

# Michael Barmada

## Apple SAN Solution Meets Massive Storage Needs

By Dustin Driver

Dr. Michael Barmada at the University of Pittsburgh uses an Xserve G5 cluster to crack the mysteries of inherited diseases.

DNA stores more data per ounce than any man-made storage media yet devised. That makes detangling it and decoding gene sequences for statistical analysis a mathematically intense undertaking.

"From a numerical point of view, if you've got 1,000 individuals in your study and you're typing 100,000 to a million markers, you've got at least 200 million points of data," says Dr. Michael Barmada, a statistical geneticist with the University of Pittsburgh. "It's a lot of data to deal with. Apart from the simple logistical problem of how to create files that contain all of the data, how do you manipulate those files, how do you analyze them, how do you get a computer to actually do all of the calculations that are required?" And how do you store and organize million-plus data files from a few dozen researchers and students?



**"[Xserve RAID] was half the price of the Dell/EMC solution and gave us four times the storage space. It was just an amazing benefit, cost-wise."**

The short answer: Use Macs. Barmada and his team have assembled the second-largest genetics cluster in the U.S. The Xserve G5 cluster crunches numbers with 282 processors day and night, storing results on an Xserve RAID linked together with Fibre Channel and managed with Xsan storage area networking software. The system is a key component of the ceaseless quest to understand genetic predispositions to common diseases like diabetes, high blood pressure and Alzheimer's.

### Sequencing Genes to Find Diseases

Barmada is part of a four-person statistical genetics team at the University of Pittsburgh's Graduate School of Public Health. He compares the genomes of people who have diseases with the genomes of healthy people, looking for any differences that may influence a particular illness. He works closely with laboratory scientists that squeeze raw data out of gene sequences for processing. "We're all interested in understanding genetic variation and human genetic disorders," says Barmada.

Initially, Barmada reviewed genetic information from local populations, usually between 50 and 300 people. The sample sizes eventually grew. "The populations that the physicians are identifying are getting larger, or we're doing collaborations with multiple groups," he says. "Instead of getting 400 patients from one physician, we'll get four or five or six thousand patients from a set of physicians around the country or 20,000 or 30,000 individuals from physicians around the world. For each of those individuals we'll analyze 500,000 or a million markers in the not-too-distant future."

The process may yield more than 20 billion points of data per study, but the work

## Science

[Why Mac for Science](#)

[Hardware for Research](#)

[Software for Research](#)

[Research Solutions](#)

[Profiles](#)

[Related Links](#)

### Michael Barmada

1. [Apple SAN Solution Meets Massive Storage Needs](#)
2. [Cost-Effective Storage](#)
3. [SAN Deployment Made Easy with Xsan](#)

### Barmada's SAN

See a [diagram of Barmada's SAN](#).



### At a Glance

Dr. Michael Barmada of the Department of Human Genetics at the University of Pittsburgh's Graduate School of Public Health sifts through piles of genetic information looking for links between genes and common diseases like diabetes, Alzheimer's and high blood pressure. During his research, he amasses millions of points of data that need to be organized, analyzed and stored. The department originally had a 48-processor Linux cluster with only one NFS server that was incapable of handling more than 40 processes at a time.

### Apple Solution

Barmada received a grant from the National Institutes of Health to build an Xserve G5 cluster. The new system features 282 processors and 3.5 terabytes of Xserve RAID storage linked with Xsan, Apple's storage area network file system. One Xsan array is held at a colocation facility six miles off campus; the other is kept near Barmada's office. Now Barmada can crunch data with blinding speed and store millions of raw data files in one centralized Xsan storage volume. If he ever needs more space, he can add another Xserve RAID without interrupting the cluster. He manages both Xsan storage arrays himself using Apple Remote Desktop 2, Xsan admin tools and Xserve RAID admin tools through a VPN.

doesn't stop there. Barmada charts external influences like diet, lifestyle and environmental contaminants, collecting even more data. In the end, a predisposition for a disease can usually be traced to more than one gene or factor. "The problem with these types of diseases — cardiovascular disease, blood pressure or schizophrenia — is that there are so many potential genetic risk factors that basically there's a linkage on every chromosome in the genome," says Barmada.

Without powerful computers and a vast storage infrastructure, Barmada's research would be impossible. Using cluster processing and an Apple Xsan storage system, he hopes to understand the genetic and environmental combinations that cause the common diseases that plague society.

[Next page](#): *Cost-Effective Storage*

#### Xsan User List

---

The scientists who run the Xserve G5 cluster at the University of Pittsburgh reconfigure their system on a regular basis. In addition to AppleCare and Xsan support, they've also found a community of system administrators who can lend a hand. The Apple Xsan user group offers sound advice and discusses creative solutions to network problems. "Anytime we have a problem now, we go to the [Xsan user list](#) on the AppleCare Support web page," says Barmada. "Within a day or two we've got an answer from the community out there using Xsan. It's a combination of IT guys and scientists who have a lot more time to play with Xsan than we do. We put in Xsan and from day one it was in production."

[Home](#) > [Science](#) > [Profiles](#) > Michael Barmada

Copyright © 2006 Apple Computer, Inc. All rights reserved.

# Michael Barmada

## SAN Deployment Made Easy with Xsan

The departmental storage system can also grow without being taken offline. Thanks to Xsan, Barmada can add one or more Xserve RAID units — and thus more storage space — without interrupting the cluster while it works. Xsan simply incorporates the added storage and its volume grows. The cluster gets more space to hold queued jobs and users get a larger volume to store their data.

To link the SAN together, Barmada chose the Emulex 355 SAN Storage Switch. The 12-port Fibre Channel switch fits neatly in a standard 1U rack and supports up to 200 MB/s port speeds.

Setting up the Apple Xsan software was a cinch. “The traditional Mac ease of use has been ported to an application that has been traditionally not so easy to use,” says Barmada. “With other solutions, I was looking at weeks of training for my IT people. They didn’t have to go through any training for Xsan. They opened the program and started fiddling with a few things here and there and they were able to put the whole thing together.”

### Vastly Improved Backup Strategy

The university has two mirrored Xsan storage pools, one at the colocation facility and another on campus. When the cluster system was first installed, Barmada had to find a way to copy several terabytes worth of files with a minimum amount of downtime. Mirroring the drives over the VPN was out of the question.

**“With other solutions, I was looking at weeks of training for my IT people. They didn’t have to go through any training for Xsan. They opened the program and started fiddling with a few things here and there and they were able to put the whole thing together.”**

“We took all of the data in our department and backed it up onto the older NFS servers there,” says Barmada. “Then we moved those to the off-site facility and mirrored them to the second SAN using Xsan. Now we have two SANs, one in our department and one off-site. We do nightly mirrors using a very careful application called Psync. It takes a long time to do a mirror, but if we’re careful about how we do it, it doesn’t generate a lot of network traffic and doesn’t bother the system too much. Now that we have everything replicated at both sites, there’s not a lot of changes that occur in one day — usually about 4,000 to 5,000. The backup process takes around six hours.”

If any one server goes down, the cluster won’t crash. The system can lose a hard drive or server without going offline — only the jobs that were running on the bad server or drive are interrupted. That’s a vast improvement over the old system. “If a server crashed, the whole system would error out,” says Barmada. “We’d have to start everything over again.” Now he can swap out one server or drive and get the system back to full capacity in no time.

### More Data for Prevention and Cures

“For human genetics the ultimate goal is that everybody has their entire genome sequenced and filed when they’re born,” says Barmada. “So when we need to go do a research study, instead of collecting 100,000 or 500,000 markers, we’ll have three billion markers, we’ll have the entire genome sequenced. Then we can start to model what is actually going on in the genome.”

To keep pace with the enormous growth and complexity of data, biologists will need technology to stay ahead of the curve. “I see the need growing exponentially,” says Barmada. “The more I add, the more they use.” The current cluster is running at capacity and the available storage is being rapidly consumed. Barmada has already ordered another 2.8TB Xserve RAID to help scale the system with demand. In an effort to build resources for the future and anticipate future needs, the professor is teaming up with other departments in the university and other schools throughout the region. In

## Science

[Why Mac for Science](#)

[Hardware for Research](#)

[Software for Research](#)


[Research Solutions](#)

[Profiles](#)

[Related Links](#)

### Michael Barmada

1. [Apple SAN Solution Meets Massive Storage Needs](#)
2. [Cost-Effective Storage](#)
3. SAN Deployment Made Easy with Xsan



exchange for a new Xserve G5 or a monthly fee, they'll get processing time on the cluster for their own projects. This system of sharing resources will foster a free flow of scientific ideas and help further research in all of the university's departments.

[Home](#) > [Science](#) > [Profiles](#) > Michael Barmada

Copyright © 2006 Apple Computer, Inc. All rights reserved.

# Michael Barmada

## Cost-Effective Storage



Barmada and his colleagues need a lot of computing power as well as storage capacity to chew through the heaps of data they acquire. Their department originally had a 48-processor Linux cluster with only one NFS-based file server. “The users learned that if they put more than 40 processes on the grid it would crash the cluster,” says Barmada. “There wasn’t enough capacity on the NFS server.” Barmada submitted a grant proposal to the National Institutes of Health for a new system with 250 G5 processors and 3.5 terabytes of Xserve RAID storage.

“I had always been on a Mac and when the Xserve G5 systems came out it became a viable option for the type of computing we needed,” says Barmada. “The large amount of RAM we were able to put into the machines, 64-bit computing and AltiVec processing were all expected to enhance the performance of our genetics algorithms.”

**“[Managing the Xserve G5 cluster] is amazingly easy. Once I get onto the system I can open System Monitor and I’ve got a graphical view of all the machines on the cluster with lights to tell me if there are any hardware issues or buffer overflows.”**

Storage was also a crucial part of the compute infrastructure Barmada was envisioning. Without a large server system to juggle multiple processing jobs at once, his cluster simply wouldn’t work. He knew that Apple could provide a relatively inexpensive server, storage and SAN solution when compared with the competition. The Apple solution was “half the price of the Dell/EMC solution and gave us four times the storage space,” says Barmada. “It was just an amazing benefit, cost-wise.”

The university got the grant and Barmada never looked back. The new cluster is light years ahead of the old one. “On an average day we have about 400 to 500 jobs sitting in the grid,” he says. “We run 200 to 250 on the grid and another 150 to 250 jobs are waiting in the queue.”

All those jobs have corresponding files. “For statistical genetics, we use flat text files,” he says. “Some can get really big when you’re dealing with 20,000 people and 500,000 bits of information. But generally, they’re on the order of a couple of hundred kilobytes or a few hundred megabytes, maybe a gigabyte at most. But to really manage the data properly we’ve had to break large files up into smaller sets. Any one analysis of a project will generate thousands of files. Several users on the system have two million files or more in their directories. There are several million files on the storage array in the grid.”

### Advantages of a SAN Environment

## Science

[Why Mac for Science](#)

[Hardware for Research](#)

[Software for Research](#)

[Research Solutions](#)

[Profiles](#)

[Related Links](#)

### Michael Barmada

1. [Apple SAN Solution Meets Massive Storage Needs](#)
2. Cost-Effective Storage
3. [SAN Deployment Made Easy with Xsan](#)

### Hardware

- 123 dual processor 2GHz [Xserve G5](#) cluster nodes with 2GB of RAM each
- 5 dual processor 2GHz Xserve G5 compute nodes with 8GB RAM each
- 3 [Xserve RAID](#) units with 14 250GB drives each
- 5 dual processor 2GHz Xserve G5 head nodes with 2GB RAM each
- 4 dual processor [Dell 2550s](#) with 3GHz Pentium IV Xeon processors and 2GB RAM each
- 4 dual processor 1.5GHz Athalon nodes with 2GB RAM each
- [Emulex](#) model 355 Fibre Channel hub

### Apple Software

- [Mac OS X Server](#)
- [Xsan](#)
- [Apple Remote Desktop 2](#)

### Open Source Software Applications

- Allegro
- Aspex
- Emulate
- Fastlink
- FastLink
- GeneHunter
- Loki
- Merlin
- PDT
- Phase
- SAGE
- Simwalk2
- Simulate
- Solar
- Superlink
- TDTPowerCalc
- Transmit

### Apple Support

Xsan consolidates all the Xserve RAID storage into one large pool. This has an obvious advantage when it comes to organization — instead of 20 or more volumes peppered with files, one volume contains all of the department's research data. In an Xsan system, one Xserve G5 is used as a metadata controller, which is like a traffic cop, directing servers and workstations to the files in the storage pool. Xsan also feeds files to each storage device in the pool more efficiently than distributed storage models, where individual users keep files on various network volumes or their local hard drives.

- [AppleCare](#) Xsan Support
- AppleCare Premium Service and Support for Xserve and Xserve RAID
- Mac OS X Server Software Support

#### **Easy Setup, Administration and Expansion**

"I'm not a computer science person or an IT administrator," says Barmada. "I'm a professor in the department." Even so, he has no trouble effectively managing the 282-processor Xserve G5 cluster and SAN — remotely no less. The system lies in a colocation facility six miles from campus. "We use the Xserve admin tools, Workgroup Manager, Apple Remote Desktop 2, Xsan admin tools and Xserve RAID admin through a VPN.

"It's amazingly easy," he continues. "Once I get onto the system I can open System Monitor and I've got a graphical view of all the machines on the cluster with lights to tell me if there are any hardware issues or buffer overflows. There are command-line tools to check the grid and see how many jobs are running. And we're installing more sophisticated monitoring tools to tell what network operations are being done on the machines, how much disk I/O there is or how much CPU time is being taken up on a per-node basis."

[Next page:](#) *SAN Deployment Made Easy with Xsan*

[Home](#) > [Science](#) > [Profiles](#) > Michael Barmada

Copyright © 2006 Apple Computer, Inc. All rights reserved.