



# Tivoli Storage Manager for Virtual Environments

Referent: **Andreas Kindlbacher** – Tivoli Storage Sales Leader



## DISCLAIMER

**The information on the new product is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product is for informational purposes only and may not be incorporated into any contract. The information on the new product is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion.**



## Opportunity Identification – TSM for VE

### Questions to ask:

- Do you use VMware virtual machines as production servers?
  - By 2012, most x86 server workloads will be running in VMs
- Do you need to improve the protection of data in virtual servers?
  - Traditional in-guest backup solutions drag down the performance of virtual servers and may not provide adequate Recovery Point Objectives for more critical data
- Do you need to improve the recoverability of virtual server data?
  - Can you easily recover an individual file, a disk volume or an entire virtual machine?
  - Would you like to recover disk volumes within a couple of minutes?

Improve  
Service

Reduce  
Costs

Mitigate  
Risks



## Pricing

- Priced by Process Value Units (PVUs) for the cores of the vSphere host running the protected guests
  - US Price = \$35.00 per 10 PVU
  - TSM or TSM EE is a prerequisite product for TSM for VE; the price is in addition to the TSM or TSM EE b/a client license covering the vSphere server
- No licensing is required for the vStorage backup server, whether it is a physical or a virtual system
- Sub-capacity pricing is available
  - <http://www.ibm.com/software/lotus/passportadvantage/subcaplicensing.html>



## Competition (see speaker notes for details, additional links)

	<b>Competitors, as of Feb 2011:</b>	<b>IBM TSM for VE</b>	<b>Symantec NetBackup</b>	<b>EMC Networker + Avamar</b>	<b>Commvault Simpana</b>	<b>Quest vRanger Pro</b>	<b>Veeam</b>
1	<b>Block Level Backup</b>	Yes	No, only file level	In-Guest and image-level backups	Granular File level recovery from a full image backup	No (Image based backup)	File level
2	<b>Flexible recovery options</b>	File, Volume, and VM image	File, Volume, and VM image	File, Volume, and VM image	File and full VM image	Instant file level recovery	File, Volume, and VM Image
3	<b>Near Instant Restore</b>	Yes	Instant restore available	Yes	Yes – only for VM	No	Yes
4	<b>Auto-Discovery of new VM's</b>	Yes	No	No	Yes	Yes	No
5	<b>Support for LAN-Free data transfer</b>	Yes	Needs an agent	WAN friendly	No	Yes	Yes
6	<b>Scalability</b>	Manages up to 2 billion objects	Need to add media server to grow capacity	Need to add media server to grow capacity	Limited	Yes	Yes
7	<b>Data deduplication</b>	Built in	Yes, extra \$\$\$	Yes, Source Based, extra \$\$\$	Integrated	No	Limited
8	<b>Unified Recovery Management</b>	Yes	Yes	Yes	Limited	No	No






# Tivoli Storage Manager for Virtual Environments

## Product Overview




## Backup used to be so easy



- Install a backup agent on production server, schedule the backup, copy the data periodically to the backup server



## It's not so easy anymore



- Proliferation of different system platforms, applications, physical locations all with different service level requirements
- Lots of different things can go wrong, all requiring different recovery capabilities
- Need to balance the needs of the business (data availability, resiliency) against the needs of the business (costs, resources, system availability)






## Virtual server environments have changed the game

- Virtual machines take advantage of the growth in processor power, memory and I/O
  - The computers get bigger, but the application workloads have remained flat
- Virtual machines can be created very quickly
  - Minutes to provision, rather than weeks to request, procure, install and test
- Tremendous savings in hardware, power/cooling, floor space and management costs
- VMs enable new levels of service management and new delivery models
  - Quickly create new servers for application development and testing, special projects, periodic workflows like financial reporting, or to expand capacity on demand
  - Virtual servers are a basic component of public and private cloud infrastructures



## Changing the game, for better and for worse ...



“Virtualized servers have been a boon to the systems side of the house, but a bane for storage managers.”

Hot Technologies for 2011,  
STORAGE Magazine, Dec 2010 [\(link\)](#)



## VMs have made data and storage management more complex



- Data continues to grow at unprecedented rates, across all industries
- Server virtualization does not reduce the data
- When a virtual server is de-commissioned or moved, what happens to its data, and the storage capacity that it consumed?
- How do you manage data management processes in virtual environments?
  - Backup / Restore
  - Disaster Recovery
  - Data Lifecycle Management



## Analyst Predictions

More than 90% of users deploying virtual machines are doing so specifically to reduce x86 server space and energy costs.

VM densities are increasing on physical servers, 1:15 today

By 2012, most x86 server workloads will be running in a VM.

Roughly 80% of the virtual machines today are utilizing “traditional” in-guest backup or host backups.

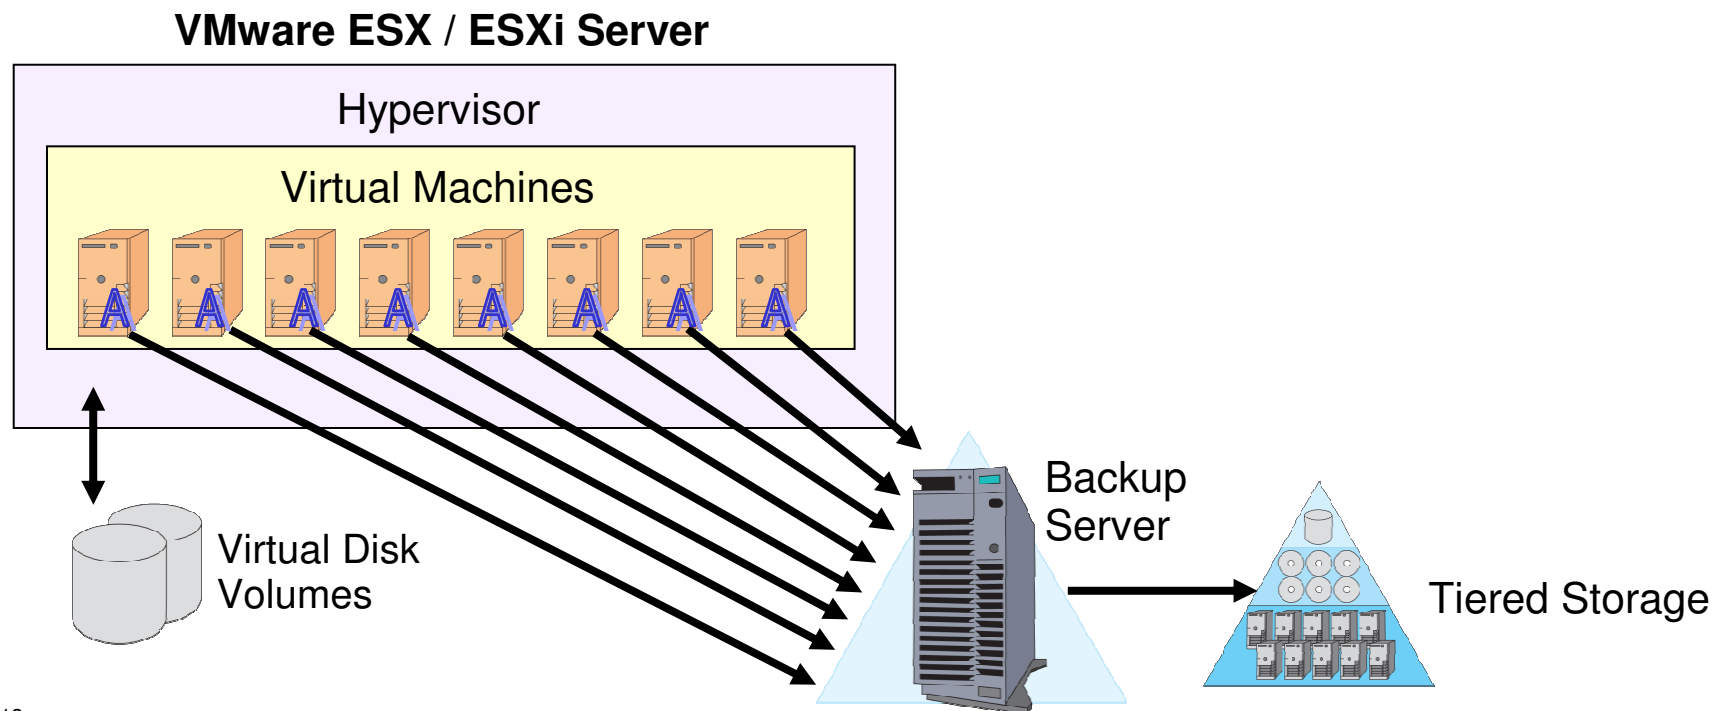
***How do you manage data protection and recovery in your virtual servers?***

Sources – Gartner, IDC, ESG




## Traditional in-guest backup and recovery model

- Install a backup agent in the guest OS, just like a physical server
- Run and manage backups just like in a physical server environment
- **Downside:** deploying, managing, maintaining ‘backup agent sprawl’
- **Downside:** can put a serious drain on processor, memory, I/O resources
  - From running multiple backups at once; file system scans during incremental backups; etc.



## Traditional in-guest backup using TSM and TSM FastBack



- Still the predominant approach
- Heterogeneous support for ...
  - Hypervisors: **KVM, Hyper-V, VMware, LPARs, Solaris Containers, HP nPartitions, etc.**
  - Guest Operating Systems: **Windows, Linux, z/Linux, z/OS, Solaris, more**
- TSM Fastback is a very good fit for this approach (on supported platforms)
  - Block-level incremental forever backup, in the background

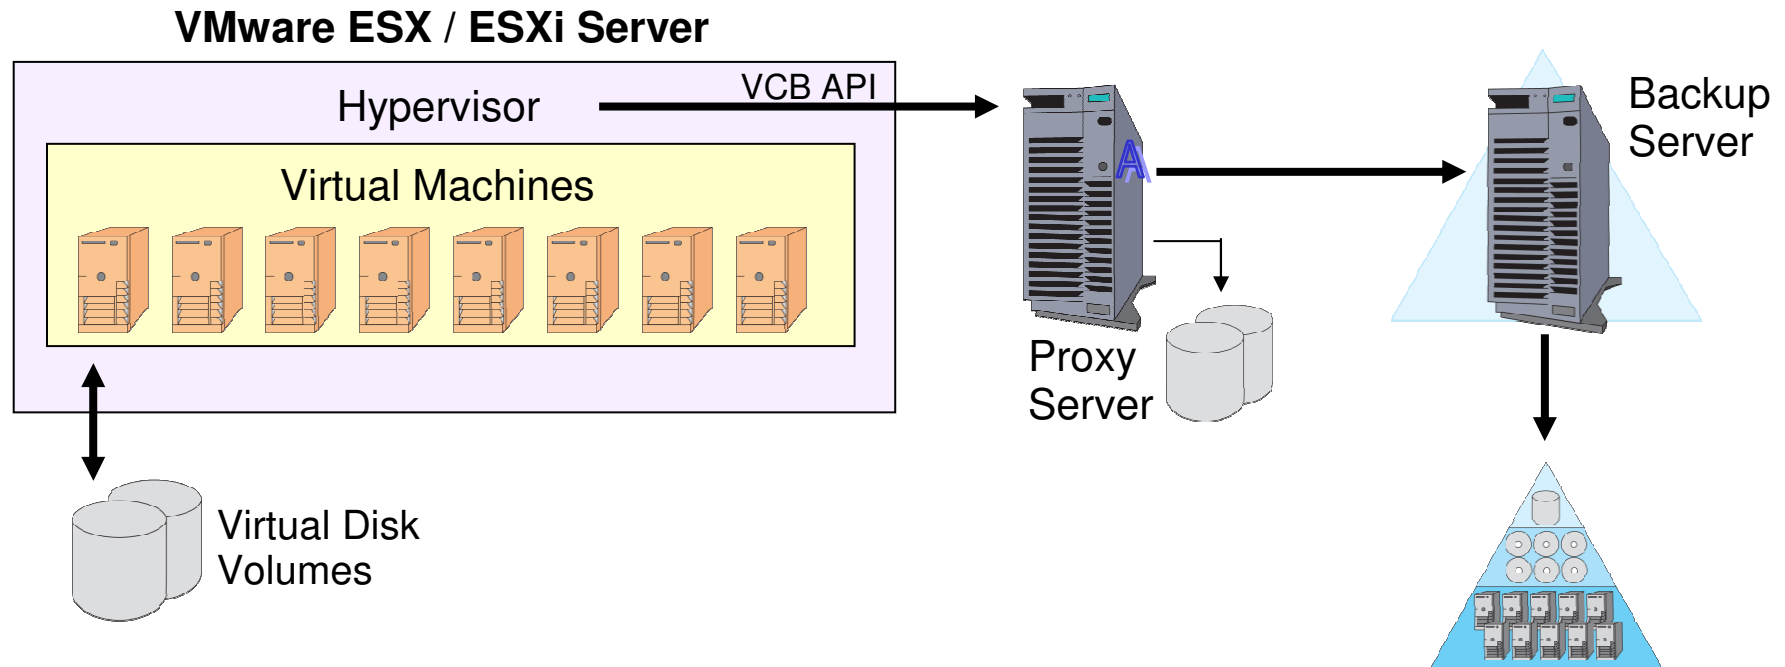
Pros	Cons
Better granularity in recovery	Multiple agents
Application awareness during backup	No VMware integration
Better recovery for applications	Management challenges
Business as usual, use existing management methodologies	Possible resource constraints

Smarter Storage Management

## VMware Consolidated Backup (VCB)

*VMware's first approach to offload backups from VMs*

- Snapshots of the VMs are taken by the Hypervisor and sent to a Proxy Server, which is then backed up by the Backup Server
- **Downside:** requires more hardware (server and storage)
- **Downside:** multi-step recovery through the Proxy Server



## Using VCB with Tivoli Storage Manager

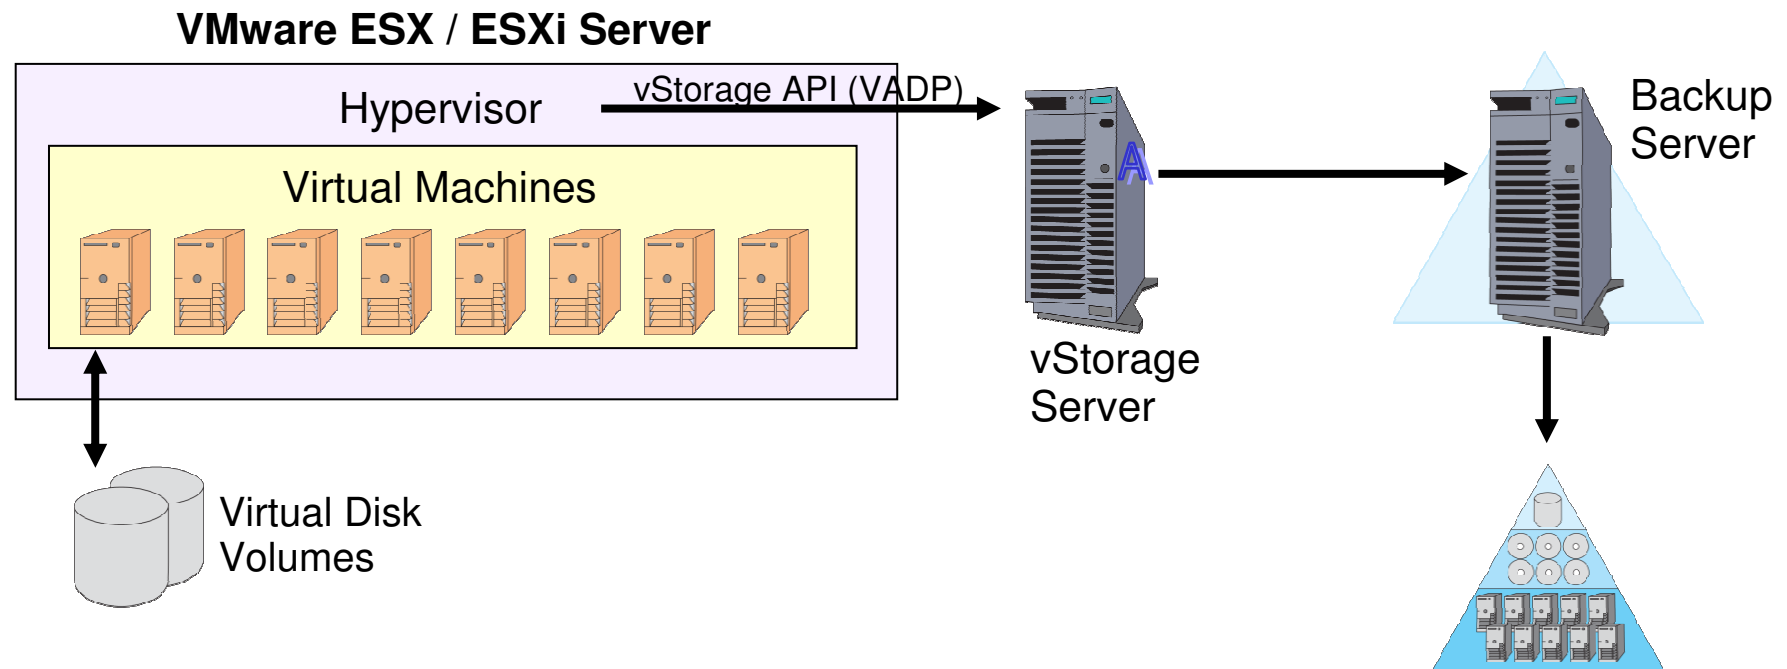
- Supported only on VMware
- This approach tries to combine the benefits of the traditional and host based approaches
- Provides an API to talk to the console and move the data through the proxy server

Pros	Cons
“LAN Free” backup	Questionable application integration (VMware triggers VSS for Windows guests)
Backup is off-loaded to the proxy server	VCB requires an additional data hop
Flexibility – supports both file level and image level backup.	Recovery might be challenging (depending on the type of backup used)
Utilizes VMware API	



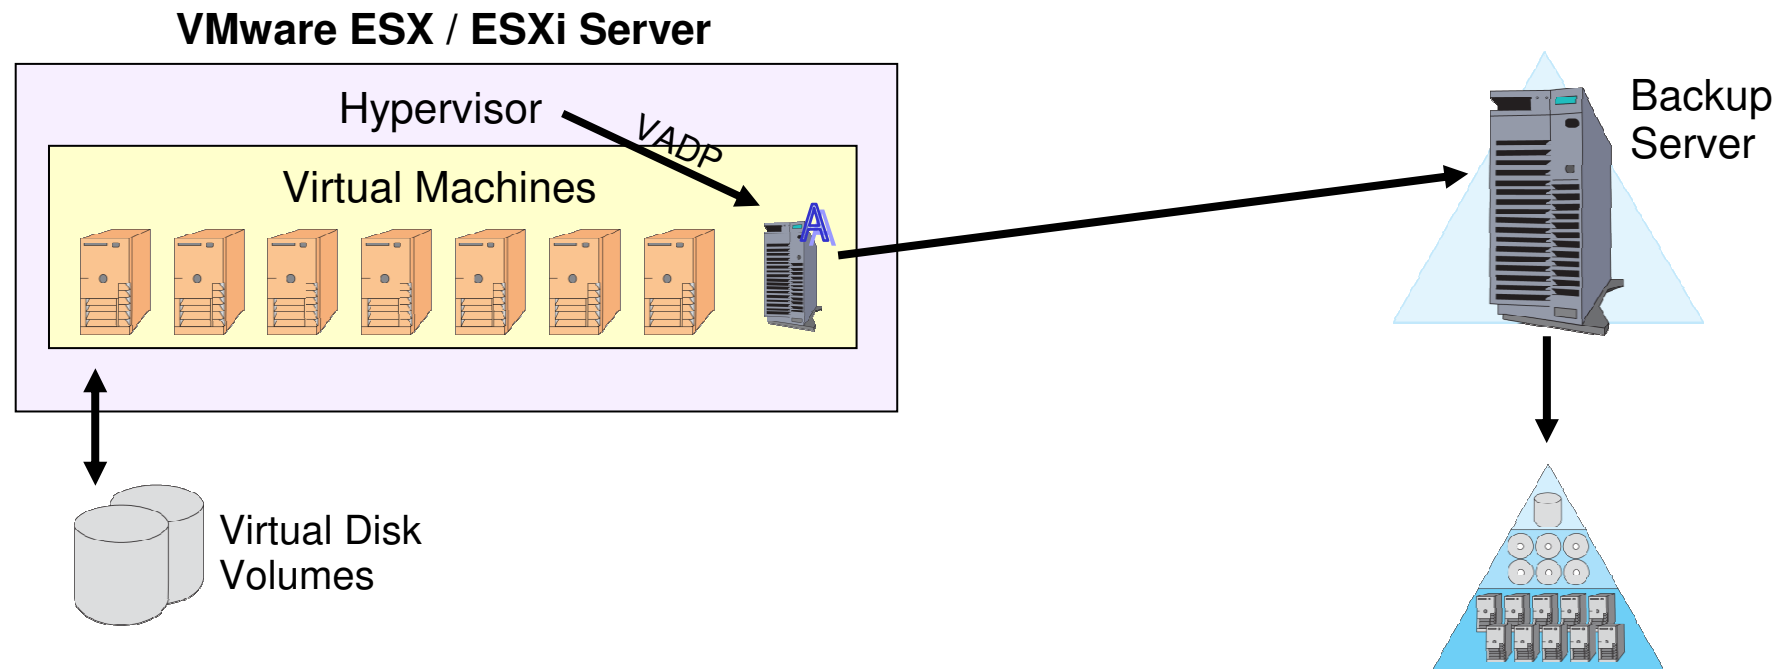
## The new approach: VMware vStorage APIs for Data Protection

- Data is accessed directly from the VM storage and passed directly to the backup server (single hop, data is not stored on the vStorage Server)
- Changed Block Tracking allows incremental backups (with periodic fulls) without forcing a scan of the guest OS file system



## The new approach: VMware vStorage APIs for Data Protection

- Data is accessed directly from the VM storage and passed directly to the backup server (single hop, data is not stored on the vStorage Server)
- Changed Block Tracking allows incremental backups (with periodic fulls) without forcing a scan of the guest OS file system
- **The vStorage Server can be a virtual machine – no additional HW needed**



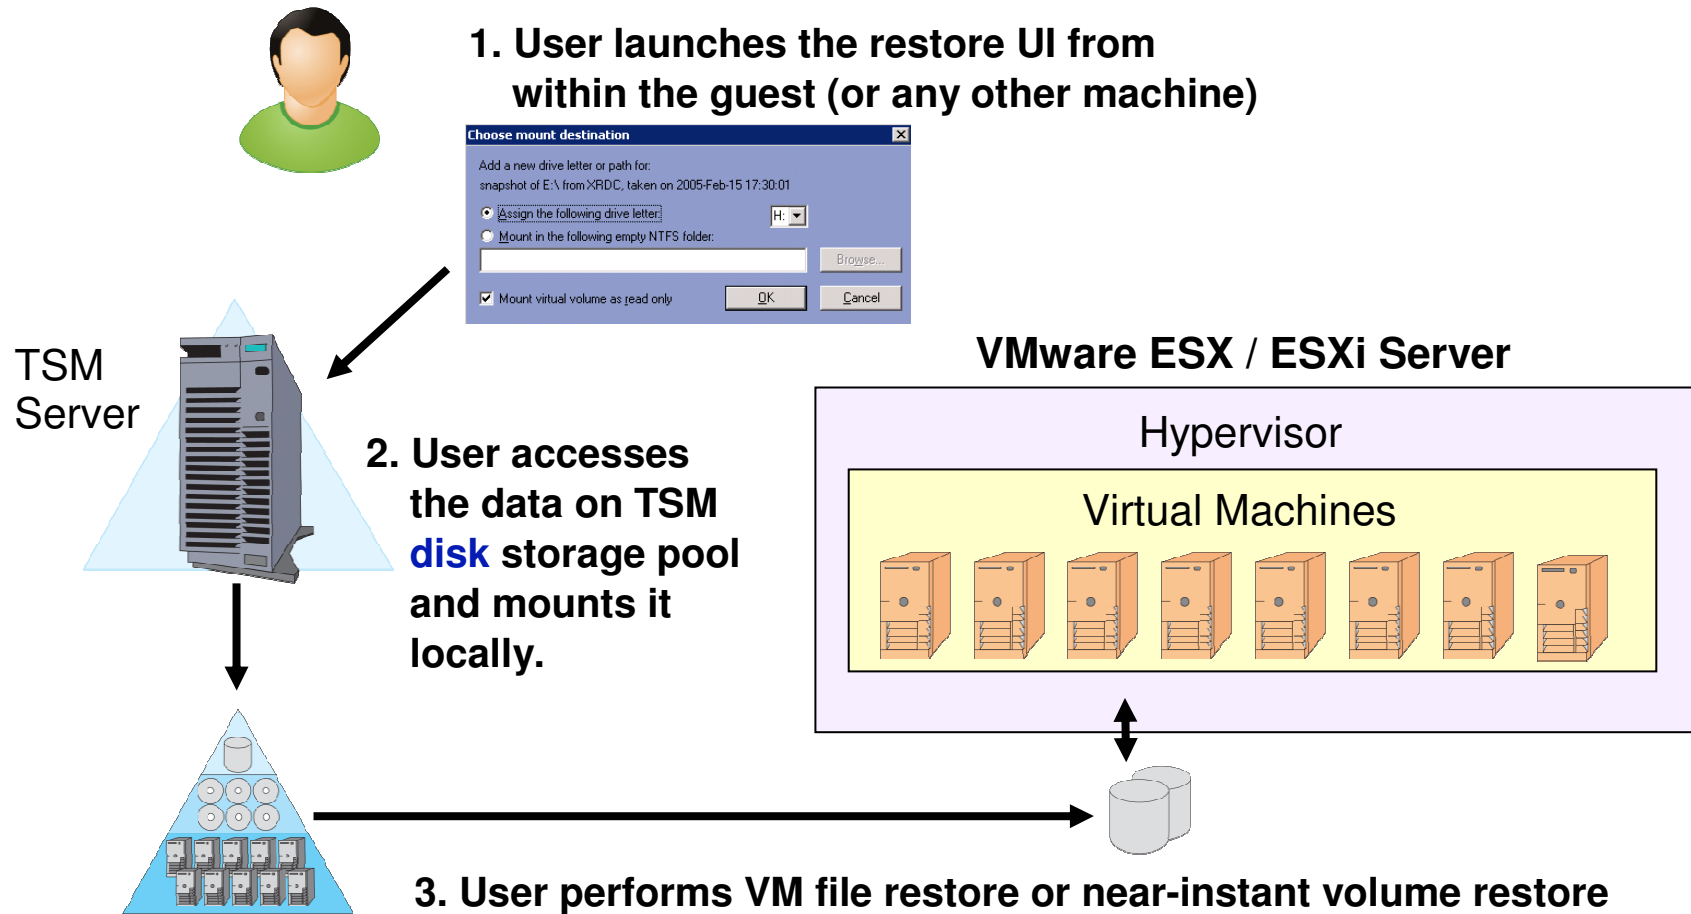
## Introducing Tivoli Storage Manager for Virtual Environments v6.2

- Announce – 22 Feb 2011; eGA – 13 Mar 2011
- Utilizes VMware's vStorage APIs for Data Protection, including block-level incremental backups based on VMware's Change Block Tracking
- **Offloads** the backup workload from virtual machines and production ESX hosts to vStorage backup servers
- Provides **flexible** recovery options – file, volume or image – from a single-pass backup
- **Near-instant restore** of Windows and Linux disk volumes
  - Data is made available immediately while it is copied in the background
- **Simplifies** day-to-day administration with the centralized Tivoli Storage Manager console
  - TSM for VE is supported on TSM Server v5.5 and above
- Automated discovery of new VMs; automatically applies backup policies



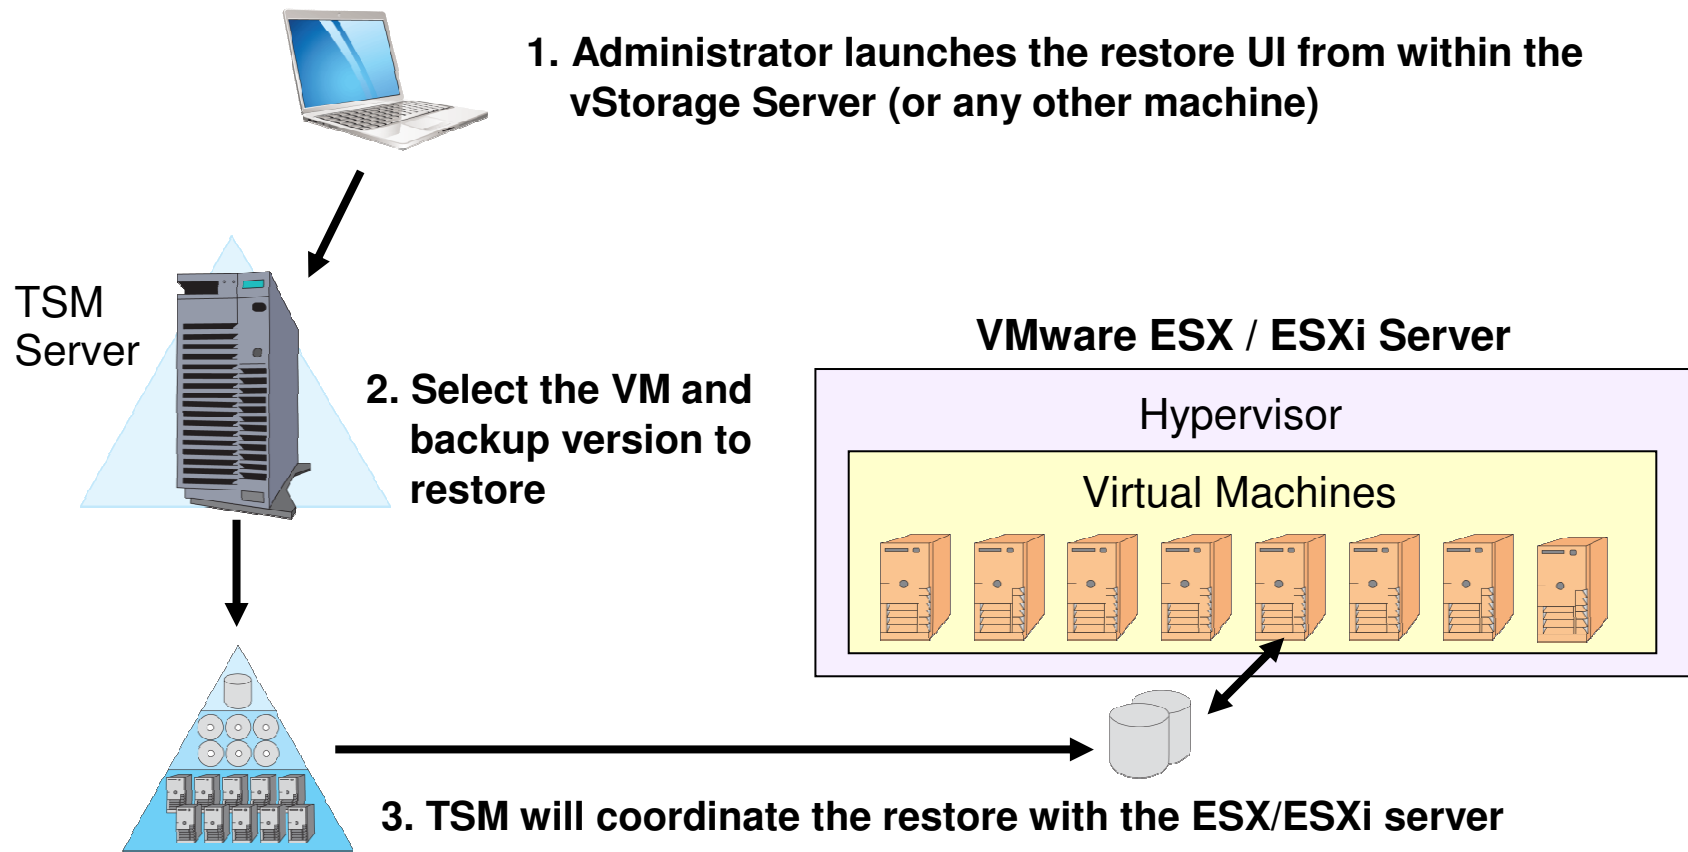
## Advanced capabilities, flexible recovery options

### Near-Instant File and Volume Recovery (Windows and Linux Guest OS)

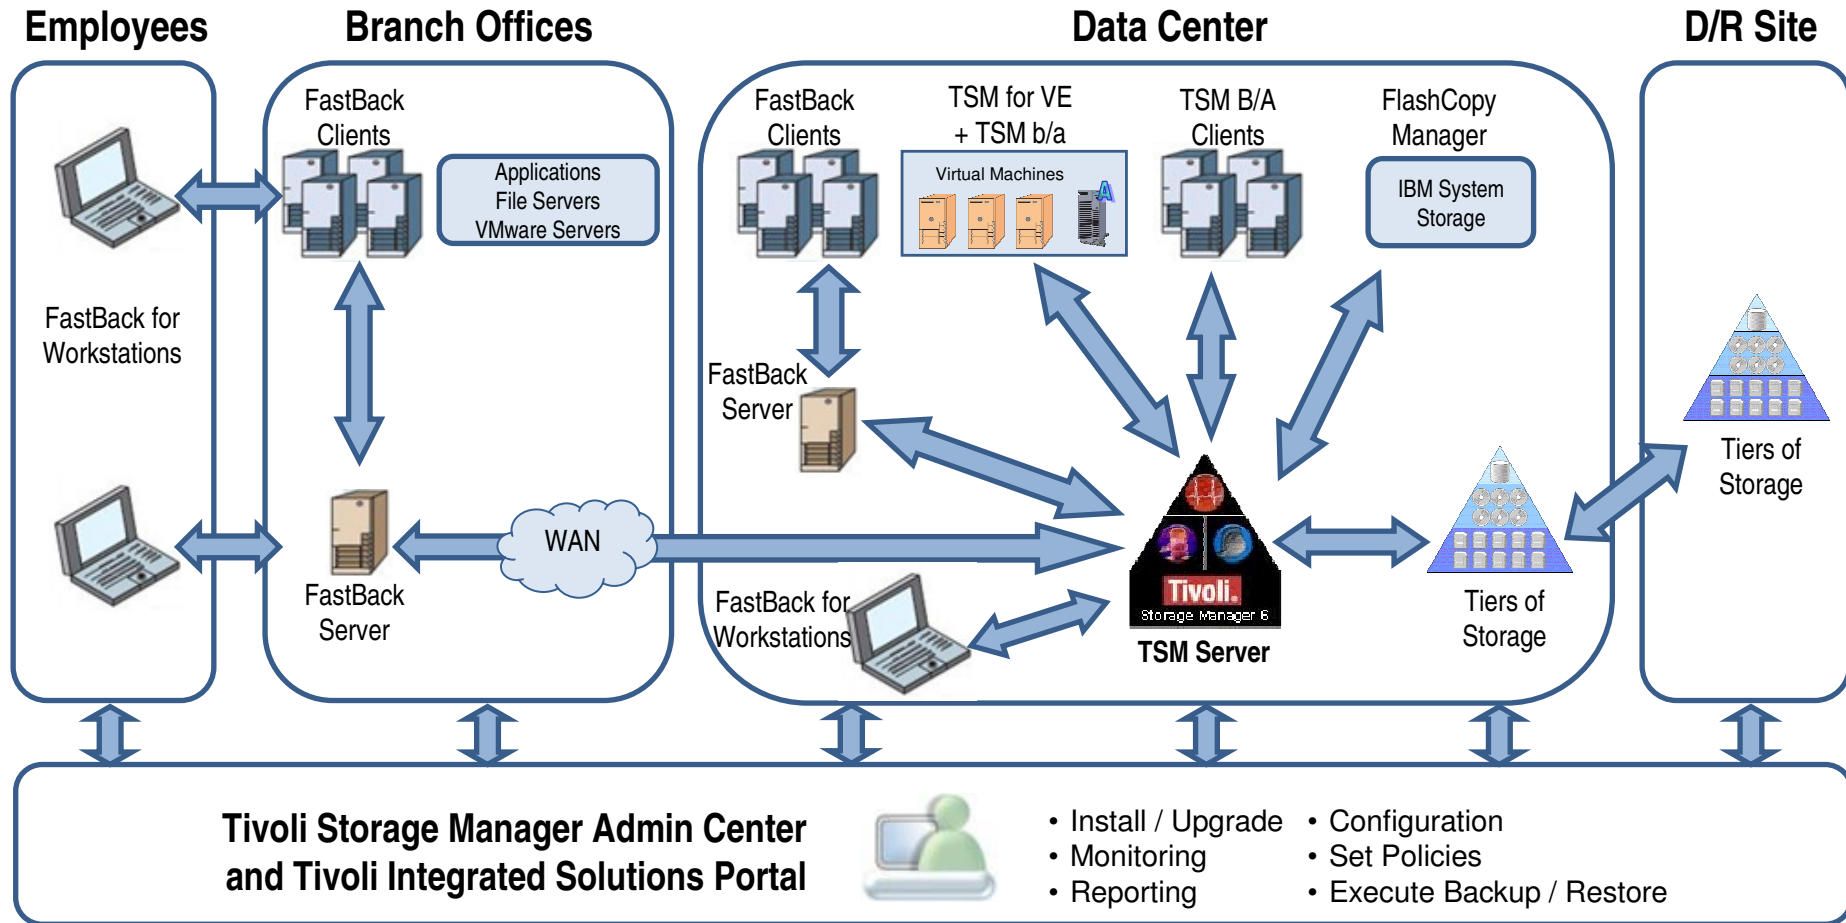


## Advanced capabilities, flexible recovery options

### Full Virtual Machine Recovery (any supported Guest OS)



# Continuing to Deliver Unified Recovery Management



## Tivoli Storage Manager for Virtual Environments - Summary

- Advanced data protection for VMware ESX and ESXi servers
- Leverages vStorage APIs for Data Protection (VADP)
- Non-disruptive, single-pass, block-level backup
- Flexible recovery options: file, volume, VM image
- Near-instant restore of files and disk volumes (Windows and Linux)
- No additional hardware required
- Simplified agent management – one agent supports multiple VMs
- Automated discovery of new VMs
- Support for LAN-free data transfer from the VMware server's storage to the backup server —preserving bandwidth for other uses
- Integrated with Tivoli Storage Manager for:
  - Unmatched scalability – manage up to 2 billion objects in a single TSM Server
  - Unified Recovery Management
  - Built-in data reduction / data deduplication
  - Policy-based tiered storage / data lifecycle management



## Why IBM

- IBM has the global resources to help solve the IT and service management problems of even the largest organizations, almost anywhere they do business
- IBM is on solid financial footing and will be here to support our customers for the long term
- IBM continues to invest heavily in research and development, providing customers with product roadmaps that will continue to add value to their relationship with IBM
- IBM offers the broadest range of products, services and financing options to help customers realize significant return on their investment





For more information  
*Links active in slideshow mode*

- [Storage Management Solutions](#)
- [Unified Recovery Management](#)
- [Backup and Recovery Solutions](#)
- [Remote Office Data Protection](#)
- [Data Reduction Solutions](#)
- [Storage Management in Cloud](#)
- [IBM Storage Blog](#)
- [Information Protection Services](#)
- [Tivoli Storage Manager for Virtual Environments](#)
- [Tivoli Storage Manager](#)
- [Tivoli Storage Manager FastBack](#)
- [Tivoli Storage FlashCopy Manager](#)
- [Tivoli Storage Productivity Center](#)
- [IBM Information Archive](#)
- [SAN Volume Controller](#)



**Khop Khun Mak**

**Cheers**

**Sukran**

**Tack IBM**

**Merci**

**Diolch**

**Gracias**

**Arigato**

**Kamsa hamaida**

**Danke**

**Salamat**

**kiitos**

**Efharisto**

**Thank You**

**Dankie**

**Grazie**

**Hvala**

**Xie xie**

**Kam ouen**

**Shukria**

**Tesekkurler**

**Köszönöm**

**Toda**

**Takk**

**Spasiba**

**Smarter Storage Management**

## Trademarks and disclaimers

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries./ Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both. IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce. ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office. UNIX is a registered trademark of The Open Group in the United States and other countries. Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others. Information is provided "AS IS" without warranty of any kind.

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

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.

Photographs shown may be engineering prototypes. Changes may be incorporated in production models.


© IBM Corporation 1994-2010. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Trademarks of International Business Machines Corporation in the United States, other countries, or both can be found on the World Wide Web at <http://www.ibm.com/legal/copytrade.shtml>.



Thank you!



Smarter software for a smarter planet

**Smarter software for a smarter planet**

Smarter software IBM smarter

**Smarter Storage Management**