



**Pulse2011**



# A Hitchhiker's Guide to Storage Clouds

What Storage Professionals Need To Know About Cloud


Trevor Foley  
Storage Solutions Specialist

## Abstract

- Cloud computing is all the rage in IT, promising to streamline services and reduce costs. But cloud raises many questions for storage pros. Is the data in the cloud adequately protected? Can or should you use public cloud services for backup and disaster recovery? How do you build your own storage cloud?
- Virtual server and storage technologies, the enabling technologies for Cloud, have dramatically changed the data centre, reducing hardware, energy and facilities costs. However, these solutions have made data and storage management exponentially more complex, creating new questions: where is the data, how do you manage it, how do you best utilize your existing storage resources?




## Storage Cloud Adoption



**Why Cloud:** An overview of the Cloud technology adoption model shows we have just completed the “early adopter” stage. Market share winners and losers will be determined as the early majority adopts Cloud solutions.

Storage Cloud Market Oppty



Sources: 1) IBM Market Insights Assessment, with IDC, Gartner and other data source inputs and IPR consultation through August 2010  
 2) STG Data Center Analysis GMV 1H11




# What is a cloud?

- Access to applications or IT resources from anywhere through connected devices and a simplified UI
- Applications reside in massively-scalable data centers
- Compute resources automatically provisioned and shared to achieve significant economies of scale
- 'Pay-as-you-go' usage model
  - Better predict and manage expenses
  - Simplify operations
  - Turn CapEx into OpEx

## Common Examples of Cloud Services

- Sales Automation: Salesforce.com
- Marketing Automation: Constant Contact
- Smartphone Apps



## Why cloud computing, and why now?

- Global economy
- Server sprawl
- Low utilization
- Rising operational costs
- Ballooning energy costs and demands
- Increasing system and network availability requirements
- Environmental compliance requirements
- Government mandates
- Heightened consumer demand for greater social responsibility
- Increasing consumer demand for more information and richer user experiences



## Common attributes of cloud computing

- Flexible pricing – Fit to needs
- Elastic scaling – One size does not fit all
- Rapid provisioning – Greater quality control
- Advanced virtualization
- A better user experience



## Private cloud

- Owned and managed by the enterprise
- Limits access to enterprise and partner network
- Drives efficiency, standardization and best practices
- Retains high degree of control, privacy and security
- Enables business to more easily customize services
- Reduces deployment time for new services
- Accessed from "inside" the firewall



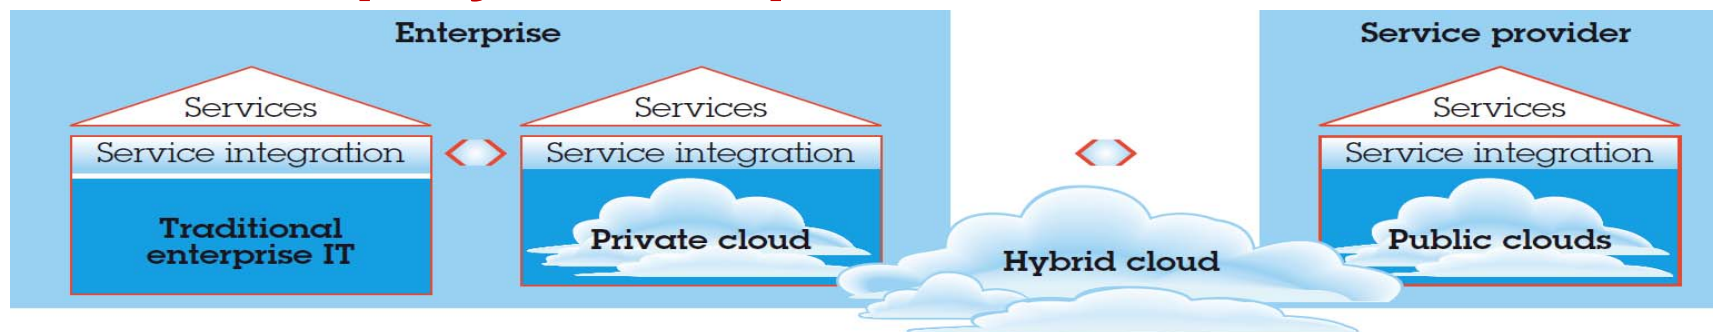
## Public cloud

- Owned and managed by service provider
- Limits access to subscribers
- Delivers select set business process, application or infrastructure services on a “pay per use” basis
- Highly standardized
- Limited customization options
- Accessed from "outside" the firewall





# Cloud Deployment Options




- Build your own private cloud
- Subscribe to a private cloud
- Subscribe to a public cloud
- All of the above (hybrid cloud)?



# Why should storage administrators care about cloud computing?

- Need a smarter way to address the growing gap between data and budget
- By 2014 more than 50% of all new storage capacity will be deployed in a public cloud\*
- 80% of new enterprise applications will be written for use in a cloud, at least as an option\*



\* Source: IDC Directions2011, March 2011



# Storage Management in Public Compute Clouds

- **Good news:**
  - You don't have to worry about storage management
  - The cloud vendor handles it all
- **Bad news:**
  - You do have to worry about data management
  - Do you have enough bandwidth to handle the added traffic?
  - Is the vendor making multiple copies for backup and DR?
  - What are the data access, retention, and recovery SLAs?
  - Is the data secure and encrypted? Who has the keys?
- **Bottom line:** Ensure your vendor is using a scalable, enterprise-class data protection and recovery solution
  - TSM can protect up to **2 billion** data objects in a single server
  - Built-in data deduplication reduces cloud service costs




## Storage Management in Private Compute Clouds

- You'll need to manage the storage systems and the data, but in a more virtualized environment
- Virtualized servers (e.g. VMware) are the enabling technology for cloud
  - They provide the elasticity and fast provisioning required by cloud computing
- Virtual servers require advanced solutions for data protection
  - Need to ensure new VMs are protected, but need to avoid “agent sprawl”
  - Traditional in-guest backup can result in severe performance problems
- Virtualized storage and centralized storage management are also key enablers



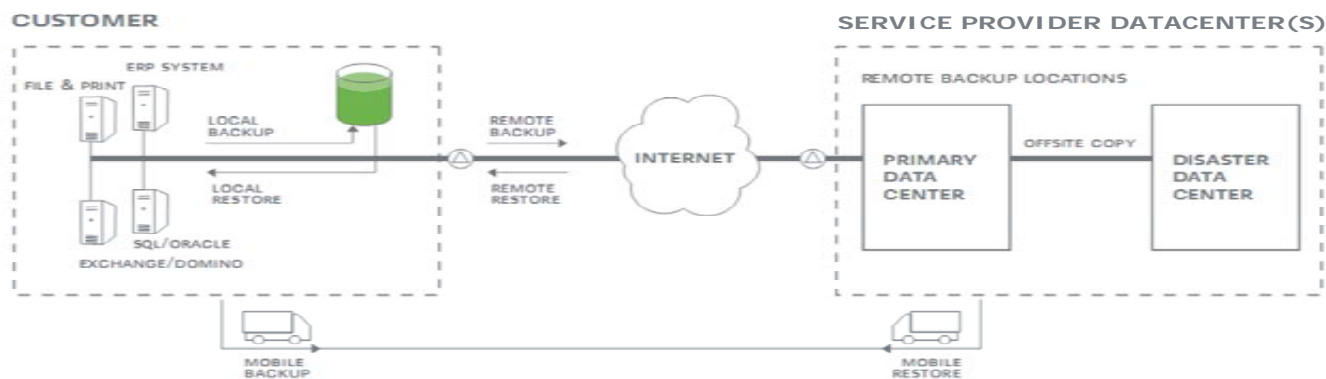
## Backup-as-a-Service Clouds

- Ideal for small-to-midsize companies and for remote / branch offices
  - Eliminates the need for some on-site technical staff
  - Pay per capacity on a monthly basis
  - Enables predictable expense control
- Considerations:
  - Support RPO and RTO for each data type?
  - Are you OK giving your data to a 3<sup>rd</sup> party?
  - Are they making multiple copies in different locations (as you would)?
  - How do you get your data back if the vendor goes bankrupt?



# Backup-as-a-Service Clouds

- Many backup-as-a-service options are available
  - IBM Information Protection Services: click [here](#) in slideshow mode
- Many use Tivoli Storage Manager as the backbone of their offerings
  - Scalability, functionality, performance, resilience, reliability
- See [www.front-safe.com](http://www.front-safe.com) and [levitar.com.au/services/cloud/](http://levitar.com.au/services/cloud/)



*"IBM delivers the technology; we deliver the business model; and our resellers can go to market with the very best backup solution there is"*  
-- Peter M. Nielsen, CEO, Front-safe A/S



## Case Study: Global Data Vault

- Challenge
  - Help clients decrease the time and cost of recovering from service outages and disasters
- Solution
  - A fully-managed secure backup service based on IBM Tivoli® Storage Manager family software
- Benefits
  - Enabled delivery of enterprise level business continuity services to small and medium-sized businesses; reduced client recovery times from days to under 30 minutes; improved governance


“The Software as a Service model provides companies with the staffing and technology needed to help ensure business continuity in the event of an outage or data loss. And working with IBM, we can deliver an enterprise-class solution that was previously out of reach for most small and medium sized businesses.” -- *Jeffrey Beallor, President, Global Data Vault, Toronto*





## Extending In-House Backup to the Cloud

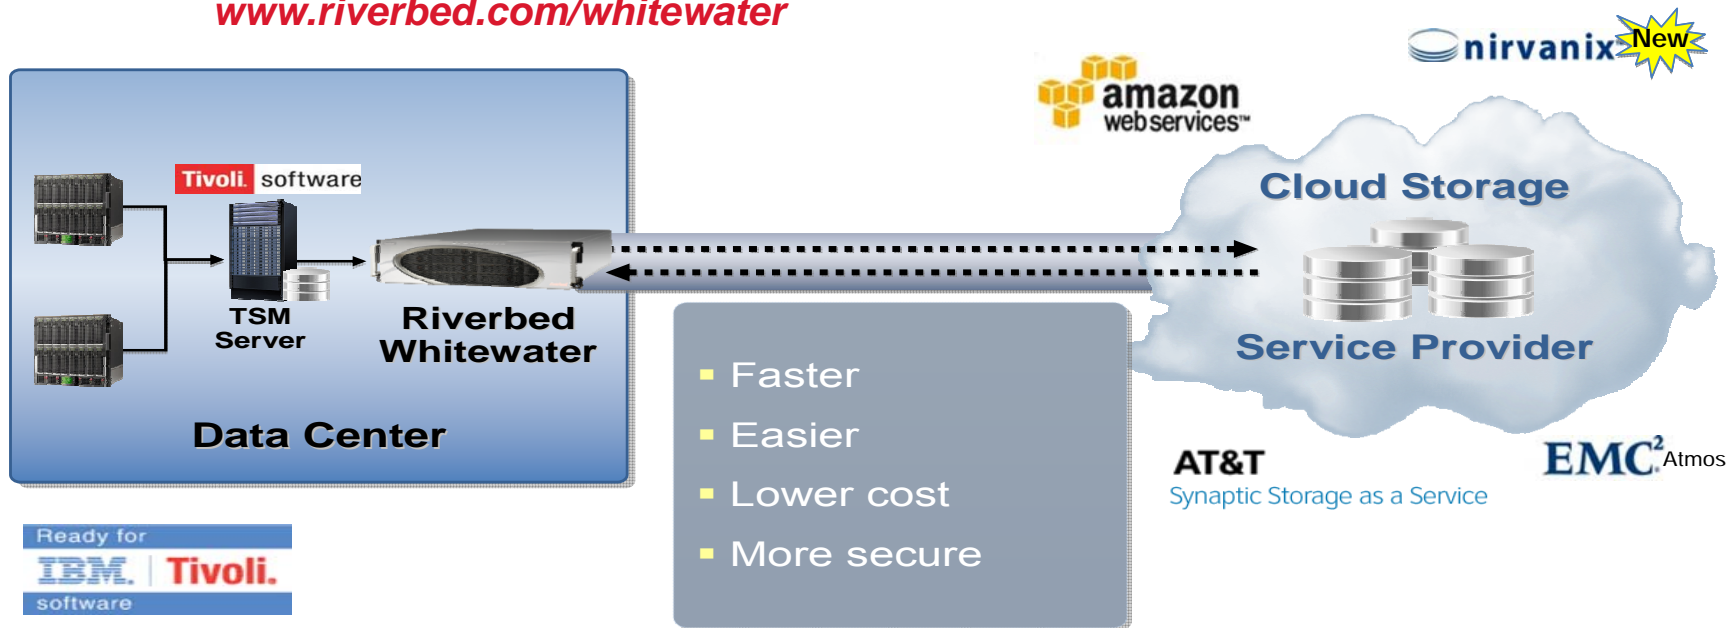
- Use public clouds as the *n*th storage tier –
  - For disaster recovery copies
  - For long-term retention and archive --- hidden benefits!!
  - Or for any other purpose supported by the available bandwidth and vendor SLAs
- Same concerns as mentioned earlier:
  - Bandwidth to handle backups *and* restores
  - Data security: access control and encryption
  - Service resilience, reliability, scalability, longevity
- Probably need a Cloud Service Gateway
  - Eliminate the ‘chattiness’ of IP-based networks
  - Locally-controlled security
  - Cache frequently-accessed data locally





# Cloud Accelerator: Riverbed Whitewater

[www.riverbed.com/whitewater](http://www.riverbed.com/whitewater)



# Server Virtualization

- Virtual Machines (VMs) are a primary technology building block of Cloud infrastructures



- They are also making their way into production application environments
  - The overall cost savings are very compelling
  - Many companies require strong justification for NOT using a VM
- But what about Storage Management?



## Virtual servers have changed the game ...



“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 [\(link\)](#)

Virtual Machines are easy to create, move and decommission.

*But what happens to the data and the storage capacity assigned to these VMs?*

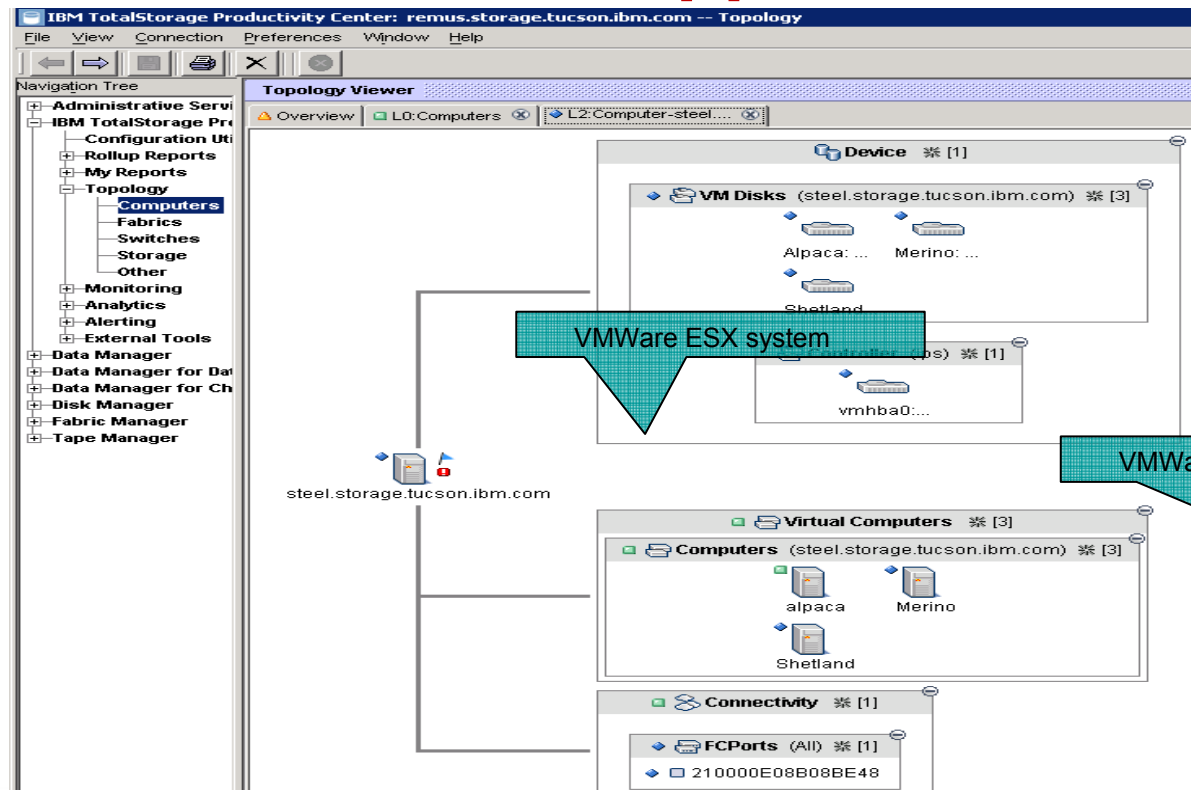


## Tivoli Storage Productivity Center

- Simplifies storage deployment, configuration and management
- Helps optimize storage performance & space utilization
- Centralizes end-to-end management of storage and SAN fabric
- Helps automate storage tasks associated with provisioning and replication
- Frees your IT staff to focus on mission critical tasks

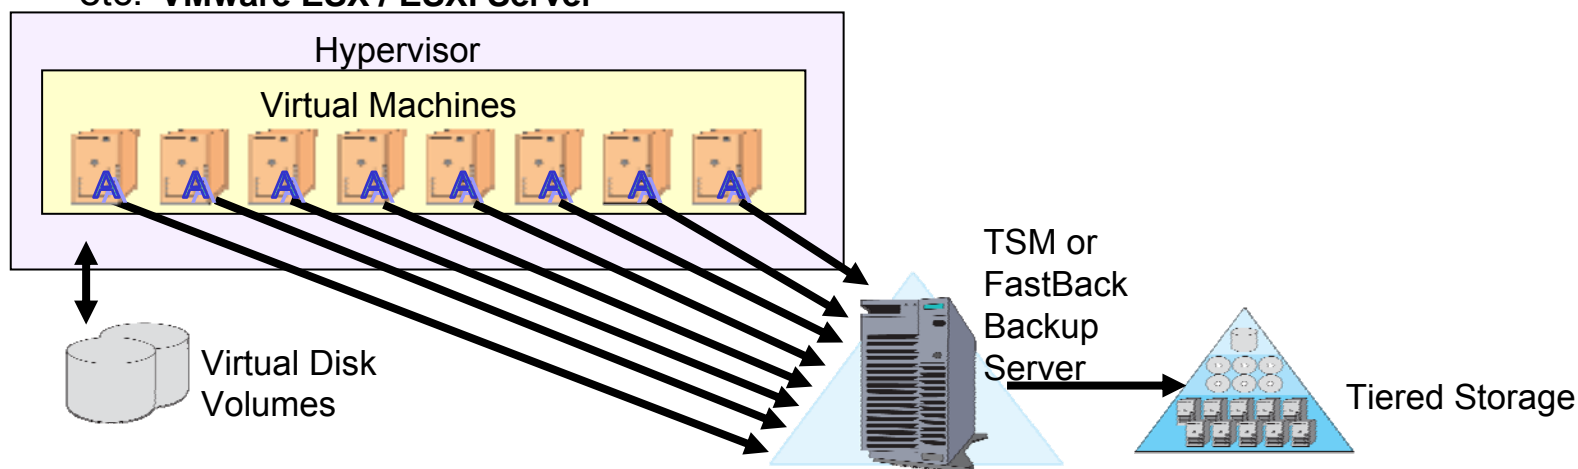


# TPC for Data VMWare Support



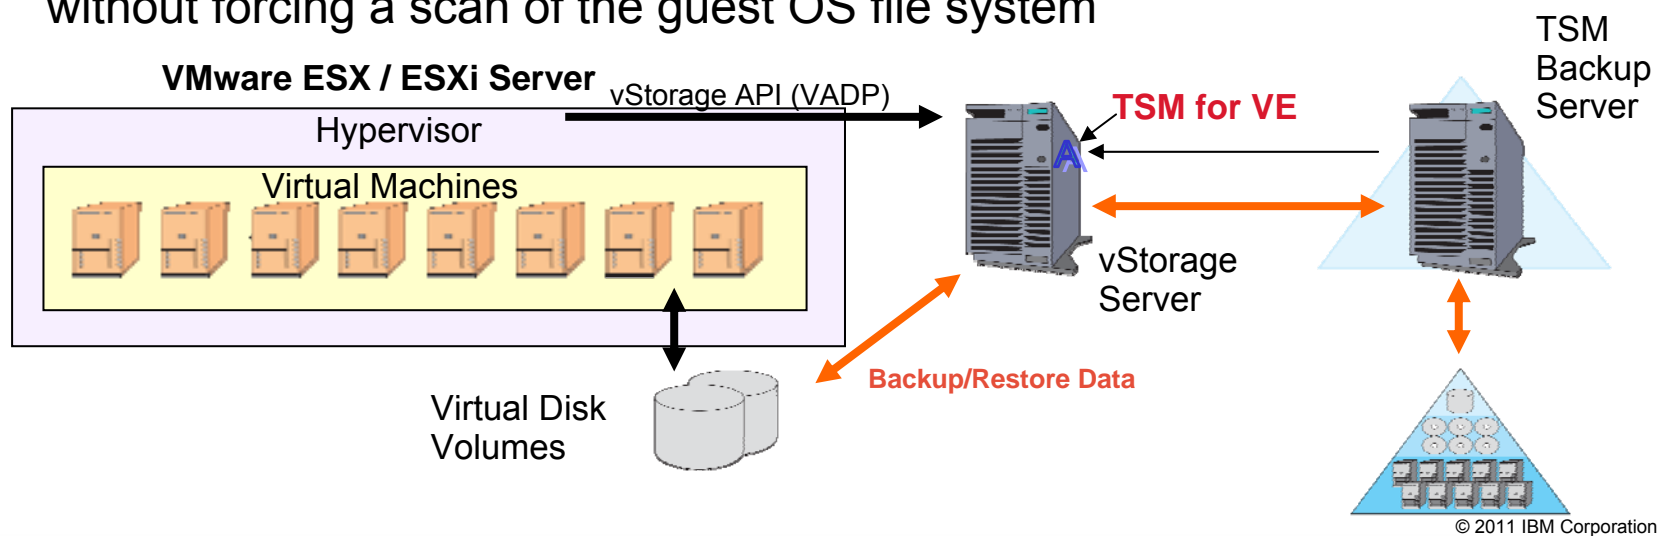
## Traditional in-guest backup and recovery model

- 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. **VMware ESX / ESXi Server**




## 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



# Questions?





## 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



## Trademarks and disclaimers

© Copyright IBM Australia Limited 2011 ABN 79 000 024 733 © Copyright IBM Corporation 2011 All Rights Reserved. TRADEMARKS: IBM, the IBM logos, ibm.com, Smarter Planet and the planet icon are trademarks of IBM Corp registered in many jurisdictions worldwide. Other company, product and services marks may be trademarks or services marks of others. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)

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.

