

SunNetwork Shanghai Keynote

June 2, 2004

Mark Canepa, Executive President, Network Storage Sun Microsystems Inc.

Mark Canepa: Good morning and welcome to all of you here in Shanghai. This is my fourth visit to the city over the last several years and it is just amazing how much more modern and more sophisticated this city becomes every time I come, so it is just really exciting to be here. I know all of you have been sitting here now for almost two hours so what I want to do is get everybody to stand up and stretch your legs a little bit because we still have a bit of time to go. Everybody get up and stretch, very good.

What I am going to do for the next 30 minutes or so is talk about the connection between the network and the services that are becoming available on the network, and the storage and the data management. I really want to focus on three things. I want to focus on how this explosion into the network that Jonathan talked to you about is turning into an explosion in the data environment. We have got to realize that

that is taking place. Therefore it is really all about being able to understand this data explosion and taming the data explosion that is going to allow you to take competitive advantage in this increasingly global marketplace. Then what I want to do is share with you breakthrough business models as well as breakthrough technologies that are going to allow you to tame this data explosion and therefore reduce the cost and complexity of using, managing, turning data into information that can allow you to take competitive advantage. So the network explosion is real; Jonathan talked very much about it. If we could look at what is taking place just in the traditional network devices, the expectation according to industry analysts is something upward of 17 billion traditional devices; PCs, hand-helds, automobiles and appliances that will be connected to the network. They are going to be communicating vast amounts of information and data. But just imagine what happens when you add the traditional devices, the sensors capability, the RFID tag and these other small devices that will be communicating literally billions of pieces of information to the network. In the next several years a lot of your competitive advantage is going to come by having those devices and most importantly learning to take advantage of the information that those devices are providing for you. This network explosion that is leading to data explosion is worldwide. As a matter of fact, some of the fastest growth in data according to analysts is happening here in Asia. You

are learning here to take advantage of this vast amount of information that the network is going to collect and you are going to need to deal with this information in new and innovative ways. So therefore a lot of times when we sit and I talk to customers about storage, many companies still think about storage as, "I just need to find a place to put the data that my applications are generating. Obviously with the basis of what we do in data management somewhere along the way you have to store the data. But think about the data that is getting thrown at you. Increasingly you have data that has to be kept for a long period of time and you have data that needs to be kept for a short period of time; you have to deal with them differently. You have data that comes to you in big sizes and you have data that comes to you in small sequential sizes; you have to be able to deal with that. You have data that is more mission critical to your business operation and there is data that is less mission critical to your business operations. Increasingly you have got to figure out how to store that information in the right place. You have to be able to manage that information. You have to know where it is. You have to be able to retrieve that information and increasingly you have to do what I call Google your business. You have to be able to search and sort and extract and make sense of that information quickly and reliably. The better you become at doing that the more you can reduce the cost and complexity of running your data center environment and the more

quickly you can add new services in order to increase your business advantage and your competitive position. A lot of what we are going to be doing together over the next ten years is really taming the complexity of the data that is going to be showing up at the doorstep of your data center. We must be prepared to deal with it.

So how are we going to be doing it? What are the capabilities that we at Sun are thinking about so that we can prepare together to absorb and deal and manage and turn data into competitive information? To Sun it is all about getting the system right. The system is not just servers, it is not just storage, it is not just operating systems, it is not just middleware, it is not just network; it is all of those. But more importantly it is all of those things working in concert to give you an advantage. So we think a lot about the right system and the right system architecture and to us it really all starts with the application. In the initial parts of Johnny's presentation he talked about how the network got put together, how the devices and the clients are on one side and the applications that create and manage the services are on the other side. It is very important to have the right mechanisms to handle applications and he talked about Solaris 10 and how Solaris 10 enables an end-to-end infrastructure, allows you to create a very efficient way of running lots and lots of applications on a server infrastructure. It is key to realize that the more services you

have in the data center implies that you are going to run a lot of different applications that are going to have to interact with each other and with the information that you have within your data center environment. Those applications want to access information the way they want to access information. This leads to the creation of an access fabric, a methodology for accessing information. But while in the last several years applications accessed a lot of kind of raw, basic data and dealt with it, increasingly applications want to deal with data at a very high level. They do not want to deal with the raw data. They want to deal with a set of data services, whether they are block kinds of services or file kinds of services. They want to deal with data in a virtualized abstracted way. Applications do not want to have to know the physical environment that is storing their data. They do not want to know what kind of disk drives are underneath there. They do not want to know whether the data is local, whether the data is remote, how often the data is being backed up. They just want to know that when they want to access a piece of information so that it can do whatever they need to do with that information in the application the data will be there the way they want it. Data services are a key innovation required to get the system right. Then underneath there is a whole physical environment, whether it is expensive mission critical disks, whether it is mid-range, whether it is low-end, whether it is near line, whether it is tape, whether it is whatever the physical storage

mechanisms local or remote. It becomes the physical underpinning of getting the system right. We have spent the last several years within Sun clearly understanding all of the pieces that have to come together and we have invented as well as acquired a key set of technologies that allow us to enable the right system architecture. We are firmly committed to ensuring that this investment in these technologies continue to drive the right system and the right system architecture. Now, once we have created the systems capability we have to tile in the right product portfolio so that we can give you the right set of products with the right cost structure and the right capability. So over the last several years we have been hard at work understanding how to tile a set of capabilities that connect to the fabric that the applications must access. I am going to fill you in on what we have been able to do. At the high end we fully realize we require high-end mission critical data center storage capability for those kinds of applications that require the data to be there all the time, in large quantities, in large sizes for big, critical applications. However, as I mentioned to you earlier not all of the information is created the same. Not all of the information wants to live in systems that have this kind of cost structure. So we have been focusing at the other end with our Sun Storage 3000 product line and I will talk some more about new members of this family today. We have been focusing on providing a low-cost, high-value data repository environment for that information

where cost of storage and retrieval is the primary variable. However, not all information is at one of those two extremes. So we have been focusing our mid-tier and kind of balancing out the requirements of the data center with the requirements of the cost of the volume play and blending the two together while adding a whole set of ease of use kinds of capabilities. In order to let you manage this important, critical and fastest growing part of the data repository space. Many of you have been asking about file capability. More and more applications use file access methodologies. We have been tiling a product line called the Storage 5000 product line and I am going to be hinting at you some of the products that are going to be coming in this area. The products together with enterprise quality storage management gives us the right portfolio to enable the taming of the complexity of the data that we are going to be faced with over the next few years.

But it is more than having the right product portfolio. It is extremely important to have the right delivery model. Jonathan talked about new and innovative ways around which we can deliver capability to all of you. He singled out an example around the developer community where if you sign up for our developer tools you can fundamentally get the underlying hardware infrastructure for free. With our storage environment there are many different models for you to be able to acquire it. Some models may work better for some of you; some

other models may work better for some others of you. There is still the traditional purchase model, the lease model, the utilization, the pay for capacity used model. But increasingly and I am going to be announcing today a set of subscription capabilities so that you do not even have to worry about owning and managing the storage environment. Sun can come in with its partners to help subscribe to a data management model and environment.

So we have talked about the right systems architecture, that system architecture leads to the creation of a given product portfolio that enables the implementation of that architecture. I stand up here at just about every one of our NC launches on a quarterly basis acquainting you with sort of the next steps we are doing towards the execution of the road map of that product portfolio. This time around it is no different. We are continuing to execute on our storage and data management strategy. What I want to talk to you about are three things today; new program, new products and then I want to talk to you about a few of the things that we are going to be officially launching at the next Network Computing launch in September. But just like Jonathan, I just cannot resist sharing with you some of the things that are going to be happening over the next several months.

So first, let us talk about new models through which you can acquire

mission critical data capability. Today we are announcing the creation and capabilities of power units for storage starting at less than \$0.02 per megabyte per year. You do not have to acquire any storage. You do not have to install, architect, configure any storage. You do not even have to manage any storage. Sun will come in with its partners and put together and manage for you an entire storage infrastructure and every month we charge you for the storage that you actually use all starting at less than \$0.02 per megabyte per year. This is upwards of three times cheaper than having to purchase, install, manage, deal with all of the complexities associated with managing your data environment. It is available through Sun and through its iForce partners starting in July and it is available on a worldwide basis; that is one. Second, we are making new product announcements in the area of our low-end storage family. We introduced the StorEdge 3000 family roughly two years ago. We focused very much on bringing high performance fibre-channel storage into the low-price market to give you very strong capabilities at volume system prices. However with the advent of new data types, with the advent of business requirements around capabilities such as compliance, around areas of near line storage, around the areas where you cannot afford to have long backup windows when your applications are unavailable. There has been a set of requirements around providing very low-cost storage for these kinds of applications.

So today we are announcing two new products. One is the Sun StorEdge 3511; it is a Serial ATA system designed specifically for near-line storage applications. The second product is the Sun StorEdge 3120. It is designed for low-end, high performance, high-density SCSI applications. Both of these products are out in the product pavilion and I strongly encourage you to stop by and take a look at the capabilities, discuss with our technical experts where these products can fit within your data center environment and capabilities. The second thing I want to announce is the release of the 4.1 version of our Storage Archival Manager File System and Quick File System software. This is the result of a technology acquisition several years ago. It started out as a large file system and hierarchical storage management for the technical computing market and it has been very successful in that area. But what we found over the last few years has been we could apply the capabilities of this software technology to other kinds of marketplaces. So a year or so ago we started shipping technologies around infinite archival for mail and other kinds of applications. I want you to look at those at our product pavilion. Sophisticated technology to manage more commercial environments. With the release of 4.1 we have enabled a set of technologies on this product to be the foundation for your information life cycle management capability. This is the technology that lets you acquire, manage and decide where to store, where to

retrieve fast quantities of information. It can literally handle trillions of files and pedabytes of data in one common management environment. We have over 2000 customers around the world that are using this technology in 8000 some odd installations of this kind of capability across a variety of markets. Now EarthSap* just happens to be one of the customers that is using this technology. Home Box Office, a division of Time Warner is using this technology in the broadcast space. They are using this technology to replace tape broadcasts and going to a fully digital disk technology for their networks. They have 32 networks that are broadcasting in the United States. We have got six of those networks converted and over the next couple of years we are going to be converting all 32 HBO channels direct live disk broadcasts and this is the technology that manages that vast quantity of information. Kangwon Land Korea Entertainment and Gaming uses this technology. Oxford uses this technology for pharmaceutical kinds of capabilities and I could go on and on. As I said, we have literally 2000 customers that use this capability.

I now want to switch gears. These are the official product announcements. You will find those products on our price list. You can go out there and avail yourselves of them. What I want to do now is talk a little bit about a product that is not quite ready but I just could

not resist. Just like Jonathan earlier I just have to share this information with you. Today when you are looking at your data center and you look at the kind of information you have you may have ERP data, you may have web data, you have OLTP data, you may have video data, you may have compliance data. You have all of this kind of data that looks very different, wants to be stored in different kinds of devices and wants to be managed differently. Today you tend to go and buy different kinds of storage for each one of those pieces of information and then you are faced with the problem of complexity of management of these various storage environments. They typically get built into different sub-nets of a storage area network. The tools are different; they manage different kinds of information. The utilization rates are not what you want them to be. You get performances that may not be exactly tuned to the kind of performance requirements that the application requires. It is a difficult thing to manage. Now imagine a system that within one management framework, that within one set of data services, that within one way of making things happen is able to combine different kinds of data, different kinds of storage classes in one system and be able to tune the performance, the scalability, the cost of the particular kinds of information you want to the system that you want to have. Furthermore, imagine the system that is simple to set up and manage. Imagine a system that gives you pooled storage technology to

increase utilization across different storage types and imagine a system that provides centralized data services. Eighteen months ago we went out and purchased a company called Pirus Networks. They were working on virtualization technology for the network but we found a bigger purpose for them. What happens if you bring network technology to storage and you apply those technologies towards building storage systems? Imagine being able to do that not at data center kinds of prices and costs but at mid-tier cost structures. We have this product in demo form in the whisper suite. I strongly encourage all 5000 of you to go out there and take a look at this product and our folks out there will be glad to share with you that information. The product is in beta today. We have beta sites across the world. We are getting incredibly positive feedback. The reason I want to share this part of the preview with you is that the product should be ready to ship sometime this month and so I do not want to wait until September. I will stand back up here in September again and talk about the product officially but I wanted you to go and take a look at it. I wanted you to generate some pilots and use this product for these kinds of systems and capabilities. In addition, I also want to preview the first set of products in our Sun StorEdge 5000 Family. The 5000 Family is going to be the designation for our NAS-based and file-based technology products. Its purpose is to deliver files across the network to be able to help radically simplify, take cost and

complexity while providing large scaleability for web services and office productivity environments and support all of the key industry operating systems; Linux, UNIX and Windows. This product should be ready within about the next 90 to 120 days so begin to take a look at the capabilities that we are able to offer.

So the talk today has really been what to do with the complexity of the data environment that you have. There are some specific things that I would like to leave with you. One is really sign up to understand enterprise class storage at \$0.02 per megabyte used; we take care of everything. It is not a lot different than your cellular phone service. You pay for the minutes you use. You let somebody else worry about ensuring that that dial tone is always there and when you send a voice mail it gets to the other end and when you make a call somebody picks it up. Build your ILM strategy using SAM FS and QFS. Get ready for the explosion of the data that your data centers are going to have over the coming years and learn to have a competitive advantage in dealing with that information and join various early access programs on the StorEdge 6920 system as well as our network-attach storage capability. Again the key take-aways; one, this network explosion is real and as the network becomes more sophisticated that network is going to present to you a lot more data than you are used to. You are going to have to go learn to deal with

that data, turn it into information, extract what you need for it and learn to take advantage of it. Secondly, tame that data explosion. It may be the only way in which you are going to be able to get significant competitive advantage no matter what industry you are in. It may be the data that presents itself through your data center that is going to provide the differentiation between what you have and what your competitor has. It may be life sciences data. It may be marketing data. It may be customer data. It may be supplier data. But it is that information that will give you the edge. Thirdly let us work together on focusing on the right breakthrough business model for your data environment and let us focus on getting the right technologies into your environment so that you can help reduce the cost and complexity of running your data center. With that I would like to thank you for listening to this presentation.

--END--