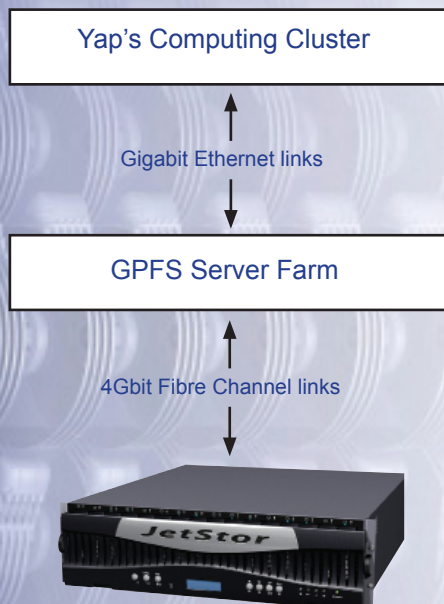




Robust Storage for High-Performance Computational Research



AC&NC JetStor® SAS 516F RAID Arrays

Yap is Perfecting Text-to-Speech with JetStor® RAID Solutions

THE ORGANIZATION

One of the more remarkable and useful capabilities of digital communications is converting spoken language into text, and no one does it better than Yap, Inc. Founded in 2006, Yap introduced the world's first fully automated cloud-based speech recognition platform. Its voice-driven capabilities let users interact more naturally with their handsets and offer service providers greater revenues per user from increased voice and data usage. Enterprises like Microsoft, Sprint, and MetroPCS use Yap's Speech Cloud™ for such applications as voicemail-to-text, mobile messaging, conference call transcription, and call mining. Moreover, when Yap was selected as a finalist in Silicon Valley's first ever TechCrunch event, often called the "American Idol" of the technology industry, its speech transcription technologies were widely acclaimed for curbing texting while driving. The firm, based in Charlotte, North Carolina, also won the VMA Innovation Award, a distinction given annually by Europe's leading wireless operators.

THE CHALLENGE

Developing software that can transcribe spoken language into text with reasonable accuracy is extremely challenging. Interpreting inflections, pitch, accents, dialects, and even the "ums" and "ahs" common in everyday speech demand programming that approaches the sophistication of artificial intelligence. Yap's solutions currently deliver an industry-leading accuracy rate, which is why they are being rapidly adopted by service providers and telecommunication companies.

This impressive performance, however, fails to satisfy Yap's team of scientists. Ryan Thomas, a research scientist for Yap, explained that a form of Moore's Law is at work in the speech transcription field. "Just as the number of transistors on a microprocessor will double every two years, the error rate for speech transcription drops by 10 percent every year," he said. "The field is very competitive and to flourish, Yap must remain in front of the curve."

As a result, Yap conducts vigorous research to improve its speech transcription algorithms. At its data center, it continuously operates a powerful computational farm to refine its statistical models by analyzing terabytes of people's speech. "The more data we process and study, the more accurate we can make our solutions," said Thomas. "Data lies at the core of our business."

Storing this data while making them rapidly available for high-performance computing is a daunting task. Yap surveyed the marketplace for solutions but rejected large storage arrays from manufacturers like EMC. "We were using the General Parallel File System, which is a clustered file system developed by IBM, so we



Yap's speech-to-text application for the iPhone®.

didn't need all the bells and whistles that drove up the costs of monolithic arrays," said Thomas. Moreover, Yap wanted systems that were modular, enabling it to quickly add storage capacity as its research demanded. "To support our work and ensure we remained competitive in the field, we needed primary storage that is fast, reliable, scalable, and cost-effective," Thomas added.

THE SOLUTION

JetStor® SAS 516F 16-bay RAID Arrays from Advanced Computer & Network Corporation (AC&NC)

SYSTEM CONFIGURATION

- JetStor SAS 516F 16 bay RAID Arrays with 450GB SAS drives
- High-performance servers from Super Micro Computer
- IBM's General Parallel File System (GPFS)

BENEFITS IMMEDIATELY REALIZED

By turning to AC&NC solutions, Yap gained a high performing storage cluster to support its research. Every JetStor SAS 516F RAID Array can house 16 450GB 15K RPM SAS disks for 7.2TB of storage. Yap's initial Jetstor installation was provisioned with 30 terabytes of capacity and linked to its server farm, comprised of powerful dual-socket, quad-core devices, with 4Gbit Fibre Channel connections. When even greater bandwidth is required, staff can upgrade the solutions to 8Gbit Fibre Channel links.

"Our JetStors store and deliver data as quickly as our servers can handle the traffic," said Ryan. Yap determined that each JetStor SAS 516F platform, when configured for RAID 5, delivers the I/O bandwidth to support three of its high-speed servers. "Unlike a large storage array, our JetStor cluster is modular. This was quite important because as we continually expand our compute farm, we can quickly and easily boost storage by installing additional disks and chassis to the cluster."

The JetStor solutions also provide the reliability to support Yap's around-the-clock research environment. "Once we set up an array, which is simple to do, it requires hardly any maintenance," said Ryan. "Our JetStors have proven to be rock-solid primary storage."

Ryan also explained that AC&NC delivered these capabilities at a fraction of the expense of many competing solutions. "A large storage array from one manufacturer cost four times more than our JetStor systems," Ryan added. "We pay only for the functionality we need and want. And whereas many providers seem to overprice their disks, AC&NC disks are very cost-competitive. JetStor RAID Arrays are affordable and deliver an unbeatable return on investment."

CUSTOMER TESTIMONIAL

HOW WE DID IT

Yap first built its research environment with SATA arrays from another vendor, but the cost of owning these systems was high because their disks were expensive. Additionally, Yap's research quickly advanced and the firm required greater performance from its primary storage. Yap then learned about AC&NC, discovering that the firm's faster SAS arrays and disks offered the sweet spot for performance, scalability, and economy.

Yap's **JetStor 516F SAS RAID Arrays** connect to the server farm at 4Gbit Fibre Channel speeds. The servers, in turn, attach to Yap's computing cluster with Gigabit Ethernet links. The **JetStor** solutions house all the voice data that Yap's computing resources access from the servers. "We're phasing out the SATA drives," noted Thomas. "The **JetStors** give us the fast storage we required to accelerate our research."

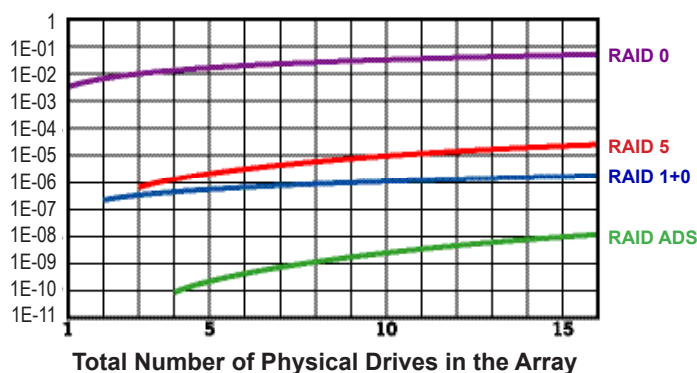
To present the **JetStor** arrays as a single storage pool to the servers, Yap deployed GPFS, a high-performance shared-disk, clustered file system. By clustering the storage devices, administrators improve data throughput, simplify scalability, and avoid managing the **JetStor** solutions as individual devices.

"The **JetStors** work well with GPFS, which is used by many of the world's most powerful supercomputers," said Thomas. "The combination of the AC&NC solutions and GPFS is ideal for a computationally-intensive environment like ours. With these technologies, we're poised to get ever closer to perfectly accurate speech transcription."

ABOUT ADVANCED COMPUTER & NETWORK CORPORATION (AC&NC)

Advanced Computer & Network Corporation designs, manufactures and markets high-performance, competitively priced data storage RAID systems. With an extensive line of storage solutions for any computing platform and operating system, comprehensive free technical support, and fast on-time product deliveries, AC&NC brings freedom of choice to all clients, from workgroup to enterprise. Our RAID systems ensure rapid, round-the-clock access to all of the information our customers rely on to move forward in the 21st century.

Probability of Data Loss by RAID Level



AC&NC JetStor® SAS 516F



USA HEADQUARTERS
Advanced Computer & Network
Corporation
5001 Baum Blvd., Suite 680
Pittsburgh, PA 15213
tel: 412.683.9010
toll-free: 800.213.2667
fax: 412.683.9070

www.acnc.com
info@acnc.com

SALES
toll-free: 800.213.2667
tel: 412.683.9010
fax: 412.683.9070
sales@acnc.com

OEM/VAR SALES
toll-free : 800.213.2667, ext. 122
partners@acnc.com

GOVERNMENT SALES
toll-free: 800.213.2667, ext. 125
gov@acnc.com
GSA Schedule GS-35F-0694J

EUROPE
Vector & Scalar Products Limited
tel: +44 (0)1628 891616
fax: +44 (0)1628 472137
sales@vspl.co.uk
www.vspl.co.uk