

**Remarks by Enrique Salem
President and CEO, Symantec Corporation
Storage Networking World Conference
Orlando, Florida
Remarks As Prepared – April 7, 2009**

Almost exactly 200 years ago, the journey of Lewis and Clark came to an end where it had begun, in St. Louis, Missouri.

When most people talk about that voyage, they talk about the journey up the river: the discoveries they made, and the tribes, plants and wildlife they encountered.

What I think is even more interesting – and relevant – for us today is the part people don't pay as much attention to: what happened on the way *down* the river, on the way back.

First off, while hunting for food, Captain Lewis was shot in the backside by one of his own men. It was an accident, but I'm sure there's a lesson in that for a CEO three days into his new job.

Secondly, by the time the group was within a couple of hundred miles of St. Louis, they were completely out of food and goods to trade. If they didn't make it back, the whole journey could have been a failure.

That's when they ran into two fur trappers hoping to find a place called Yellowstone. Lewis needed supplies, but had nothing to trade. And then he realized he had something of great value. He drew the traders a map of what they could expect upriver.

He wasn't carrying a lot in his hands, but he was carrying a lot in his head. What he had was information: information that took him two and a half years to gather. That information is what allowed the journey of Lewis and Clark to be a success. And 200 years later, it's still all about information.

Information is often the most valuable thing you possess.

That's why we're all here today.

My journey hasn't been quite as epic as Lewis and Clark's. But it, too, has brought me full circle.

I come from a storage background. I used to write software for removable media. I worked on software for QIC-40 and QIC-80 tapes for one of the early Symantec backup products. Maybe some of the older folks in the room have fond memories of when 80 megs on a quarter-inch-cartridge seemed pretty cutting edge.

I began my technology career as the eighth software developer at Peter Norton Computing. I was acquired by Symantec in 1990 for the first time, and I stayed here for nine years, holding multiple roles. I was the first chief technology officer at the company and I also ran the security business.

I left to work for some other companies, and in 2004, the company I was working for – Brightmail – was also acquired by Symantec. I became known as the “Boomerang Baby” – I kept ending up back with Symantec.

That’s when I figured that if you can’t beat ’em, join ’em. And if you have to keep joining ’em, you might as well lead ’em.

Throughout my career I’ve had the opportunity to see the technology and the business challenges evolve.

I remember when the Michelangelo virus hit in 1992 spreading through the shipment of infected computer components. At more than a dozen companies worldwide data was wiped out doing thousands of dollars worth of damage. Even back then we could see how critical it was to secure and manage our information.

Now, more than 15 years later that idea couldn’t be more true.

Today, the average medium-to-large enterprise experiences data growth rates of 50 percent a year. That means that every two years, the amount of information you need to secure and manage more than doubles.

A lot of this information comes from applications – and four of every ten data center managers report having more than 1,000 applications they have to deal with. Half of which are mission critical.

There’s also the unstructured data – email, spreadsheets, and instant messages – that needs to be managed. Many of our customers say this unstructured data is a major driver of rising IT storage costs.

More and more information resides outside of your organization – in software-as-a-service applications or with partners.

While information continues to grow, so do the risks to information.

Security threats today are designed to steal confidential information. Attackers aren’t going for glory. They’re going after information with goal of selling it in the underground economy.

Last year, more than 650 documented breaches exposed more than 35.5 million records. Not only do these breaches cost an average of \$200 per lost customer record – they often trigger audits and additional reporting requirements. Some of these breaches resulted from something as simple as a lost backup tape.

You’re also dealing with an increasing number of data retention requirements – whether associated with external regulations or internal business practices.

And, you often need to find information quickly in the event of a legal issue. A recent survey we conducted with CIOs and IT managers found that the majority of those surveyed felt confident in their overall retention practices. That makes sense since a lot of these CIOs are saving everything many times over.

But there's a gap, because only 15% were confident enough to "bet their pay check" that they could indeed produce all relevant information for a legal discovery event. Given that more than eight out of ten of you will be involved in at least one legal discovery effort – and many of you will be involved in significantly more – that's not the betting line you want.

So take a step back and look at the overall picture. We're in an environment where everything seems to be falling – budgets, staffing levels, acceptable downtimes.

But, the amount of information you have to store is increasing.

And the cost of storing it is increasing.

Energy costs in a typical data center double every five years. Some estimate that the cost of energy is the second largest operational expense for the data center. Today, many of you will spend 50 percent of your hardware investment on energy and cooling for your data center.

In that sense, the cost isn't just to your bottom line, there is an environmental cost associated with using too much storage.

So when you look at the array of challenges, I bet some of you feel like the merchant in the ancient story. This merchant went to the oracle and said his abacus counters couldn't keep up with the workload -- but he couldn't afford to hire any new workers. What should he do?

"Each abacus counter must grow another finger on each hand," said the oracle.

"That's very wise," said the merchant. "But how do I get them to do that?"

"Ah," said the oracle. "I only make policy. Implementing it is your job."

Substitute "consultant" for "oracle" and I'm sure you've dealt with a similar scenario.

Implementing is your job. But you're not alone. Everyone in this room is in the same boat. And we're designing solutions that can help.

In this economic climate, when one third of CIO's are expecting their IT budgets to remain flat, and many are expecting to see them reduced, the most important measure of your success isn't the ROI.

It's ROY – Return on Yesterday.

How can you take what you bought yesterday, and get more out of it today.

During the good times, most organizations added more storage as file systems grew and new applications came online. This tactical, siloed approach to buying storage led to too much capacity. As a result, the average storage utilization these days is only 35 percent.

You often have more room to grow than you realize. We had one customer who told us that they were managing massive information growth. They estimated they were using 50% of their storage. We showed them they were only using 20%.

That leads to my key message for you today: Stop buying storage. You heard me right. Stop buying storage.

How? Reduce the amount of information that you store in the first place, and make the storage you do have more flexible and efficient.

If you want to stop buying storage, there are four key opportunities that we think everyone should be looking at right now.

Storage resource management.

Thin provisioning.

Data deduplication.

And, intelligent archiving.

Let me say a bit about each:

First, storage resource management.

If you're like most companies, you have a lot of "orphan storage." It's time to give your orphan storage a loving family.

In many companies there are differences in storage hardware, and often islands of storage. One department might have plenty of free storage while another is adding arrays.

You need to identify and reclaim what you've bought but aren't using. Find that orphan storage, and bring it home.

The hardware vendors will tell you they can show you how your existing storage is being used. Remember, their ultimate goal is to sell you more hardware.

And the truth is they can only tell you how their storage is being used. If you're using more than one hardware vendor, and I'm guessing a lot of you are, you'd have to look across all of your vendors to get a complete picture.

Even then, how much storage is in an array and how much has been allocated to a server is only half the story. You need information not only on a single storage array, but across all of your data centers.

What you need are storage management solutions that work across different platforms and give you an end-to-end view of the storage you have, the storage you're using, and the storage you're not.

You need a better way to peek into your systems so you can make informed decisions about how your existing storage assets are being utilized. Storage resource management technology can help identify problem areas and consolidation opportunities and create a priority list of solutions.

How does managing storage utilization help you stop buying storage?

Let me tell you about our work with a large financial services company. When we first started working with them in 2006 they were using 12 percent of their storage. We worked with them on a zero growth storage initiative, and, two years later, they are at 41 percent.

That's progress. But the real measure of progress is \$70 million. By deferring new storage purchases, we helped them save \$70 million.

If you want to see a happy CFO in a tough economy, tell them you've found a way to save them millions of dollars.

Second, make use of a technology the market didn't have in its lexicon just a few years ago: thin provisioning.

So, at risk of sounding like a diet guru: it's time to get your storage fit and thin!

In traditional systems, storage that is dedicated in the storage array is allocated to a server. That server, in turn, makes that storage available to the application. This is one factor that has contributed to the low storage utilization in data centers.

Every day, you're getting requests to add new applications, with each new app demanding more storage than it will likely need to use.

What if you could allocate storage on demand?

Thin provisioning allows you to allocate storage as its needed, instead of putting it all in place on day one.

But the challenge with today's thin provisioning arrays is that they aren't intelligent enough about what the server and applications are actually using.

Symantec is leading the industry by providing the only solution on the market that enables you optimize thin provisioning to "get thin."

Our Storage Foundation SmartMove functionality enables you to migrate online to thin storage – without any disruption – and automatically reclaim unused storage during the process.

As any dieter knows, getting thin is only half the challenge. The other half is staying thin. We're working with the hardware storage vendors to help you keep your storage thin over time. Our Thin Reclamation API makes sure that the storage knows when information is deleted. That way unused storage can be added back to the pool and utilization rates remain high.

So, how does thin provisioning help you stop buying storage?

It's a simple concept – you buy the storage you need and can save money from over provisioning storage.

Third, get serious about deduplication.

Eight is Enough may have been a great TV show. But eight is seven copies too many when it comes to storage.

How many copies of any given database or file do you have in your data center? In 2006, InfoStor claimed it could be as high as 25 copies. In practice, we've seen customers storing 500 or more copies of data.

As organizations have been migrating from tape to disk-based backup, deduplication technology is becoming a more accepted and popular tool to streamline the backup process. Deduplication technology improves backup performance and reduces both the storage and bandwidth consumed by backups.

If you want to see why deduplication is so important, look the example of a simple PowerPoint presentation. Let's say three people have the same 5 MB presentation on their computer.

This is a conservative example, because, really, when's the last time you've seen a presentation only emailed to three people?

If you perform a full backup once a week in three weeks alone you'd have 9 versions of the same presentation backed up. In a year you'd have 156 copies and in 10 years you'd have 1,560 copies.

One copy is enough.

But deduplication goes even further than that.

Each one of those slides in the 5MB PowerPoint contains the company logo. If you assume there are 20 slides in the presentation, then you just backed up the company logo 31,200 times.

Try justifying the storage purchase for that one.

Our deduplication technology is intelligent enough to look inside a file and know to store only one company logo. Because one copy is enough.

Remote offices can deduplicate the data at the source before sending it to the data center. In some cases it can reduce the bandwidth needed to move backup data by as much as 99 percent. This improves disaster recovery, enabling remote offices that once thought they had too much data to back it up over the network.

We also provide deduplication in the data center. That allows you to protect all your applications and servers – including physical and virtual ones – as you do now, and deduplicate the data in the back-end storage.

The real differentiator for Symantec is this: we can manage a single pool that deduplicates data across the entire network. That means, for example, we can deduplicate at a remote office to make the data set smaller before moving it across the network to the data center, and then, deduplicate it again to make sure that any duplicate data between the remote office and the data center pool is as efficient as possible.

Again, one copy is enough.

So, how does deduplication help you stop buying storage?

We've been working with a bank with 52 remote branches and have helped them reduce storage costs by 25 percent. Since the backup process is now automated they've been able to improve productivity of the storage team by 50 percent.

And they're seeing what we all know: save space, and you're also saving time and money.

Fourth and finally, archiving needs to become more intelligent.

Today most organizations are thinking about how to address three common management challenges:

How can you reduce the cost of storing and managing unstructured information, like email?

How can you automate and enforce consistent classification, retention and deletion policies – in other words, how can you keep what you need and delete the rest?

How can you quickly retrieve and preserve data for litigation and regulatory investigations?

It's an age old challenge – you have to separate the wheat from the chaff.

The key is to store less: not just by deduplicating, but also by deleting information that has reached the end of its useful life. Yes, you need to abide by legal hold requirements, but you also need to abandon the “keep everything forever” mantra.

Already, we're seeing companies automatically delete information that has reached the end of its useful life using retention folders in the archive. One of our customers retires two terabytes of data each day with this approach.

Intelligent archiving solutions are making the concept of tiered storage a reality, enabling organizations to move older less-frequently used or less important files from high-cost disk to lower cost storage.

With content awareness technology, email can be given a unique retention category in the archive. Based on the content it can be retained and migrated across tiered storage, including tape, when appropriate.

How does archiving help you stop buying storage?

One law firm we work with has seen dramatic results. They reduced the amount of data in tier-one storage by 70 percent and reduced total storage costs by 40 percent.

I began by talking about Lewis and Clark exploring – and I'll end by talking about skiing.

I love skiing, but I picked it up pretty late in life. I was never one of those four year-olds barreling down the slope with perfect form, too young to know what it's like to be sore the next day when you fall.

So when I was learning how to ski, and I started to lose control a little, I'd give into my first instinct. I'd lean back. I felt like leaning back would put me further away from everything I was trying to avoid, and give me some control.

I quickly learned that if you want greater control, you actually have to lean forward. That's what lets you dig in your edges and take charge.

The same is true for all of you today. There are challenges coming at you from all angles. The worst thing you can do is lean back.

But if you lean forward.

If you stop buying storage and instead: Embrace storage resource management solutions, Commit to thin provisioning, Actively employ deduplication, Start using intelligent archiving.

You'll not only save money and time, you'll not only reduce your company's environmental footprint, you'll have taken your challenges head on, controlled them, and overcome them.

Thank you.

###