Intel IT Center

intel Look Inside.™

HIGH-PERFORMANCE COMPUTING

Managing Vital Resources with Intel® Xeon® Processors

California Department of Water Resources creates a centralized HPC platform that doubles high-end modeling performance, quadruples capacity, and reduces operations costs for modeling by 45 percent

Delivering Business Value, Reducing Costs

Like many resource-strapped agencies and businesses, the California Department of Water Resources (DWR) handled its high-performance computing (HPC) modeling needs through a mix of high-end workstations, departmental servers, and outsourced services. But with modeling requirements rising and California's water challenges becoming critical, the mixed environment was becoming unsustainable.

Michael Hom
Data Center Chief, California Department of Water Resources

Collaborating closely with DWR's end-user scientists, the department's mission-focused IT organization decided on a new approach. "IT's job is to enable the business," says Tim Garza, DWR's chief information officer. "Our previous approach was becoming a constraint instead of an enabler to the aspects of the business related to high-end computer modeling."

DWR created a centralized, in-house resource optimized for HPC modeling. To power the environment, they chose HP ProLiant* SL and DL series servers, the Intel® Xeon® processor E5-2600 product family, Red Hat Enterprise Linux*, and an InfiniBand* fabric interconnect.

DWR's new environment is the first state-of-the-art, in-house HPC modeling environment for any California state department. And it's delivering strong value to California's scientists, citizens, and businesses. The platform increases DWR's high-end modeling performance by an average of 200 percent, expands its high-end modeling capabilities and capacity by 400 percent, and reduces a range of cost metrics, Garza observes.

By delivering that performance and capacity, DWR empowers its scientists and engineers to more effectively manage California's water resources. In doing so, they're helping to preserve California's delicate ecosystem, support the state's massive agriculture industry, and ensure clean water for every citizen. DWR's move from departmental to centralized HPC resources represents an approach that can benefit many other departments, agencies, and businesses.

> *"We needed a platform that could run diverse applications and scale up to meet the modeling needs of our program areas. There is no room for compute or memory bottlenecks, so we needed high compute performance, memory capacity, and throughput for large data sets."*
>
> Michael Hom,
> Data Center Chief,
> California Department of Water Resources

At a Glance

Project

- Create a new HPC modeling environment to support growing requirements, improve performance, and reduce costs

Accomplishments

- Increased DWR's high-end modeling capabilities and capacity by 400 percent
- Boosted high-end modeling performance by an average of 200 percent
- Dramatically reduced the modeling timeline, in one case from four months to 40 hours
- Captured annual cost saving/avoidance of USD 650,000
- Reduced operations cost associated with computer modeling by 45 percent
- Eliminated outsourced processing fees of USD 30,000 or more per model

Lessons Learned

- Recognize when it's time for a more serious approach to HPC. Educate end users on the potential benefits of a shared environment.
- Collaborate closely with HPC users and vendors to understand application requirements and design an optimal solution.
- Select vendors that have deep HPC expertise and are committed to your success.
- Choose high-performance, scalable technologies supported by an effective management stack.

Enhancing Water Management

Water management is critical for California, where a rising population, an ongoing drought, and a fragile natural environment have led to what Garza calls a slow-moving emergency. While California is a leader in water efficiency and conservation, its most recent drought is severe enough to threaten the state's USD 2 trillion economy. California generates almost 13 percent of the U.S. gross domestic product, including USD 45 billion in annual farm revenues. More than half the fruit, vegetables, and nuts grown in the U.S. come from California, with industries from winemaking to recreation to power generation depending heavily on clean water supplies.

One of the largest departments within the California Natural Resources Agency, DWR is central to California's water management strategies. The department supplies and manages the state's water delivery systems, including the California State Aqueduct, which delivers water to 28 million residents and 750,000 acres of farmland. DWR also helps coordinate the state's integrated water management strategies, inspects 1,200 dams, and improves and manages California's levees. DWR's scientific and engineering leadership and the magnitude of its water challenges have made the department a destination employer for job seekers in hydrology, environmental science, engineering geology, and other disciplines.

'A Major, Major Business Tool'

HPC is a vital resource for DWR's scientists, engineers, and technologists—and their needs have been growing. Over the last several years, DWR has segmented into multiple business units supporting areas such as water delivery, environment science, irrigation management, flood control and emergency operations, and energy production. Each line of business uses complex HPC modeling to make water management decisions that impact agriculture, the environment, and individual citizens.

"High-performance computing modeling is a major, major business tool for our lines of business to deliver their mission," explains Garza. "They're accumulating masses of data from many sources, including aerial imaging, water flow and land sensors, sonar, radar, environmental data metrics, and climate and other remote systems. They use complex computer models to make a variety of crucial water management and water delivery system recommendations and decisions. They also use HPC models to forecast and project items such as the impact of climate changes or drought conditions, flood patterns, and many other critical events and items."

Outgrowing the Piecemeal Approach

As both the department and the water challenges grew, DWR's piecemeal HPC approach began to constrain its modeling and decision-making activities. "Souped-up PCs and workstations place severe limits on performance, storage, and scale," Garza says. "Departmental servers are general-purpose systems that aren't designed for complex and specialized HPC modeling."

> *"Because the modeling teams are no longer constrained by the costs of outsourced services, the whole modeling timeline is faster."*
>
> Steve Croft,
> Chief Technology Officer,
> California Department of Water Resources

Outsourced HPC services had their own shortcomings. "The model process runs cost between USD 30,000 and 50,000 to execute, which really limited the number and size of the runs we could perform," Garza says. "They also didn't provide the flexibility and timeliness to meet our business demands and deadlines. In many cases, it took four to six months for scientists to create updated models and perform the required process runs."

Collaborating for Scale-Out HPC

Working closely together, DWR's modeling teams and technologists set out to create a centralized, on-premises environment that could satisfy DWR's diverse modeling needs and put HPC on a sustainable footing. Modeling experts from DWR's Bay-Delta and Flood Modeling offices—both of which rely heavily on advanced computational fluid dynamics and other complex modeling applications—were closely involved. The Sacramento-San Joaquin Bay Delta is a large, complex ecosystem and an important source of water for California's economy and citizens.

DWR identified usability, accessibility, flexibility, and performance as key criteria for the platform. "We needed a platform that could run diverse applications and scale up to meet the modeling needs of our program areas," says Michael Hom, data center chief at DWR. "There is no room for compute or memory bottlenecks, so we needed high compute performance, memory capacity, and throughput for large data sets."

After reviewing available technologies, DWR decided on the HP platforms and the Intel Xeon processor E5-2600 product family. "We had HP platforms and Intel® processors in our environment," says Garza. "We were familiar with the technologies, and they had already proven themselves. Because of where Intel sits in the industry and how it pushes the performance envelope, we were confident the chip would never become the bottleneck. So it was a strategic decision for us."

DWR also invited HP and Intel into the collaboration to help optimize the solution's design and the transition to an in-house modeling environment. "HPC is a different animal," says Steve Croft, chief technology officer at DWR. "It takes a different way of looking at things. So, this was not your typical IT project. There was a lot of collaboration and innovation going back and forth between multiple DWR modeling application teams, our data center operations people, and the HPC folks from Intel and HP. We would not have been able to design and implement an HPC environment that can handle this range of high-end workloads and deliver such high performance without their collaboration."

Screaming Throughput and Easy Management

The resulting platform offers power and flexibility for DWR's diverse and dynamic modeling workloads. "It provides the unique combination of high-performance computing but in a very highly scalable environment so that we can easily scale up to the modeling needs of our program areas," Croft says. "It is also very flexible in that it can support multiple modeling software packages and incorporate visualization as well as rendering at the highest density levels."

> *"This solution was a great blending of science, engineering, technology, and people. It all works together to deliver what is needed for our program areas."*
>
> Francis Chung, Ph.D.,
> Principal Engineer and Chief,
> Bay-Delta Modeling Support Branch,
> California Department of Water Resources

Key Technologies

- 55 HP ProLiant* SL and DL series servers with 588 Intel® Xeon® processors E5-2600 product family and Intel Xeon processors 5600 series, configured with 4 TB of memory
- HP Server Management Services*, HP Integrated Lights-Out* (HP iLO*)
- Red Hat Enterprise Linux* 6.1 operating system
- 36-port InfiniBand* switches
- 46 TB of NAS and SAN storage
- Intel® Fortran and C++ Compilers

Did the scientists get the throughput they wanted? "The throughput screams," Hom adds.

The environment has also proven easy to access, configure, and manage—a crucial consideration for a centralized HPC resource that must handle diverse workloads. "In high-performance computing, no two modeling jobs are the same," says Croft. "So while the infrastructure and the physical components don't change between jobs, our users can basically tear down and reconfigure the system in software to get the resources they need for their modeling job and results set. It's much more dynamic than you have in the transactional world."

DWR uses HP server management tools, including HP Integrated Lights-Out* (HP iLO*), to perform these tasks, optimize system utilization, and give each job the resources it needs for top performance. "It's basically a self-managed platform," Croft says. "The power users can do it themselves using the HP management stack. Our operating expense for the system is very minimal."

Savings Enable More Science

DWR's modeling platform is having a transformative impact. With significantly more capacity and performance at their fingertips, scientists can build more detailed models, simulate more complex and longer-range scenarios, explore more alternatives, and get results more quickly. They can enhance understanding, decision making, and communication by visualizing and rendering their results in finer detail.

The savings DWR is achieving from the new environment amplify these benefits. DWR is saving/avoiding USD 650,000 on modeling-related costs. It has eliminated outsourced processing fees and lowered its operations cost associated with computer modeling by 45 percent.

"Because the modeling teams are no longer constrained by the costs of outsourced services, the whole modeling timeline is faster," says Croft. "They can provide more models in real time. They can create more effective models because they're able to try more scenarios. So the system not only resulted in a rapid return on investment from savings on capital expenses, it also increased the department's ability to anticipate, forecast, and manage the impact of changes in water levels or in the environment. These results allow us to make better business, scientific, and engineering decisions to better serve California."

Changing What's Possible

For one of the Bay-Delta team's crucial modeling activities, the new platform reduced the modeling timeline from four months to 40 hours. Another of the team's complex models now runs at a 1-to-100 real-time ratio, allowing scientists to simulate a full year of activity within four days.

And several modeling analyses that simply were not possible before are now practical. These include complex climate modeling, 3-D groundwater model simulations, and daily time steps of the State Water Project and the Central Valley Water Project systems models.

"The work we are doing with this environment greatly helps advance the quality and quantity of the California Water Delivery System, the Bay Delta Conservation Program, Climate Change Program, biological opinion remand studies, and sediment and mercury investigations," says Dr. Francis Chung, principal engineer and chief of DWR's Bay-Delta Modeling Support Branch. "All these efforts help to better serve the people of California."

Win/Win

The shift to a shared, in-house environment for high-end HPC modeling is a win/win for DWR's scientists and technologists, and one that DWR's experts say other organizations may want to consider.

"So often, you'll have an application that starts small, so you put it on your workstation," says Garza. "Then, it grows and becomes a critical tool in your analysis. At that point, you're really limited in your ability to scale. More importantly, if this tool is part of your mission-critical delivery system around your business intelligence and your ability to predict and forecast, do you really want to leave that up to a workstation? Shouldn't you treat it as you would any other mission-critical resource and put it on more robust infrastructure? For us, it was a clear decision."

Collaborating for Success

Teamwork has been central to the project's success. "It's because of the collaboration with our users, HP, and Intel that we've been able to remove the business constraints," says Garza. "We have an environment that not only meets current needs but will scale to meet future needs as the business demands around water delivery, flood protection, and other program areas grow. This platform isn't about IT. It's about our science and engineering users—enabling them to do more and do it as quickly as they need to. It's about technology as a business enabler."

Chung reinforces Garza's point. "The DWR high-end modeling environment is an excellent example of collaboration between our specialized business domain expertise, technology, and private-sector providers and collaborators," says Chung. "This solution was a great blending of science, engineering, technology, and people. It all works together to deliver what is needed for our program areas."

Printed in USA 0614/LJ/TDA/XX/PDF ♻ Please Recycle 329576-001US