

## SUCCESS STORY

### ACCELERATE: LIFE SCIENCES

Children's Mercy Kansas City Reduces Time to Achieve Major Medical Breakthroughs for Critically Ill Children by Nearly 2x with Rapid Genome Sequencing Powered by High-Performance DDN Storage



### CHALLENGES

- Compute- and data-intensive genomics required high-performance, highly scalable storage
- Accelerate time to discovery while simplifying both clinical and research workflows
- Traditional scale-out NAS lacked scalability in performance and capacity to address data ingest and access demands

### SOLUTION

An end-to-end solution of DDN GS7K<sup>®</sup> transforms the lives of children by speeding the identification and treatment of rare diseases

### RESULTS

- Reinforced leadership in pioneering rapid pediatric genomic testing while strengthening goals to redefine pediatric medicine

### CHILDREN'S MERCY KANSAS CITY IS TRANSFORMING CHILDREN'S LIVES AND REDEFINING

pediatric medicine with an incredibly innovative genomics approach that speeds both the identification and treatment of genetic diseases. Consistently ranked among the leading children's hospitals in the nation, Children's Mercy also operates the world's first whole genome sequencing center in a pediatric setting.

At the hospital's Center for Pediatric Genomic Medicine, physicians collaborate with clinical laboratory scientists, molecular geneticists, bioinformaticians and software engineers to sequence and analyze rare inherited diseases. In 2012, Time magazine named the Center's rapid genome mapping test, STAT-Seq, a top medical breakthrough of the year for decoding a genome from start to finish in about two days—a process that traditionally takes six weeks or longer. And now, with the help of DDN Storage, they hope to reduce the decoding time again, dramatically shortening the critical time to diagnosis, and more importantly, treatment!

According to Shane Corder, HPC Systems Engineer, Center for Pediatric Genomic Medicine at Children's Mercy, rapid genome sequencing can be the key to diagnosing critically ill patients and prescribing viable treatments faster and more effectively. "Genomic sequencing, coupled with high-performance computing, gives us a whole new look at medicine by quickly shedding light on DNA variations that can explain a child's condition," he explains. "Our focus on developing clinical tests for next-generation medical treatments improves outcomes for patients at Children's Mercy and around the world."

In addition to expediting patient diagnosis and disease treatment, Children's Mercy relies on genomic testing to help doctors avoid unnecessary and painful tests, such as muscle biopsies. "It's not uncommon for children to endure dozens of clinic visits and scores of painful tests in search of a cause to their rare illness," Corder adds. "With genomic testing, we can get answers without repeated visits, painful tests or undue financial burden on the parents. Our goal is to provide a faster path home for patients and their families."

## THE CHALLENGE

Genetic sequencing is very compute- and data-intensive, which puts ever-increasing pressure on the Center's IT team to deliver ample processing power and highly capable data storage to support testing based on both whole genome sequencing and whole exome sequencing. While clinical sequencing is the Center's primary focus, Children's Mercy also pushes the envelope when it comes to researching genetic diseases.

With more than 6.4 billion bases in a person's DNA, encompassing 22,000 genes that code for 100,000 proteins, it's easy to understand the escalating demand for high-performance computing and storage. "Clearly, informatics is the bottleneck in genomics," says Corder. "Our goal is to keep pace with the data deluge in both our clinical and research environments so we can quickly analyze data to produce meaningful insights."

## BUSINESS BENEFITS

- Elevated support for next-generation genomics testing to accelerate medical breakthroughs for critically ill children
- Scalable storage also enables the Center to push the envelope in terms of researching genetic diseases

“

*Genomic sequencing, coupled with high-performance computing and DDN storage, gives us a whole new look at medicine by quickly shedding light on DNA variations that can explain a child's condition. Our focus on developing clinical tests for next-generation medical treatments improves outcomes for patients at Children's Mercy and around the world.* ”

“

*It's not uncommon for children to endure dozens of clinic visits and scores of painful tests in search of a cause to their rare illness. With genomic testing, we can get answers without repeated visits, painful tests or undue financial burden on the parents. Our goal is to provide a faster path home for patients and their families.* ”

### Shane Corder

HPC systems engineer, Center for Pediatric Genomic Medicine at Children's Mercy Kansas City

Time is of the essence when decoding genomes of seriously ill newborns to find the genetic causes of their illness and initiate viable treatment options. Likewise, bringing closure to a long and arduous diagnostic odyssey for children with hard-to-diagnose ailments is equally important. “There are roughly 8000 known genetic diseases in the U.S. alone and 1 in 30 children are known to have these,” Corder says. “One in six children are admitted to a Kansas City hospital for a genetic disease, and unfortunately, these genetic diseases cause 1 in 5 deaths here. The hard facts are that we know the cause of less than 5000 diseases, yet with our genomic testing we are making major progress in finding the causes and discovering new treatments.”

To that end, the Center embraces advanced technologies to help accelerate rapid genome testing. “We're trying hard to break the 48-hour window required for the crucial STAT-Seq test by applying advanced compute, storage and accelerated sequencing technologies,” Corder says. “The goal is to complete the entire process from enrollment to delivering a final report to the physician in 26 hours.”

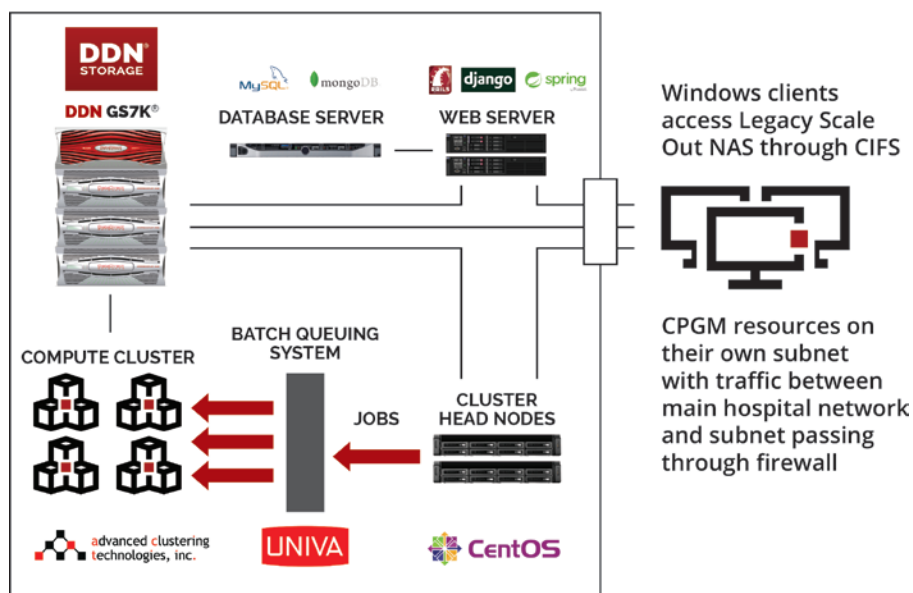
On the compute side, Children's Mercy has invested in a high-performance compute (HPC) cluster with 40 Linux nodes totaling 1,300 cores and an entire suite of software tools to meet a variety of sequencing demands while maintaining both its CLIA certification and CAP accreditation. The Center had previously deployed EMC Isilon storage, which initially met their storage requirement for handling both clinical and research workflows. Over time, however, Children's Mercy's traditional scale-out NAS storage lacked the scalability in performance and capacity to address demanding data creation and access needs. “We needed a highly scalable platform that could scale up or out to handle massive data ingest, processing, storage and collaboration,” adds Corder. “Ultimately, we needed a more flexible, powerful approach than traditional scale-out NAS could deliver.”

## THE SOLUTION

In moving from enterprise-class storage, the Center sought a more powerful platform that could handle data ingest and analysis in the same platform. The team also wanted to accommodate new instruments and modern techniques for accelerating sequencing and test results. Additionally, finding a high-density, massively scalable solution was crucial as Children's Mercy faced serious space constraints in its data center.

With all major technology purchases, achieving the ideal cost-performance ratio was important. “Our legacy storage was pretty pricey,” notes Corder. “It was clear we needed to optimize performance in the most cost-effective manner.”

## GENOME CENTER NETWORK



## TECHNICAL BENEFITS

- With DDN SFX, nine out of ten random read I/O requests are served from cache, which delivers across-the-board performance improvements
- Handling random data access or reads of millions of small files on dedicated solid-state modules removes the performance penalty on data-intensive storage requests
- DDN's storage versatility has resulted in the addition of a new EXAScaler appliance to meet the needs of a new compute cluster

“

*Clearly, informatics is the bottleneck in genomics. Our goal is to keep pace with the data deluge in both our clinical and research environments so we can quickly analyze data to produce meaningful insights. We needed a highly scalable platform that could scale up or out to handle massive data ingest, processing, storage and collaboration. Ultimately, we needed a more flexible, powerful approach than traditional scale-out NAS could deliver. ”*

“

*Not only did DDN's high-density weigh heavily into the final decision, the stellar sales, post-sales and technical staff was just amazing. It felt like they gave us the highest priority. Density is a huge factor for us. The opportunity to scale non-disruptively when adding more DDN storage is a major benefit. We're keen to do more with DDN. ”*

### **Shane Corder**

HPC systems engineer, Center for Pediatric Genomic Medicine at Children's Mercy Kansas City

After surveying the HPC storage landscape, the Center was most impressed with DDN's portfolio of industry-leading solutions for accelerating genomics workflows. In particular, the team saw an immediate opportunity to remove informatics bottlenecks with strong performance in a scalable solution that met both budget and density requirements. As a result, they felt DDN could keep pace with the diverse and dynamic demands of the Center's Illumina HiSeq sequencers. Moreover, DDN provided solutions for addressing both high-throughput genomics and global collaboration.

For example, DDN's GRIDScaler® platform, which integrates the industry's leading performance and density storage with IBM® Spectrum Scale™-based parallel file system, was ideal for supporting the Center's full range of clinical and research workflows. DDN delivered the complete end to end data management of the high performance and collaboration environment seamlessly. Another plus: DDN's ability to couple fast, scalable storage with the Center's planned deployment of the Edico DRAGEN Bio-IT processor, which is an FPGA-based genomic analysis acceleration technology.

Following a thorough review, Children's Mercy deployed the DDN GS7K® with 1PB of scale-out storage to meet performance, parallel I/O and capacity requirements. “Not only did DDN's high-density weigh heavily into the final decision, the stellar sales, post-sales and technical staff was just amazing,” says Corder. “It felt like they gave us the highest priority.”

## THE BENEFITS

Thanks to its industry-leading DDN storage and robust HPC environment, Children's Mercy now is achieving dramatic improvements in data alignment, analysis, discovery and collaboration. For instance, the DDN GS7K scales easily and seamlessly to support the processing needs of the Center's software tools, called the Nordic suite. These tools, which are used with different types of sequencing, testing and specimen options, include SSAGA, which maps genes to a symptom; RUNES, for variant characterization; and VIKING, which integrates outputs from the two previous tests for dynamic filtering and trio familial analysis.

With DDN, the Center has further accelerated their production pipelines, which can make a major difference in transforming the life of an acutely ill child. For example, the time to perform its STAT-Seq rapid genome test has reduced by about 43% from 2.5 to 1.75 hours. “Ingest from our Illumina sequencers also has increased significantly. Together, these performance improvements will prove invaluable in helping us to decode a full genome in less than 26 hours,” notes Corder.

The DDN solution provides benefits across other areas, as well. For example, traditionally Edico processing requires the transfer of data to local solid state drives (SSDs) on the DRAGEN server. DDN provides sufficient performance to run the Edico processor directly from shared storage, eliminating the complexity and time involved with the data transfer, and ultimately eliminating the need for expensive local SSDs.

DDN's GS7K also provided Children's Mercy with the ideal form factor, delivering high-density storage within its small data center footprint, and the ability to scale capacity in a simple fashion. “Density is a huge factor for us,” Corder adds. “The opportunity to scale non-disruptively when adding more DDN storage is a major benefit.”

The Center also is exploring opportunities to expand its DDN storage to reinforce its leadership in pediatric translational research. As such, the Center hopes to deploy DDN's WOS to facilitate secure sharing and collaboration with other clinics and researchers worldwide. “We're keen to do more with DDN,” concludes Corder.

## **ABOUT DDN®**

DataDirect Networks (DDN) is the world's leading big data storage supplier to data-intensive, global organizations. For more than 15 years, DDN has designed, developed, deployed and optimized systems, software and solutions that enable enterprises, service providers, universities and government agencies to generate more value and to accelerate time to insight from their data and information, on premise and in the cloud. Organizations leverage the power of DDN technology and the deep technical expertise of its team to capture, store, process, analyze, collaborate and distribute data, information and content at largest scale in the most efficient, reliable and cost effective manner. DDN customers include many of the world's leading financial services firms and banks, healthcare and life science organizations, manufacturing and energy companies, government and research facilities, and web and cloud service providers. For more information, visit our website [www.ddn.com](http://www.ddn.com) or call 1-800-837-2298.