

# A Vision of NetApp as a \$2.5 Billion Company

Dave Hitz  
November 2, 2003

## Introduction

NetApp's ultimate goal is to drive the reorganization of the storage market from a vertical structure, dominated by computer systems vendors, to a horizontal structure, dominated by focused storage vendors, and to have the #1 market share in that restructured market. This paper is about the next step towards that goal, which is to achieve \$2.5B in revenue by FY2007 or sooner.

NetApp has been wildly successful. We created a new market for NAS appliances, and we are the market share leader. We've taken NAS from a low-end solution for engineering workgroups to an enterprise solution for mission critical Oracle databases.

And yet, I worry that our success carries the seeds of our own downfall. NAS used to be a separate and protected niche within storage, but we have made NAS so mainstream that every storage vendor now wants their piece. NAS and SAN used to be quite distinct, but we developed such good enterprise NAS devices that we now compete head-to-head against SAN for much of our business. The convergence of Ethernet and Fibre channel performance and reliability is a foundation for these other trends.

In summary, I worry that these trends could transform us from the secure market leader of a \$1.5B "NAS Storage Market" into – three years down the road – a small and vulnerable player in a much more unified \$15B Open System External Storage market.

Nevertheless, I am optimistic.

Despite my (healthy?) paranoia, I believe that the convergence of SAN, NAS, DAS, Fibre Channel and Ethernet creates great opportunity. Market restructurings and economic downturns often propel newcomers into market leadership, and right now we have both. Also, I believe these trends will take time to develop. We have time to act. Our entry into the SAN market and our focus on large enterprise accounts are exactly the right first steps.

My goal in the rest of this paper is to describe, as clearly as possible, the company that I think NetApp needs to be by FY2007:

In “**Market Position and Growth Requirements**,” I’ll use IDC data to paint a picture of our current position in the market, and how I believe it needs to change over the next three years. In summary, I believe that today in the overall external storage market there are “big players” and a pack of question marks. We are currently a question mark, but we can become a big player if we grow to \$2.5B by FY2007.

Reach \$2.5B by FY2007! Are we crazy? In “**Sanity Check**,” I’ll explain why I think growth rates like this are sane to consider. One reassuring observation is that when the dot-com crash hit, only 30% of our revenue was non-technical. The rest was technology or Internet. Today, we are back to a \$1B run rate, but 70% of our business is non-technical. This non-technical (or enterprise) business is our current key target, and in the past three years we’ve grown it by 133%. To put this in perspective, at a 25% growth rate, we’d be over \$2B in FY07, and at 35% we’d be almost to \$3B. Aggressive, but not crazy.

In “**A Vision of NetApp at \$2.5B**”, I’ll describe what I think NetApp will “feel like”. Who will our customers be? What will our market mix and product mix be like? What factors will have lead to our success? I believe that the vast majority of our growth will come from selling the products we are already shipping to customers we already have. I believe that the key to success will be learning to satisfy roughly 150 Named Accounts and 30-50 Strategic Accounts.

My first goal in this section is to tell a “future history” story of what NetApp will need to have accomplished in order to get to \$2.5B. My second goal is to encourage people throughout NetApp to tell more detailed “future history” stories about what their own group or department must do in order for NetApp to reach \$2.5B.

In “**Strategies and Thinking Tools**,” I will pull out some important themes from the previous section.

In “**Beyond \$2.5B**,” I’ll talk briefly about a longer term vision for NetApp, and why I think that getting to \$2.5B is the right focus for now.

Throughout this paper, I have two goals. The first is to share some thoughts on what I think we need to do. Second, and more important, I want to create some useful tools for thinking, and a framework for discussion and argument. Even if specific ideas are wrong, having a common framework should be useful.

This paper focus primarily on non-technical strategy issues, because I believe that new technology innovation will not be the key factor in our growth from \$1B to \$2.5B. On the other hand, I believe that our innovation over the next 3 years will be critical to our long-term success in the market after we have reached \$2.5B. My next

project is to write a companion paper on “**NetApp’s Technology Vision for 2003-2008**”.

I want to thank James for steering me as I wrote this. I started by whining about the challenges we face in the market, and James steered my whining into a vision of how we can succeed. In addition, both Dan and Mark Santora used early drafts of this paper to guide discussion at strategy offsites, and that provided many refinements and new ideas.

Rob Salmon and Steve Gomo have growth scenarios that show \$3B in FY2007, which would mean that we reach \$2.5B somewhat sooner. I considered changing the target in this paper, but I think that as we move past \$2.5B, some of the arguments here will no longer hold true. Needless to say, the sooner we pass \$2.5B the better!

## 1.0 Market Position and Growth Requirements

Where are we, and where do we need to be? In this section, I use numbers from IDC’s 2002 market report to understand our position in the market relative to our competitors.

This chart shows the big picture.

<u>Market</u>	<u>Share</u>	<u>Position</u>
NAS	34.8%	#1
Networked Storage	7.8%	#4
Open External Storage	4.1%	#8

Source: IDC with some massaging. See the companion spreadsheet.

If you believe that NAS is a secure, long-term niche, then we have the #1 share, which is a great position to be in. If you believe, as I do, that DAS and SAN and NAS are converging over time, then the more natural market to consider is Open Systems External Storage – storage for Windows, UNIX and Linux. In that larger market, we are the #8 player, with only 4.1% share.

To have a safe share in the Open External market, we must be much bigger than we are now. In the long term, of course, I want more than a “safe share.” I want to be #1. But strategically, I think a safe share is the next place for us to focus. From there, we can consider the steps beyond.

### 1.1 Our Market Position

In a companion spreadsheet (Hitz\_BusVis03\_IDC.xls), I’ve got a detailed analysis of IDC’s 2002 Storage Report. If you want to dig into the numbers, read that. Here, I’ll give a brief summary of what’s going on in the market.

To understand the market in which we play, I examined three segments from the IDC report. The DAS (Direct Attach Storage) segment, the SAN (Fibre Channel SAN) segment, and the NAS (Network Attached Storage) segment. Together, I call these the Open Systems External Storage, or OE (Open External) for short. It is this Open External market in which we rank #8 with a 4.1% revenue share.

You can argue whether the more natural comparison market for NetApp is Open External as a whole, or just Open Networked (SAN + NAS). But to understand the dynamics of the market, you must include DAS. This is because DAS is rapidly shrinking as people convert to Networked.

Between 2000 and 2002, several big trends stand out:

- The OE market lost about \$8B, almost entirely from DAS. SAN and NAS stayed relatively flat. Today DAS represents about 50% of OE, SAN about 40% and NAS 10%.
- DAS shrunk dramatically, from \$14B to \$6B and from 71% to 48% of OE. IDC expects that to drop to \$3.7B and 23% by 2007. I believe this remainder will be very low-end. Also, the system vendors are generally maintaining or gaining share in DAS, but the storage vendors are losing. DAS may be the last stronghold of the “vertical market structure” in which system vendors own the storage.
- EMC lost big. In the OE market, they went from \$4.7B to \$2.2B. However, if you look at percentage TB sold, EMC actually went up. Nobody stole EMC’s share – they just cut their prices in half.
- Many vendors converted DAS revenue to SAN. For instance, IBM converted about \$500M of DAS revenue to \$500M of SAN. By TB share, EMC converted about 8 points of DAS share into 11 points of SAN share. Likewise, HP converted about 7½ points of DAS into 4½ points of SAN.
- Sun is a big loser in every segment, by both revenue and TB.

In the Open External market overall, there are 4 big players (EMC, HP, IBM and Hitachi), and a pack of middle-sized players (Sun, Dell, Fujitsu, NetApp and NEC). Here are some revenue numbers for 2000 and 2002:

	<u>2000 (\$)</u>	<u>2002 (\$)</u>	<u>2000 (%)</u>	<u>2002 (%)</u>	<u>Rank</u>
HP	4,575	2,827	22.4%	22.4%	#1
EMC	4,771	2,200	23.4%	17.4%	#2
IBM	1,352	1,467	6.6%	11.6%	#3
Hitachi	1,688	1,403	8.3%	11.1%	#4
Sun	2,499	972	12.2%	7.7%	#5
Dell	319	683	1.6%	5.4%	#6
Fujitsu	605	559	3.0%	4.4%	#7
NetApp	674	519	3.3%	4.1%	#8
NEC	313	268	1.5%	2.1%	#9
total	20,403	12,640	100.0%	100.0%	

Source: IDC's 2002 Storage Market Report.

Note: Includes Hitachi's OEM revenue (\$308M in '00 and \$509M in '00). See "Misc Math" sheet in companion spreadsheet.

Some observations about the big guys:

- Sun is half way between "big player" and "pack", but it's shrinking fast, so I put it in the lower group.
- Roughly, the "big players" have 10-20% share, and the "pack" has 1-5%.
- EMC's revenue dropped a lot, but I believe that is because they cut their prices. Now they are competitive, so I don't expect this dramatic shrinkage to continue.
- To be a big player takes roughly \$1.5B in storage subsystem revenue, with 12% market share.

Things aren't that much different in Open Networked Storage (NAS plus Fibre Channel SAN), as this chart shows.

	2000 (\$)	2002 (\$)	2000 (%)	2002 (%)	Rank
EMC	2,050	1,946	34.4%	29.4%	#1
HP	1,303	1,506	21.9%	22.8%	#2
IBM	71	687	1.2%	10.4%	#3
Hitachi	93	655	1.6%	9.9%	#4
NetApp	674	519	11.3%	7.8%	#5
Sun	606	426	10.2%	6.4%	#6
Dell	174	331	2.9%	5.0%	#7
Fujitsu	44	84	0.7%	1.3%	#8
total	5,951	6,615	100.0%	100.0%	

Source: IDC's 2002 Storage Market Report.

Note: Includes Hitachi's OEM revenue (pro-rated by % of Hitachi overall revenue in SAN vs. DAS). See "Misc Math" sheet in companion spreadsheet.

Again, there are "big players" and "pack." The big players remain the same (EMC, HP, IBM and Hitachi).

NetApp is almost big enough to qualify, but going from 11.3% share to 7.8% isn't a healthy trend. Two factors caused this. First, this period includes the dot-com crash. Second, other vendors had DAS install bases to convert to SAN, and we did not. (See especially IBM and Hitachi.)

This time, the "big players" have 10-30% share, and the "pack" have 1-8%.

## 1.2 Growth Requirements

How big must we be in order to be a "big player" instead of being in the "pack"? I believe that we need to be \$2.5B in FY2007. Sooner would be safer!

Several approaches gave me roughly this same answer. One technique was to ask James and Dan. The table below shows the results of several other methodologies.

	<u>2006 (FY07)</u>	<u>2007 (FY08)</u>
12% OE	2,815	2,996
15% NW	2,497	2,889
#1 NAS&iSCSI/FC "pack"	2,064	2,783

All of these numbers are based on various cuts of IDC's forecasts. I adjusted IDC's numbers to account for the fact that only 65% of our top-line revenue drops to "storage systems". The rest is service, software and other products like NetCache and gFiler. (See the companion spreadsheet for details on all of these numbers.)

In the first line, I simply took 12% of IDC's forecast for the Open External market. Today, 12% would be sufficient for third place, and even if the shares adjust

somewhat, it should still be enough to be a “big player”. In the second line, I took 15% of the forecast for Networked Storage. I used a slightly higher percentage because the market is more consolidated, and I worry that big players will convert even more DAS business to SAN.

The last line represents the scenario I think we should target: to be #1 in both NAS and iSCSI, and a pack player in Fibre Channel SAN. That would make us #1 in IP Storage, and one of the “big players” in Open External Storage overall. At \$2.5B in total revenue, this gives roughly the following targets for storage system revenue:

\$1.0B	Be #1 in NAS
\$0.3B	Be #1 in iSCSI
<u>\$0.3B</u>	<u>“Pack” Player in FC-SAN</u>
\$1.6B	Total Storage Systems

The remaining \$0.9B comes from service, support, software and non-storage product lines like NetCache and gFiler. (IDC counts NearStore as storage system revenue.)

## 2.0 Sanity Check

In “**Growth Requirements**,” I argued that we would be in a vulnerable and unsafe position if we were smaller than \$2.5B by FY2007.

That’s a fine desire, but what makes us think we can grow so quickly? We hit \$1B in FY01, and it’s taken us almost three years to get back to the same place.

Let’s put these figures in perspective. Looking at annual numbers, we grew to \$1B, shrunk down to \$800M, and now after almost three years we’ve gotten back to a billion dollar run rate. But if you look at quarterly numbers, our comeback has been much more dramatic. Our drop from peak quarter of \$270M to low of \$190M is a 30% loss, or the equivalent of dropping from \$1B to \$700M. And yet, our most recent quarter (Q1FY04) was our second best ever, and next quarter should be our new best.

Even more dramatic, consider that when the crash hit, only 30% of our revenue was non-technical – 70% was tech or Internet. Now, at roughly the same size, 70% of our business is non-tech. This non-tech or “enterprise” market is currently our key target, and in the past three years we’ve grown 133%! So I think these numbers are aggressive, but not insane.

The chart below shows possible sizes for NetApp at various growth rates, assuming that we do \$1.1B in FY03.

	FY2004	FY2005	FY2006	FY2007	FY2008
10.0%	1.1	1.2	1.3	1.5	1.6
15.0%	1.1	1.3	1.5	1.7	1.9
20.0%	1.1	1.3	1.6	1.9	2.3
25.0%	1.1	1.4	1.7	2.1	2.7
30.0%	1.1	1.4	1.9	2.4	3.1
35.0%	1.1	1.5	2.0	2.7	3.7
40.0%	1.1	1.5	2.2	3.0	4.2
50.0%	1.1	1.7	2.5	3.7	5.6

To put this in perspective, remember that our recent quarters have been running over 25% year-to-year growth rate, and that's been accelerating. The growth rate in our enterprise business has been almost 35%.

Our fiscal years don't align exactly with calendar years, but the difference is small. At a 30% growth rate, a fiscal year that's slipped out by one quarter is less than 7% bigger than the calendar year. Thus, it's reasonable to directly compare FY2007 with IDC's 2006 projections.

### 3.0 A Vision of NetApp at \$2.5B

The goal of this section is to ask: What will NetApp be like at \$2.5B? Who will our customers be? What will our market mix and product mix be like? What factors will have lead to our success?

My goal for this section is to write some "future history" stories, and to encourage people to write "future histories" for their own departments and groups.

To give a sense of what I mean by future history, let me give an example of a present day history:

In 1993, NetApp started selling to engineering workgroups, progressed to high-end technical and scientific customers, and then to commercial data centers with enterprise database applications. At the dot-com crash, 70% of our business was with technical and Internet customers, but by 2003 we had transitioned to 70% of our business being with enterprise customers. Our success with Oracle databases was a key component of our transition into this enterprise/commercial market.

I have some guesses about things that will turn out to be important in the story of NetApp growing to \$2.5B. In particular, I think the following areas will each have important stories:

- Simplicity

- Products, Customers and Markets
- NAS, SAN and iSCSI
- Partners
- Strategic and Named Accounts
- Total Customer Experience (CS, PS, Eng, ...)
- Vision and Technology
- Business Units and Product Lines

In the following sections, I'll tell the stories.

### 3.1 Simplicity

NetApp created a uniquely simple solution to the problems of storage and data management. We not only simplified our products – we also simplified our customers' environments, and we made NetApp simpler to do business with.

Our success improved as enterprise customers came to understand that simplicity was a key driver for other things they valued:

- Simplicity drives down total cost of ownership
- Simplicity reduces user error and downtime
- Simplicity reduces IT staff confusion and overhead
- Simplicity enables faster application deployment
- Simplicity is the absence of things that you used to do
- Simplicity is enterprise

In some ways, this was a return to NetApp's roots. Our first products were too simple to operate in enterprise environments, but once we finished adding the features required to play in the enterprise, we realized how important it was to re-focus on simplicity.

Simplicity requires consistency, repeatability, and regularity. NetApp learned that programmatic or process-based solutions were critical to providing consistent simplicity. NetApp learned to innovate around process, and to treat NetApp processes as "products" that need to be designed and implemented in order to meet customer requirements.

### 3.2 Products, Customers and Markets

To get to \$2.5B, NetApp needed \$1.5B in incremental revenue. This came almost entirely from products, markets and customers that NetApp already had in 2003.

NetApp did create new products, enter new markets, and sell to new customers. Doing this was critical for NetApp's strategic position as it reached \$2.5B, but few – if any – of these new areas made a significant revenue contribution.

Brand new products didn't contribute much revenue, but enhancements to existing products were important. In particular:

- An improved release model that fit enterprise requirements
- Cost reductions
- Features that were already in the pipeline in 2003
- Fast-turn-around client side features
- Features related to quality and serviceability
- Features related to simplicity, manageability and data management

(See “**Vision and Technology**” below for a discussion of why innovative new products and technology were important, even though they didn't directly contribute much revenue.)

### **3.3 NAS, SAN and iSCSI**

By FY2007, NetApp was widely recognized as one of the “big players” in the Open External Storage market. It not only retained its #1 share position in NAS, but it had become a major player in block storage as well. At \$2.5B, about 60% of NetApp's storage system revenue was from NAS, and about 40% was from block-based storage, split evenly between Fibre Channel and iSCSI.

NetApp began its entry into the blocks market with Fibre Channel SAN, but the size and maturity of that market made it difficult for NetApp to break out of the pack of mid-sized players and become one of the big players. Fibre channel reached 20% of NetApp's revenue, even though NetApp only had 4% market share, because the FC market was still so much larger than NAS and iSCSI in FY2007.

NetApp's success in Fibre Channel was critical to its credibility in iSCSI. iSCSI first took off in the Windows market, but also started growing quickly in UNIX a couple years later. NetApp capitalized on the unexpectedly fast growth of iSCSI and become the market share leader. NetApp kept the lead by combining its experience in Fibre Channel SAN with its experience in IP-based storage. In addition, the advantages of unified storage were more pronounced for iSCSI than for Fibre Channel, because blocks and files shared a common networking infrastructure, and separate storage systems made even less sense.

### **3.4 Partners**

NetApp's relationship with a handful of big partners was critical to its growth.

Oracle's relationship with NetApp remained tight as Linux based “grid compute farms” became common in the commercial space. This relationship gave NetApp great credibility, not only for Oracle related storage, but in the Enterprise in general.

Veritas was driven even closer to NetApp as EMC continued to attack Veritas's core markets. This relationship resulted in improved interoperability between NetApp and Veritas solutions, as well as improved customer confidence.

NetApp's success in iSCSI helped improve the relationship with Microsoft. NetApp's success in home directory consolidation created tension with some parts of Microsoft, but iSCSI solutions helped create real value for other parts of Microsoft, like Exchange and SQL Server.

In addition to the big partners, NetApp depended on many smaller partners for its success in providing solutions. NetApp couldn't develop every capability that the customer required, but by working closely with partners, it could take responsibility for providing a complete solution to the customer's problem. Perhaps the clearest example of this was in ILM (Information Lifecycle Management). EMC made a series of acquisitions to become the "one stop shop" for ILM, but NetApp assembled a collection of partners to successfully compete.

[Note: There is an on-going discussion about what areas we need to own ourselves, either through development or acquisition, and what areas can be covered by partnerships. Some people think that partnerships are almost always the better path, while others argue that we will be competitively handicapped unless we have direct ownership of more components of the whole solution. My intuition isn't strong enough yet to tell a future history story, but I think it is an extremely important question.]

### **3.5 Strategic and Named Accounts**

Between 2003 and 2007, NetApp's fastest growth came from its 200 largest customers – 50 or so Strategic Accounts and 150 or so Named Accounts.

The Strategic Accounts program focused on NetApp's very largest accounts, and it achieved its goal of getting 30-50 customers each generating over \$30M in annual revenue. In addition, a handful of accounts were doing \$75 to \$100M. In 2003, NetApp's average penetration into its Strategic Accounts was only 3%, but by 2007 it had increased this to 10%, and was well on its way to 20%.

Our biggest success with these customers came as they chose to "bet the farm" on NetApp – that is, as they chose to trust NetApp with some of their most business critical storage applications.

At first, it was mostly smaller companies in the Named Account program that "bet the farm" on NetApp, but NetApp focused on making these customers successful, on making them comfortable with the bet they had made, and these lessons helped NetApp to convert many more customers – including many of the largest Strategic Accounts – into "bet the farm" customers.

### 3.6 Total Customer Experience (CS, PS, Eng, ...)

NetApp focused on improving the “Total Customer Experience” for all of its customers, but the critical success factor was making sure that the 200 largest Strategic and Named accounts were happy.

NetApp succeeded by improving all of the contact points with customers, not just the obvious contact points in Sales and CS, but contact points throughout the company. For each of these customer contacts points, NetApp asked, “How can we make the customer’s experience as simple as possible.”

Success came from fundamental changes in many groups at NetApp. All groups had two design criteria in common. First, simplify the customer experience in every way possible. Second, use innovation to create a uniquely "NetApp experience".

#### Customer Support

NetApp’s focus on improving both the perception and reality of its Support and Services capability became a source of strength in the market, especially with major enterprise customers. Enterprise customers made decisions primarily based on 3 criteria: features, support and price. Also, enterprise customers had a different definition of “support” than our earlier customers, and the new definition caused NetApp to reverse its support strategy in many ways.

In the old model, CS used a centralized approach that relied almost exclusively on major Global Support Centers (GSCs). CS was structured to handle the flow of new cases: How should cases come in? Where should they be routed? How should they be handed off from one GSC to the next? Ultimately, how should they be closed? The focus was on one case at a time, flowing through the system. That model worked when NetApp was small, and it was okay for smaller channel and territory customers, but enterprise customers needed something different.

Enterprise customers expected more personal, more face-to-face support. NetApp changed to a customer centric model that focused more on each customer’s overall experience and less on dealing with cases one at a time. How many cases has the customer had recently? How is their environment doing overall? Is this most recent case the first after a long period of quiet, or is it the straw that breaks the camel’s back? In other words, we treated each account as a project to manage, and not as a series of unrelated events.

This new model was delivered primarily in the field, with support from the GSC and its worldwide knowledge base.

Additionally, in early 2004, NetApp began deploying new remote management technology that lived on-site with the customer and improved supportability of complex customer environments.

## **Professional Services**

Enterprise customers also asked for NetApp to play a larger role in the deployment and operation of their storage infrastructures. To meet this need, NetApp grew its Professional Services organization. Early in this period, the role of Professional Services was defined and discussed often throughout the company. That role has never changed: to increase the absorption of NetApp technology in the market by focusing on the success of our customers.

The year NetApp first did \$2.5B in business was also the year Gartner Group described NetApp as "the role model for tech companies that want to excel in enterprise customer care."

## **Engineering**

Engineering contributed in two key areas.

First, Engineering developed industry leading product quality, supportability and manageability features. By 2003, NetApp products already had much of core functionality required to play in the enterprise. What became increasingly important was for engineering to understand how enterprise customers operate large data centers scattered around the world, and to make sure that the products fit well into that environment. NetApp Engineering developed "enterprise simplicity".

Second, Engineering created a compelling vision and a track record of delivering innovative technology. This gave our large customers confidence that NetApp would remain a market leader in the long-run. (See "Vision and Technology" below.)

The first area was a prerequisite for the second, because – without quality products that worked well in the Enterprise – the vision and innovation would have been irrelevant.

Engineering struggled with the tension between providing conservative enterprise quality and supportability, on the one hand, and continuing with vision and innovation on the other. In the end, it delivered both.

## **Plus More**

Manufacturing made "bet the farm" customers happy by making sure they had priority in receiving equipment they needed in order to run their business.

Finance learned more about the unique business requirements of key customers and verticals and offered financial solutions tailored especially for them.

As NetApp became more and more global, many organizations had to develop functions in place around the world so that customers didn't have to wait for Sunnyvale to wake up in order to get their business done.

Plus more.

### **3.7 Vision and Technology**

NetApp's growth from \$1B to \$3B depended on a market perception that NetApp had a winning vision and a strong track record of delivering cool new technology.

Even though technology that wasn't already shipped or in the pipeline in 2003 had little direct impact on revenue, it was critical for maintaining customer confidence. For large enterprises, the cost of switching vendors is very high, so their continued investment in NetApp depended on their confidence that NetApp was an innovative technology leader and was going to remain one. The evidence that convinced them of this was NetApp's vision and roadmap, as well as a track record of delivering cool new products as promised.

In addition, the new technology developed during this period positioned NetApp for continued success after it reached \$2.5B.

### **3.8 Business Units and Product Lines**

In this paper, I've focused almost exclusively on our filer product line. On the one hand, this is appropriate since it generates the bulk of our business.

On the other hand, I think that our ultimate success will depend on a coordinated strategy that encompasses all of our business units and product lines. Each business unit needs its own future history story, and in addition, we need to understand how they play together to accelerate our overall growth.

NetCache, for instance, has often been the first product to be sold into a new enterprise account. This business unit was probably the hardest hit in the dot-com crash, but the Internet itself certainly isn't going away. In addition, I continue to believe that caches are a form of "self managed distributed storage", and that the ideas of caching blend well with the ideas of storage. How can we best leverage this?

I expect gFiler to be another product that opens the door to new accounts. Many potential customers already have a preferred storage vendor, but gFiler will allow us to sell to them anyway. An important question around gFiler is how important the role of "Fibre Channel Virtualization Engine" will become.

NearStore may have the trickiest story of all, because I believe that in 3 to 5 years, many customers will be using ATA drives as primary storage for mission critical applications, including databases. Will we move ATA to the Filer line, or will we call all ATA-based storage NearStore, in order to distinguish it from Filers as they move even further up-market? In addition, ATA technology is central to two giant new markets: disk-based backup and information lifecycle management (leveraging different classes of storage). How will we engage those markets?

### 3.9 Future History Conclusion

This is just a sample of the “future history” stories that could be told. Maybe my stories are accurate, and maybe they are not. It doesn’t matter for my bigger goal, which is to foster a style of thinking. I want people to ask what NetApp will be like – how NetApp will feel – when it hits \$2.5B.

I know I’ve left out many important stories, and I encourage people to flesh out the details I haven’t covered. Here are a few future histories I wish I knew:

- Which applications and workloads were most important to our success?
- Which vertical markets were most important to our success?
- Which 5 or 10 customers first reached \$50M?
- Who were the top channel partners? How and why did they succeed?
- Which competitors were most problematic? Which did we beat easily?
- What solution areas mattered most to our customers?
- Which new market areas did we begin that positioned us well for growth beyond \$2.5B?

And never mind the company as a whole! How will my group feel, or my department, when NetApp hits \$2.5B? What investments matter most in reaching this goal, and what investments can we postpone to make room for more important ones?

### 4.0 Strategies and Thinking Tools

Some common themes emerged in the previous section, and my goal here is to bring them into clearer focus.

#### **Total Customer Experience and “Bet the Farm” Customers**

My instinct is that our overall success on the way to \$2.5B is tightly linked to the concepts of “Total Customer Experience” and “Bet-the-Farm customers”.

I believe that our biggest growth opportunity will come from giant customers deciding that they trust us with their most mission critical data. Put differently, if we can’t get giant customers to trust us, and serve them well when they do, then this will be our biggest barrier to growth.

My favorite definition of Total Customer Experience is that it is the stuff we must do throughout NetApp to make “Bet the Farm” customers confident that they have made a safe decision. At every point of contact between customer and NetApp, do they have an interaction that meets their needs and gives them confidence?

It’s one thing for us to decide that an account is strategic, probably because they have a giant IT budget and we want our share. It is an entirely different thing when a customer decides that we are strategic to them because they use us for their most business critical applications.

We need to take especially good care of these customers, because their experience – good or bad – will influence other companies as they consider whether to “bet the farm”. These customers have especially high expectations, because the impact of failure is so high. If we can learn to make them happy, it will serve us well in the rest of our customer base.

I admit that few companies are likely to make such a strong commitment to NetApp that they would literally fail if we don’t do well for them, so “bet the farm” is probably an overstatement. Still, I like the term because it emphasizes how important it is for us to serve them well.

At first, most bet-the-farm customers will probably be from our Named Accounts program, but our ultimate goal should be to get our largest Strategic Accounts to this point.

### **Focus on 30-50 Strategic Accounts**

I believe it is critical that we grow significantly with our current 26 Strategic Accounts. Today, we have roughly 3% average penetration in their storage purchases. We are so small that they barely notice us. If we can’t drive our penetration to 10 or 20% in our biggest most important accounts, what makes us think we can succeed in new accounts?

Some people have argued that to triple the revenue from Strategic Accounts, we should triple the number of Strategic Accounts. I think that is missing the point! If we triple the number of accounts in order to triple the revenue, then we will **still** have an insignificant percentage of each account. I want a pool of giant customers that have made such a big investment in NetApp that they care strongly about our success.

As we start having accounts that pass \$75M or even \$100M in revenue, we need to become even more focused on what they want. Today the NetCache and NearStore product lines are about this size, and for them we have a

handful of dedicated product market people and 30 dedicated engineers. Should we have the same for the “Verizon market”? And GE? And Nokia?

I don’t know the answer. Perhaps that’s not what they need. But my instinct is that we ought to be thinking fundamentally differently about customers of this magnitude. We ought to be structured to listen carefully and respond quickly, not just in Sales and CS, but also throughout the company.

### **“Getting to \$2.5B” versus “Winning after \$2.5B”**

I think we need to distinguish between two very different styles of strategic thinking.

The first style asks, What do we need to do in order to reach \$2.5B? I’ve argued that it will be primarily issues about our existing customers, our existing products, and our existing markets. What must we do to optimize revenue in these areas? What programs have the best ROI? Where are the big wins?

The second style asks, What must we do now in order to be well positioned to win after we reach \$2.5B? At \$2.5B, we will be the smallest of the “big players”, so we need to be well armed for further growth. This will be less about revenue and ROI, and more about market position, partners and vision.

### **Focus Beyond Technology**

I worry that we tend to overestimate the importance of technology, at the expense of other critical factors. For startups, great technology can be the primary fuel for growth. At \$1B and beyond, other market forces become much more important.

Which big industry players want us to succeed, and which want us to fail? Do we have allies (partners and customers) who want us to win because they depend on us? Do we have enemies who are focused very specifically on killing us?

Don’t get me wrong. I believe that great technology is critical for our long-term health, critical for maintaining the confidence of big customers, critical for success after we have reached \$2.5B. (Plus, the great technology that we’ve already developed will be critical in getting us to \$2.5B.)

My concern is that our instinct at NetApp is to focus on our technology and on how our technology will help us win, and to neglect factors that – in our immediate goal of reaching \$2.5B – are actually more important.

## 5.0 Beyond \$2.5B

This whole paper is focused on how to get to \$2.5B by FY2007, which should position us as one of the smaller “big players” in the overall Open External Storage market.

Why is this the right goal? Shouldn't we be targeting the #1 position?

Absolutely! Our ultimate goal should not only be to reach the #1 position, but also to drive the restructuring of the storage market from a vertical configuration dominated by systems vendors, to a horizontal configuration dominated by focused storage vendors. We don't just want a #1 position in a market we share with systems vendors – we want a #1 position after we've driven them out of storage.

However, that's a long way from where we are today, with our #8 position and 4.1% share.

I believe our most immediate strategic concern is to outgrow the vulnerability we face as DAS, SAN and NAS converge. We need critical mass in the unified Open External Storage market that is emerging. We must keep a strong focus on this important next step.

Nevertheless, we must also remember the bigger picture. Several times in this paper I've referred to how we will be positioned to compete in the market after we have reached \$2.5B. For this longer time-horizon, I believe that the new technology we develop over the next 3 years will be critical, and my next project is to write a companion paper to this one on “**NetApp's Technology Vision for 2003-2008**”.