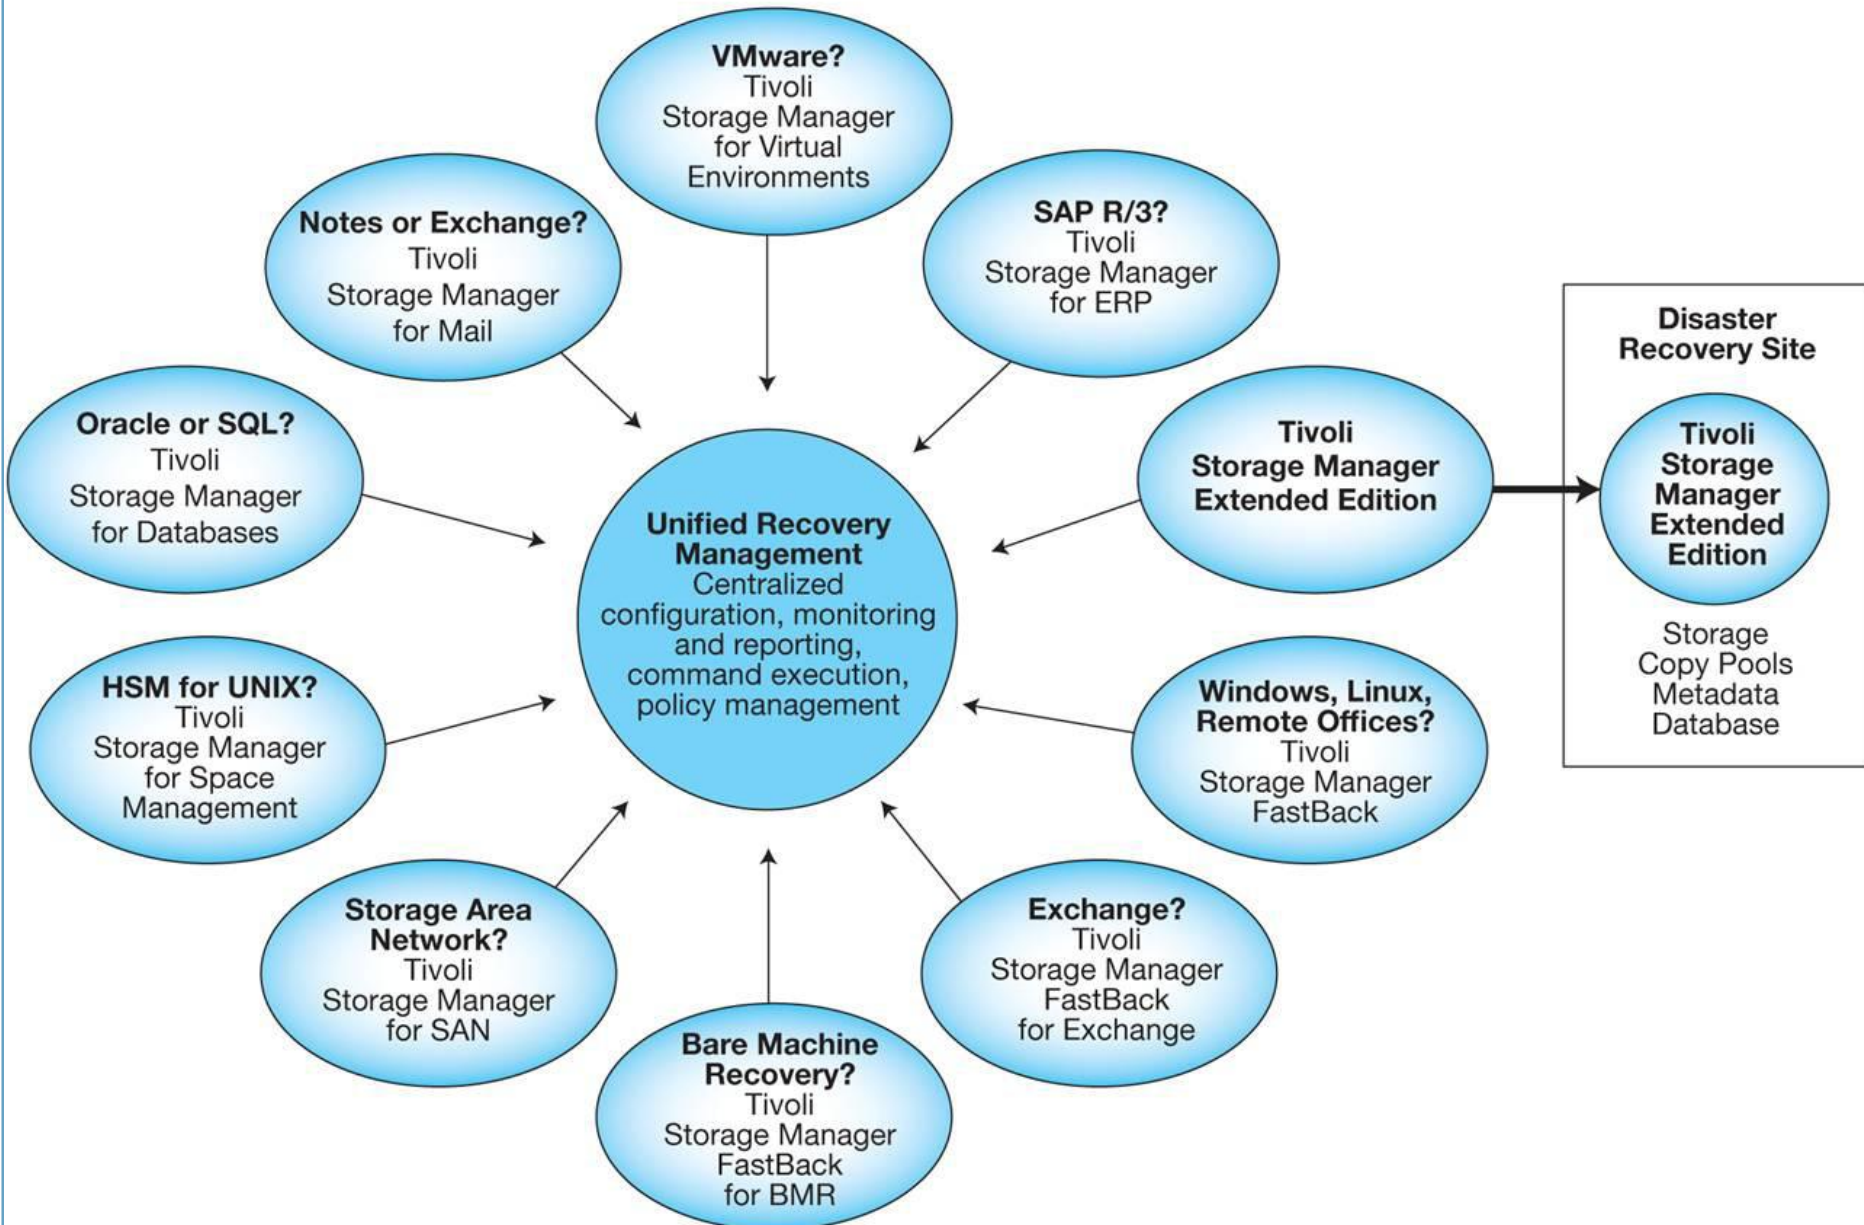
- Enable z/OS clients to take advantage of latest TSM features while leveraging existing mainframe-attached storage resources

Tivoli Storage FlashCopy Manager v3.1

- Improved coverage
 - HP-UX
 - VMware vSphere 4 and vSphere 5
 - Improved functionality with Microsoft Exchange and SQL



IBM Tivoli Storage Manager Suite for Unified Recovery



2. Categorize the Data Migration & Deletion



Business without **LIMITS 2012**



Questions ????? SRM Helps Answer

How much storage do I have available for my applications?

How is my storage being used?

How do I know the storage is not the bottleneck for user response time issues?

When do I need to buy more storage?

How well are my storage systems performing?

How do I maintain storage service levels?

How reliable is my SAN?

How do I simplify and centralize the management of my storage infrastructure?

How do I monitor and centrally manage my replication services?

How can I automate the provisioning of my storage systems, databases, file systems and SAN?

Which applications, users and databases are the primary consumers of my storage?

Is the storage infrastructure available and performing as needed?

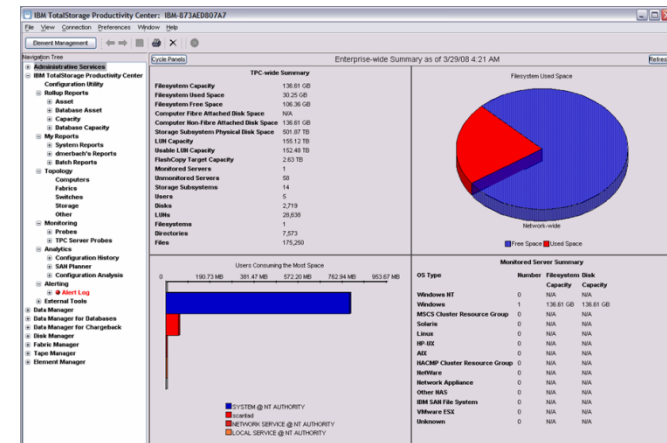
Which files must be backed up, archived and retained for compliance?



Determine What You Have BEFORE You Try To Fix It

IBM Tivoli Storage Productivity Center (for Data)

- Your files systems are probably bursting from data that is old and rarely accessed
- Some data can become a liability after it's outlived it's useful lifetime
 - Think e-discovery: do you know what was saved 5 years ago?
- Categorizing and then leaning out this old data from production systems will:
 - Reduce capacity requirements and lower CAPEX and OPEX
 - Improve backup and restore performance
 - Help meet data retention and expiration mandates
- TPC identifies data for migration and deletion:
 - By date saved or last accessed
 - By location and owner
 - By file type and size



Tivoli Storage Productivity Center – Summary of Offerings

What You Need to Manage

Servers

- ESX servers
- Apps, DB's, file systems
- Volume managers
- Host bus adaptors
- Virtual HBAs
- Multi-path drivers

Storage Networks

- Switches & Directors
- Virtual devices

Storage

- Multi-vendor storage
- Storage array provisioning
- Virtualization / Vol. mapping
- Block + NAS, VMFS
- Tape libraries

Replication

- FlashCopy
- Metro Mirror
- Metro Global Mirror

TPC Can Help

Start Here



TPC 5.1


- ✓ Single management console
- ✓ Heterogeneous storage
- ✓ Health monitoring
- ✓ Capacity mgmt.
- ✓ Provisioning
- ✓ Fabric management
- ✓ FlashCopy support
- ✓ Storage System Performance Management
- ✓ SAN Fabric Performance management
- ✓ Trend Analysis
- ✓ DR & Business Continuity
- ✓ Applications & Storage
- ✓ Hyperswap Mgmt.

... and Mature

IBM SmartCloud Virtual Storage Center

All this and more...

- ✓ Advanced SAN Planning and provisioning based on best practices
- ✓ Proactive configuration change management
- ✓ Performance optimization
- ✓ Tiering Optimization
- ✓ Complete SAN fabric performance mgmt.
- ✓ Storage Virtualization
- ✓ Application Aware flashcopy management



Business without **LIMITS 2012**



3. Automate the Migration, Archival and Deletion



Business without **LIMITS 2012**



Move data to the right place


70% of data

hasn't been accessed
in **90 days** or more

**Why are you storing it on
expensive primary storage?**



Data lifecycle management



Regularly scrub production systems of old/stale data

- Automated processes based on business-driven policies

Reduce the growth of primary storage – reduce CAPEX

- Faster backup / restore



Migration / Hierarchical Storage Management (HSM):

- Leaves a pointer; enables transparent access to migrated data

Archive

- Completely removes files from production systems
- Supports data retention and version control policies



Is Long Term Preservation a Real Issue?

Finance

Rule 17a-4 requires broker-dealers to retain account record information for six years. The **six-year period begins either at the time the account is closed or when the information is replaced or updated**

Life insurance policies have to be kept for **life of policy plus 6-10 years**

Healthcare

Medical records should be preserved for the **life of the individual and beyond**

OSHA requires employers to keep records of both medical and other employees who are exposed to toxic substances and harmful agents. Employers must maintain these records for **30 years**

Government

Land registry records, social security records, etc. **Life of individual to forever**

X-rays are often stored for periods of **75 years**

The retention requirement for the [medical] records of minors is **20 to 43 years of age**

Aerospace

Aircraft designs records have to be retained for the lifetime of aircraft (**60+ years**)

Pharma

Pharma needs off-line electronic data storage for **50 to 100 years or longer**

Petroleum

Oil-field data is used over life of field (**50+ years**)

Scientific and Cultural

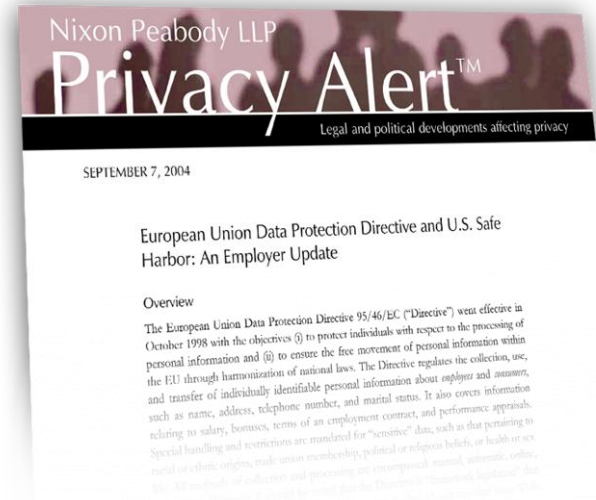
Satellite data is kept **for ever**

We would like to keep Libraries and Art data **for ever**



Most Important – It's The Law

- Regulatory and corporate governance issues are requiring firms to re-evaluate their current compliance and risk management infrastructure
- Companies may face multiple regulations which share or drive different requirements
- Corporations are mandated to:
 - **Ensure records are unaltered and undeleted**
 - **Adhere to internal & external records management policies**
 - **Quickly discover and produce business records for audit purposes**
- Sample regulations include:
 - **Sarbanes- Oxley, USA PATRIOT Act, Basel II, SEC 17a-4, HIPAA, Many others**



\$1.5 million

The average cost to defend a corporate lawsuit per case (Gartner, 2006)

The average e-discovery request can cost from

\$150K to \$250K

Sources: CIO Magazine survey 2007; IBM Tivoli Market needs and profiling study 2005;



IBM Strategic Archiving & Retention Offerings

Data content layer

- IBM CommonStore, IBM Content Manager, IBM Content Collector
- IBM FileNet P8 Content Manager, Image Services, SAP Connector
- IBM Records Crawler
- IBM Optim
- Grid Medical Archive Solution (GMAS)

Policy Management layer

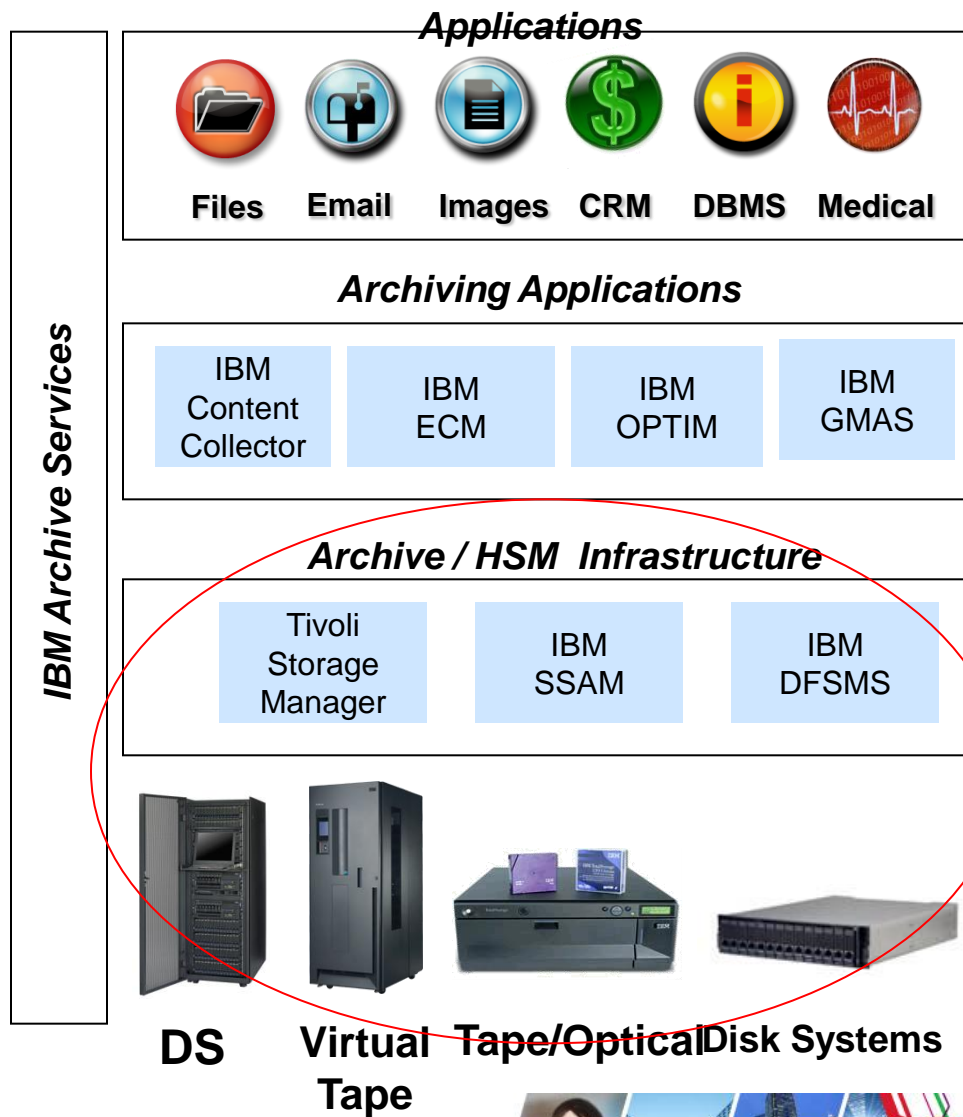
- Tivoli Storage Manager
- System Storage Archive Manager

Storage Media layer

- IBM Disk and Tape Systems

Services layer

- IBM Enterprise Archive Services



4. Deduplicate and Compress What's Left



Business without **LIMITS 2012**



Enhanced Data Reduction in Tivoli Storage Manager 6



Built-in Data Deduplication (Server & Client side)

- Tivoli Storage Manager 6 Extended Edition includes data deduplication for disk storage pools, *at no extra charge (Natively)*
- Improves recovery times as many more recovery points can be stored on disk; or reduces the amount of capacity needed
- Effective with data from any source including: API, backup, HSM, archive
- Builds on automatic data compression
- [ESG Lab Report](#) confirms 95% reduction ratio after just 11 weeks of backups
- Complementary with the IBM TS7650G ProtecTIER Gateway
 - TSM 6 built-in data deduplication and the TS7650G ProtecTIER Gateway can coexist to meet different customer needs

Tivoli Storage Manager 6, combining progressive incremental data capture with data de-duplication, does a **better job** of reducing storage capacity requirements than pure data deduplication solutions



IBM is delivering deduplication across its entire storage portfolio

Addressing the full spectrum of customer requirements for data reduction

Complementary Solutions Today!

TS7650G ProtecTIER Gateway

10+TB/night, >100MB/s performance

Easy management (via an appliance model)

Minimal overall operational impact (via no post processing)



Dedup data from multiple backup vendors including TSM

Deduplication is only done once

- For customers who need to**
- ✓ Lower operational costs and energy usage
 - ✓ Simplify and accelerate information protection
 - ✓ Reduce amount of disk and improve manageability

But can be done at multiple enterprise locations

No impact during backups/ingest

Fully integrated into TSM



No specific hardware dependencies

Dedup TSM legacy data

Tivoli Storage Manager



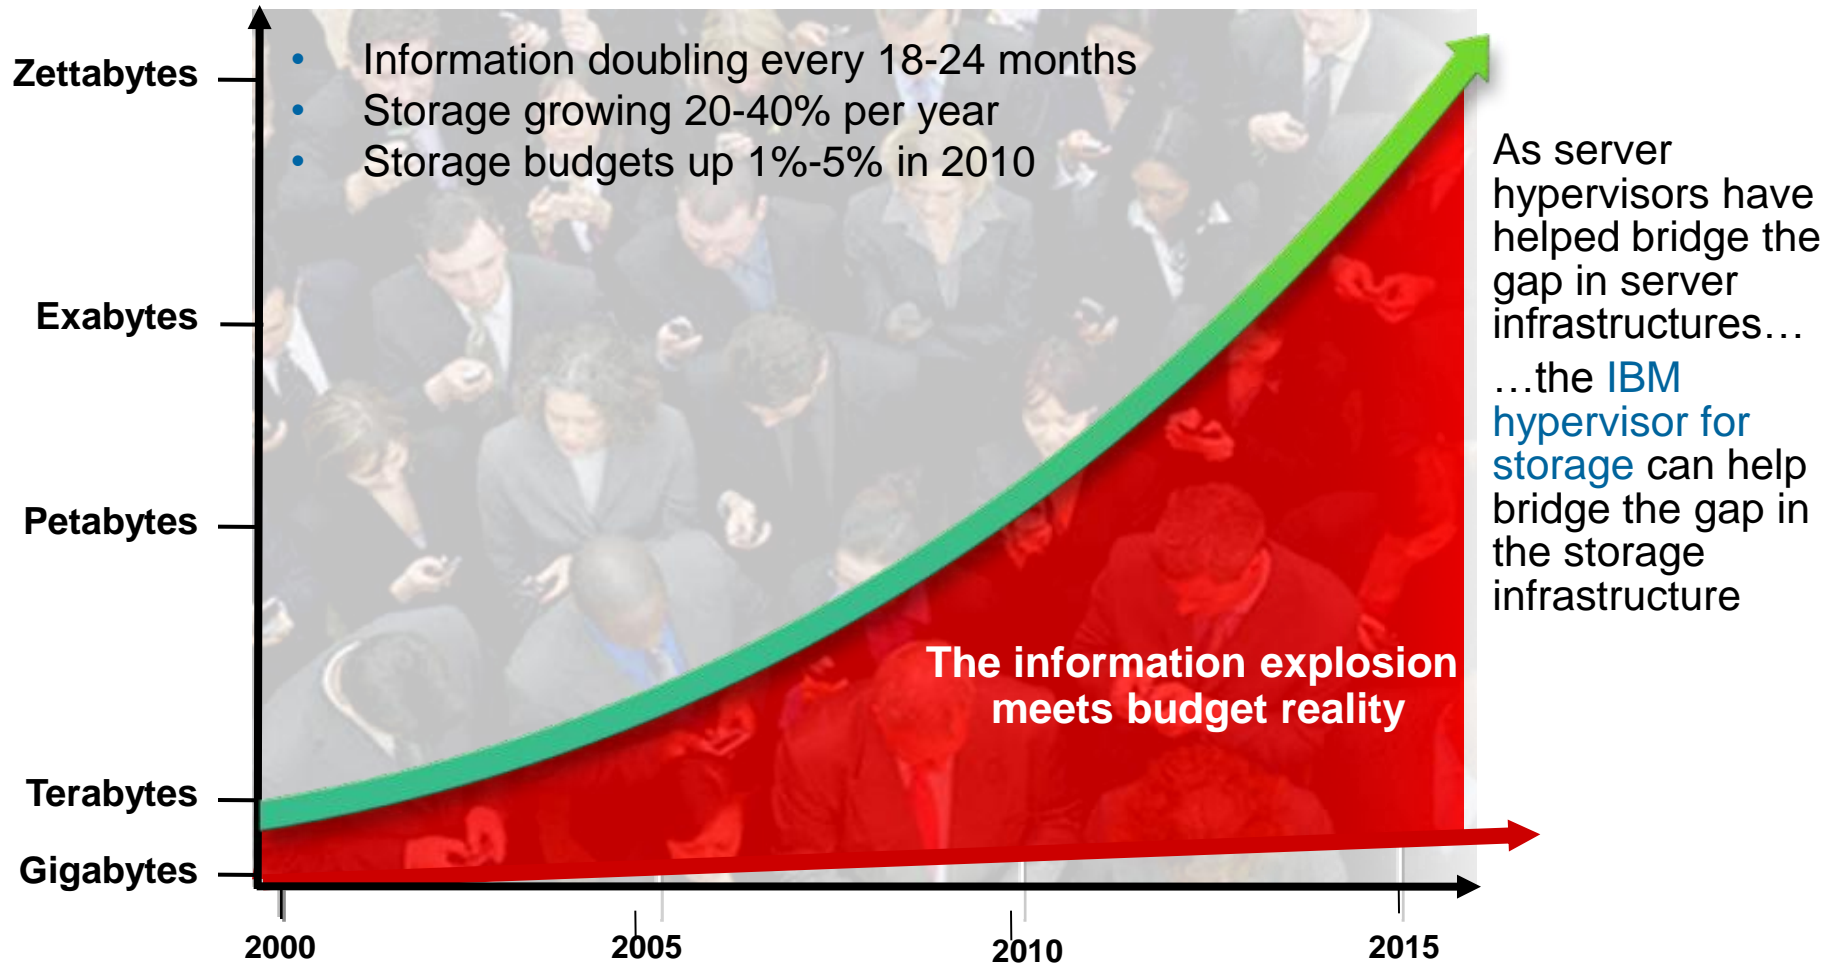
5. Maximize Storage Utilization



Business without **LIMITS 2012**



The Information Explosion



What's a Storage hypervisor?

- **Server hypervisor** has a well understood value connotation
 - Pooled physical resources are consumed by virtual machines resulting in high asset utilization
 - Virtual machines are mobile giving CIO's their choice of physical server device
 - A common set of value capabilities and centralized management are provided for virtual machines regardless of what physical server they are sitting on
 - ...cost savings and flexibility!

On Intel systems

server virtualization platform (VMware vSphere) and
server virtualization management (VMWare vCenter)

On Power systems

server virtualization platform (IBM PowerVM) and
server virtualization management (IBM Director VMControl)

- **Storage hypervisor** is a rapidly emerging way of describing the same value points in a storage context

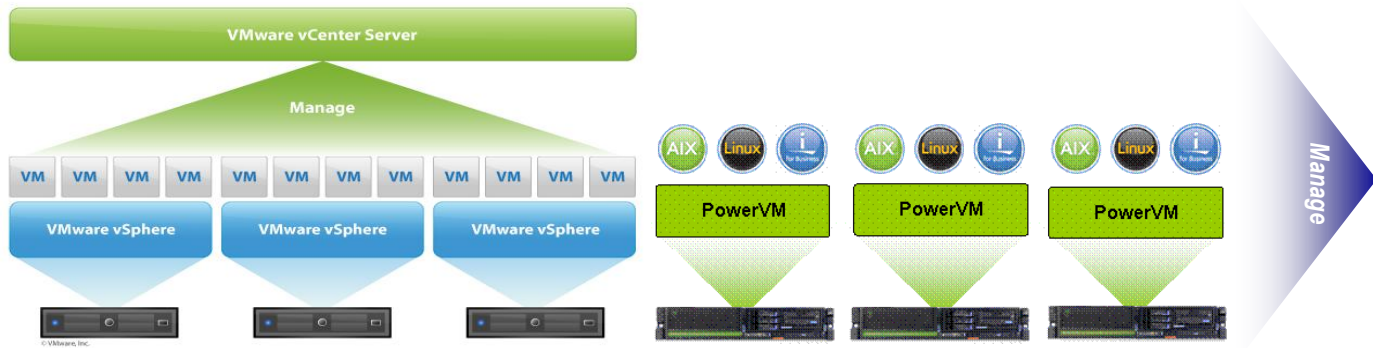


IBM's Storage hypervisor -



Storage Hypervisor


Virtual Server Infrastructure



Virtual Storage Infrastructure



(SAN Volume Controller)



VM Control

IBM Systems Director

Tivoli Storage Productivity Center

- **Virtual Storage Platform - SAN Volume Controller**
 - Common device driver - iSCSI or FC host attach**
 - **Common capabilities**
 - I/O caching and cross-site cache coherency
 - Thin provisioning
 - Easy Tier automated tiering to Solid-state Disks
 - Snapshot (FlashCopy)
 - Mirroring (Synchronous and Asynchronous)
 - **Data mobility**
 - Transparent data migration among arrays and across tiers
 - Snapshot and mirroring across arrays and tiers

- **Virtual Storage Platform Management - Tivoli Storage Productivity Center**
 - **Manageability**
 - Integrated SAN-wide Management with Tivoli Storage Productivity Center
 - Integrated IBM server and storage management (Systems Director Storage Control)
 - **Replication**
 - Application integrated FlashCopy
 - DR automation
 - **High Availability**
 - Stretch Cluster HA



IBM InterConnect 2012

October 9 - 11 | Resorts World Sentosa, Singapore



Turning Opportunities into Outcomes

In this era of interconnected industries, businesses and consumers, a new kind of leadership is required to turn opportunity into business outcomes. Smarter businesses are capitalizing on information as an indispensable resource and using technology as the catalyst for unleashing innovation. They are expanding the digital world of the back-office into the front-office.

Given this new reality, Business and IT leaders are collaborating to better align business and technology investments in order to respond to three business imperatives:

- **Re-invent relationships and uncover new markets**
- **Manage the velocity of business change**
- **Implement the new economics of IT to fund new innovations**



Hot Topics

- Changing the Economics of IT with IBM PureSystems
- Defending Against Cyber-Threats with Security Intelligence and Behavioral Analytics
- Rethink IT. Reinvent Business with Cloud Computing
- Transforming Critical Business Processes
- Unlocking Opportunities with Smarter Analytics and Big Data
- Gaining Competitive Advantage Through Software Innovation
- Creating Exceptional Experiences By Combining Social and Commerce Best Practices
- Speeding Innovation and Extend Reach Securely with Mobile Enterprise
- Transforming IT for Insight and Efficiency with Smarter Storage
- Enabling Growth with Enterprise Systems



Additional Resources

- [Introductory Demo](#)
- Product Information:
 - Tivoli Storage Manager [Data Sheet](#) [Web Page](#) [Tech Info](#)
 - Tivoli Storage Manager Suite for Unified Recovery [Data Sheet](#) [Web Page](#) [Tech Info](#)
 - Tivoli Storage Manager for Virtual Environments [Data Sheet](#) [Web Page](#)
 - Tivoli Storage Manager FastBack [Data Sheet](#) [Web Page](#) [Tech Info](#)
 - Tivoli Storage Manager FastBack for Workstations [Data Sheet](#) [Web Page](#)
 - Tivoli Storage FlashCopy Manager [Data Sheet](#) [Web Page](#)
- White Papers:
 - [Ten ways to save money with Tivoli Storage Manager](#)
 - [ESG: Tivoli Storage Manager for Virtual Environments](#)
 - [Unified Recovery Management](#)
 - [Extend Tivoli Storage Manager to the Cloud](#)
 - [Migrate to IBM Tivoli Storage Manager Easily and Confidently](#)
 - [Leverage the IBM Tivoli Competitive Advantages in Storage Management](#)
 - [Using IBM Data Reduction Solutions to Manage More Data with Less Infrastructure](#)
- Solution information:
 - [Data Protection and Storage Management](#)
- DeveloperWorks: [TSM Wiki](#)



Thank You!



Business without **LIMITS 2012**



Acknowledgements, Disclaimers and Trademarks

© Copyright IBM Corporation 2012. All rights reserved.

The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

References in this publication to IBM products, programs or services do not imply that they will be made available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth, savings or other results. All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information concerning non-IBM products and services was obtained from a supplier of those products and services. IBM has not tested these products or services and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products and services. Questions on the capabilities of non-IBM products and services should be addressed to the supplier of those products and services.

All customer examples cited or described are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer and will vary depending on individual customer configurations and conditions. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

Prices are suggested U.S. list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

IBM, the IBM logo, ibm.com, and other IBM products and services are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml

