


BY J. K. BUCSKO

PACS Storage Isn't Rocket Science



Since at least the 1980s, information technology has more or less been ruled by the so-called Moore's Law—the axiom that every 18 to 24 months, the computing power available in the same space doubles. The economic corollary that computer prices fall by half during the same time period has also become a piece of widely accepted popular wisdom—apparently everywhere except in the diagnostic imaging department.

Many administrators are feeling pressured to take on enterprise PACS storage solutions that actually threaten to add unnecessary layers of technology and complexity to what should be a relatively simple application. That was the message Jacob Farmer, chief technical officer of Cambridge Computer Services, Inc., presented at an American Healthcare Radiology Administrators conference in October. Farmer, who works with corporate clients to design and troubleshoot information systems and technical strategies, brought more than 17 years' worth of storage technologies experience to his session, titled "Demystifying Data Storage for PACS."

"Storage is actually one of the simplest parts of computing," he said, "and it can be demystified into its component parts for hands-on understanding. All these acronyms that PACS administrators are now being showered with are derived from old, familiar technology."

Who Needs an Enterprise Solution?

According to Farmer, too many radiology department and imaging center administrators feel driven by marketplace trends to take on much more computing power than they really need. "The only thing special about PACS is that it's a hungry application—it consumes an endless amount of storage," he said. "But there's a lot of gross exaggeration going on [in the market now]."

Many hospitals and imaging centers are getting locked into high-priced systems that may not be necessary for radiology, said Farmer. While they may find the high

level of technical support comforting, many administrators feel out of their league with the rest of the software. The box below shows how Farmer categorized and compared the needs of enterprise solutions and "plain" PACS storage.

For diagnostic imaging, said Farmer, "it's not always about performance; it should be about access and backup, and now, in today's climate, about security. If you have data that's priceless, you have to protect it or you're out of business." So the key words to creating your best system are "off-site," "offline," and "affordable."

But information technology (IT) personnel tend to think of the radiology department or even affiliated imaging centers as identical cogs in the network. They'll attempt to reproduce a smaller version of the same system in every department. "IT people are inclined to bring their corporate standards to radiology," he said, "when they could really use a different system [with a simpler communications interface]."

His version of the "perfect" PACS storage system is one that is fast enough (but not necessarily the fastest); is easy to back up; supports live data copy at an off-site location; is easy to expand without shutting down the system; and, very importantly, is inexpensive—meaning your facility can purchase, maintain, and add storage devices as necessary.

You can even build your own expandable, affordable PACS storage system with primarily off-the-shelf components, according to



Farmer, who said he's helped clients do just that. And you can put together your own network attached storage (NAS) and storage area network (SAN)—although you may not need either one, despite what vendors tell you.

Simplifying Storage Options

To begin with, Farmer said, think of PACS as "buckets of storage" and you'll

Enterprise vs. PACS: Compare and Contrast

Enterprise systems are designed to handle multiple complex operations, including the following:

- computers of different vintages;
- computers running different operating systems;
- different applications with different needs;
- regular moves, adds, and changes;
- full weekly backups instead of incremental daily backups;
- premium management tools, which are built into the price; and
- storage management can cost more than the storage itself.

"Plain vanilla" PACS systems represent a problem of capacity, not complexity. Here are the common key characteristics of such systems:

- specifically designed to handle large amounts of storage;
- don't require intelligent storage devices;
- use relatively few servers accessing data; and
- build backups incrementally through accumulation, not erasure or modification.



Photo courtesy of NASA

be less likely to want to clap your hands over your ears next time a vendor rep starts spilling about SAN, NAS, CAS, SCSI, IDE, SAIT, and the rest of the storage alphabet. The box below lists the most common acronyms and their definitions, but remember that

the blanket term “storage networking” pretty well sums up the main concept as it applies to PACS.

File servers are great for PACS storage, said Farmer. Network file serving protocols are very mature, and every vendor supports them. More importantly, by using network file servers, you eliminate interoperability hassles. The common protocols are Network File System for UNIX, Common Internet File System for Windows, and NetWare for Novell systems—but if your PACS software works with one file server, it will work with any file server.

“This really is the way to go for long-term data,” Farmer emphasized. “File servers are totally industry standard; there’s no, ‘Sorry, it’s not supported.’ Plus they’re scalable and fast enough for PACS applications.”

And while the current PACS buzz is all about NAS, you don’t need to go into debt and tie yourself into a single-vendor system for the next five (or more) years. If you know your way around file server systems, you can make your own.

You’ve Got NAS

According to Farmer, the term NAS was created many years ago to differentiate very powerful, all-in-one file servers from conventional servers. Today, however, “NAS is a generic term,” he said. Proving the truth of Moore’s Law, PACS file server appliances have come to oper-

A PACS Storage Terminology Glossary

Block: A unit of data storage; the least common denominator in conventional storage technologies.

File: An object consisting of multiple blocks.

Meta Data: Data that describe other data—eg, file attributes, permissions, indices, and logs. PACS meta data comprise DICOM and modality information, location of images on disk media, etc.

Network Attached Storage (NAS): Specifically designed to be a dedicated file-sharing appliance, NAS supports no other network tasks except delivering data to the user. “Think of it as a funny name for a file server,” says Farmer. An NAS unit can be located anywhere in a local area network, and multiple NAS devices can be linked.

Storage Area Network (SAN): As the name implies, a subnetwork or special-purpose, usually high-speed network device for interconnecting different types of storage units working together in a larger network. So SANs were originally designed to operate maximally within wide area enterprise networks. Farmer, however, claims that today, “SAN is just a different way to plug in disk drives.”

Content Addressable Storage (CAS): A method of encoding and retrieving files and objects using a unique data address derived from what’s actually inside. Because all digital information, such as MRIs and x-rays, is fixed, CAS streamlines and speeds up the process of retrieving these large image files from offline archives, regardless of the actual physical location. “CAS is just another funny name vendors like to use for recognizable technology,” says Farmer.

Hierarchical Storage Management (HSM): This method for long-term archiving assigns and moves files to slower but typically less expensive storage media according to the need for access. For example, an image would first be available on disk; then, as it aged, moved to optical storage; and finally to tape. Once programmed, HSM file migration should be automatic. Farmer calls HSM “a dirty word in some languages—and a blessing in others.”

Redundant Array of Inexpensive Disks (RAID): A series of hard disks plugged together using shared logic to act like a single large disk. A RAID enables sharing or replicating data among multiple low-cost, low-intelligence devices, allowing users to increase capacity and speed cheaply. “With just a little extra logic to get the drives to read as one,” according to Farmer, “and you’re good to go.”

ate much like any other file server. “You have to ask, ‘How is this NAS any better than an ordinary file server? How can it justify its price?’” said Farmer.

Vendors typically answer that their NAS comes preloaded and preconfigured with many integrated utilities, such as replication and snapshots. But you may not need all that—and even if you do, file servers are one of the easiest computer components to configure. You can usually purchase utility systems from a number of other software companies (eg, many now come free with Windows and other standard applications). Or as Farmer noted, “You can find all kinds of consultants in the Yellow Pages,” usually for less than the NAS vendor will charge.

So how much should you pay to make PACS storage even easier? The answer absolutely depends on whether or not you’re an enterprise-computing customer, said Farmer. If you are, then a proprietary NAS, with all hardware and software from a single source, will likely prove cost-effective and efficient in the long run.

If, however, you don’t anticipate the need for several different servers touching the storage (multiple licenses), you want to be able to use readily available backup software (proprietary NAS systems are notoriously harder to back up than conventional units), and, most importantly, you don’t want to commit to buying all your storage from a single source at today’s prices, Farmer offered this alternative: Buy an

Disk Storage Cost Projections 1992-2010

Year	Cost Per 10TB
1992	\$40,000,000
1995	\$6,655,000
1998	\$1,107,226
2001	\$184,215
2004	\$30,649
2007	\$5,099
2010	\$848

— Original sources: Steve Gilheany, www.archivebuilders.com, and Ed Grochowski, IBM research division, Almaden Research Center, San Jose, Calif.

ordinary PC. Install a network card. Plug it into your network. Delete Solitaire, Minesweeper, and other games. Turn it sideways. Start storing files. “Voila,” he said, “it’s a file server.”

If NAS is a better, cheaper way of managing files, then the second-most-mentioned PACS storage option, SAN, is primarily “a better way to plug in hard drives,” said Farmer.

Because SAN presents “a more flexible, robust storage infrastructure,” it easily supports centralizing storage administration—and thus centralized backup. That’s a plus if you have a lot of data to manage, as you would in a very large networked system. But beware: You’ll likely have to buy your SAN and backup systems separately. And in many cases, you can centralize backup without buying an expensive specialized SAN device.

Many PACS administrators struggle with the question of NAS vs. SAN, when, said Farmer, “you don’t really need either. Or,” he added, “you could use both. They are not mutually exclusive.” An NAS/SAN hybrid can prove a simple, efficient, and cost-effective solution for many imaging facilities.

Why Mess with Success?

Regardless of what configuration or vendor you use, stressed Farmer, don’t allow yourself to be locked into an expensive contract for unnecessary storage or other IT overkill. If your current plug-and-play system already works fine, you may just be adding layers of unnecessary complexity—changes that could ultimately be costly in money, productivity, and accessibility.

Remember, too, Moore’s corollary—the cost of computing power continues to drop. That’s certainly true with storage as new technologies continue to refine your options. And as media prices change, the economics of storage management and backup will likewise change. Since 1989, according to research by IBM and others, disk storage costs have steadily fallen by a factor of three every two years, as Farmer illustrated with the chart to the left. “That chart has been right every year, so don’t lock yourself into today’s prices. Ride the wave,” he advised.



— J. K. Bucsko is a freelance healthcare writer and editor based in Westville, N.J. She is a frequent contributor to **Radiology Today**.