

R Systems NA, Inc. Case Study

R Systems uses vSMP Foundation for SMP and Cloud to provide the most flexible, high performance system application services to its customers



ScaleMP's vSMP Foundation platform for server aggregation solution provides low cost, dynamic virtualization solutions for flexible high performance workloads



"vSMP Foundation for SMP and vSMP Foundation for Cloud were definitely tools I wanted in my arsenal," said Keller. "In my world, our customers aren't looking for capabilities from a system that's smaller than what a single server can provide. They're usually looking for creative ways to combine the power of dozens or hun-

dreds of machines, and they don't want to worry about understanding the underlying system. ScaleMP hides a lot of that complexity and make it look like one single computer even though it has 128 processing cores. Additionally, the support team at ScaleMP is incredible and it's great to work with technology visionaries who see where our pain points are and can make them disappear gracefully. I look forward to continuing to use their products as we acquire more large memory systems. In many ways, it's up to the researchers to catch up to ScaleMP and figure out how they can optimize their programs now that they have this much memory available. We're here to make sure they have access to those systems, from proof of concept all the way through deployment."

BUSINESS BENEFITS

- Reduced capital expenditure outlays due to the avoidance of purchasing additional SMP systems
- The ability to provide customers with systems that specifically meet their needs so they only pay for what they use
- Shorter wait times and lower costs for research customers
- Research customers can avoid the unnecessary recoding of applications

Objective:

R Systems is a service provider offering cloud and traditional computing resources. The company continually seeks out innovative, leading edge technologies to provide researchers with the best supercomputing resources available at a significantly lower cost than in-house systems. To do this, R Systems needed to find technology that provided increased flexibility when dealing with its hardware resources and unforeseen demand. The company also needed to be able to isolate and secure research organizations working in clusters, so includes this as a requirement for all new technology implementations.

Approach:

After learning about ScaleMP from a Beowulf cluster list, Greg Keller, principal of R Systems, initially deployed vSMP Foundation for SMP since it was the only solution available that provided full access to all of the RAM and any processor in a given system. Since then, R Systems

has also adopted vSMP Foundation for Cloud because of its ability to dynamically allocate isolated virtual SMP systems that meet the exact technical specifications of its customers.

Customer Background:

R Systems is a service provider offering cloud and traditional computing resources to the commercial and academic research community, including the oil and gas, manufacturing and bio science industries. R Systems provides its customers with dedicated hosting, shared systems, virtual private clusters and support for off-site/remote facilities. The company's service offerings and simplified business model allow it to provide the power of world-class computing — previously available only to elite researchers and corporations with massive funding — at a fraction of the cost of other supercomputing resources.

IT IMPROVEMENTS

- A reliable, high-performance, scalable cloud compute infrastructure
- The ability to deploy private symmetric multi-processing (SMP) systems with specific compute requirements on demand
- The ability to deploy multiple private SMP systems on a given hardware system

A dedicated team of six employees manages the lifecycle of these high performance computing technologies for R Systems customers – acting as technology advisors and helping meet the unique needs of each researcher. The typical researcher comes to R Systems for a hosted SMP system. By and large, these researchers seek cloud resources because their in-house IT team is unable to meet resource demands or does not have the appropriate facilities. These researchers are generally not particular about the underlying platform used. Additionally, R Systems works closely with business managers who want to stay on budget and IT departments looking to partner with the company to implement technologies in-house.

Prior to ScaleMP, R Systems relied entirely on parallel file systems and standard distributed Beowulf clusters. On the distributed computing/traditional cluster side, R Systems uses Infiscale's Perceus, a RAM-based root file system. The company also runs both Windows and Linux and is hardware, OS and technology agnostic.

Customer Challenges

One of the ongoing challenges R Systems faces is providing cutting-edge tools for researchers that help them avoid recoding applications and allow them to focus on their research rather than the underlying technology. Before ScaleMP, R Systems had a number of researchers dealing with scale out issues and rewriting their OpenMP code as an MPI application because a single server was not big enough.

“One of the sayings in the industry is that hardware is temporary, software is forever,” said Keller. “Companies and researchers don’t want to spend a lot of time rewriting software to deal with a particular hardware architecture, because in three to five years, it will all have changed. Heavily optimizing for any specific hardware platform is futile in the long run.”

Another key challenge R Systems faced was the need to isolate each company, especially those that were potential competitors. R Systems wanted each enterprise and organization to have access to a remote and private system. Previously, if R Systems had a 16 node SMP system and two customers wanted to use it, they would be placed in a queue to use the system rather than being able to use it at the same time.

“I’ve got two data centers filled with computers and at any point in time, some subset of them will be partitioned off and carved up for different users of different companies,” said Keller. “When a customer shows up and says, ‘I really need an SMP system today,’ it’s very hard for me to predict which systems will be available on that day.”

A related issue was that the single 16 node system could only be used by one customer at a time due to confidentiality issues. Because R System was unable to dynamically configure the larger system, they would have to build a queue of customers that may only need four or eight node systems, and have them use the larger system one at a time – resulting in system underutilization.

Additionally, R Systems provides customers with more than just server resources as a commodity (like those provided from Amazon Web Services) and as a service provider offers more interaction and coordination to keep resources running. When a customer shows up and needs a system immediately because its own in-house resources are down, the company must quickly partition what the customer needs to get through the crisis. R Systems needs to have the right tools in place to be able to manage their internal cloud and these service-level expectations.

Solution Selection Process:

R Systems learned about ScaleMP two years ago from the Beowolf cluster list. From there Keller decided to learn more about ScaleMP's offerings and, once he understood the technology, he adopted vSMP Foundation for SMP in July 2008. He chose ScaleMP because it was the only solution that gave them full access to all of the RAM and any processor in a system. There were projects like MOSIX and others that allowed a single system image, but any given application could only consume as much RAM as was in the server on which the process ran. In addition, a vSMP Foundation for SMP-enabled system can run both OpenMP and MPI codes unchanged. Regardless of how an application is coded it will run with the same expected high performance manner natively.

R Systems initially deployed a four node system using vSMP Foundation for SMP and ultimately upgraded to a 16 node system. R Systems also uses vSMP Foundation for Cloud, which makes managing resources on-demand much simpler – so that R Systems can create the right size system for the right customer at the right time.

“Our approach is to provide customers access to the hardware, OS or applications at whatever level they want to take over responsibility,” said Greg Keller, principal of R Systems. “Many of our customers just want a system that will run MPI code and their applications. They don’t want to understand the OS or network or anything behind it. However, other customers want root access and they want to be able to change things at the OS level. There’s not a lot of interaction we have on a day to day basis with the system until the customer or group of hardware transitions to the next use case.”

With vSMP Foundation, R Systems gains the following key benefits:

- **Avoids recoding of applications:** vSMP Foundation allows R Systems’ customers to continue to get their research done instead of taking months to convert their application over to a distributed memory system. ScaleMP avoids that recoding effort so that researchers do not have to try to optimize for distributed clusters today, which a couple of years from now may be obsolete. Because vSMP Foundation supports both OpenMP and MPI codes equally, software doesn’t have to be rewritten in order to take advantage of SMP system performance.
- **Isolates systems for security purposes:** vSMP Foundation enables R Systems to carve up their resources into smaller virtual SMP systems if needed – allowing two customers to work on the same system and yet have complete isolation and autonomy.

- **Eliminates underutilization of systems:** vSMP Foundation for Cloud enables R Systems to use systems that would otherwise be left idle so that an SMP system is always ready in case a customer needs it. Now, R Systems can use any node as part of a standard cluster and, as customer demand requires, can dynamically grab nodes that are available and pull them up as a virtual SMP system – making it easier to manage customer expectations and not underutilize systems.
- **Enables more flexibility:** Instead of customers choosing between larger systems than actually needed because the configuration is locked in at 16 nodes, or figuring out how to squeeze their application into something smaller vSMP Foundation for Cloud helps R Systems avoid those painful decisions and keep resources much more fluid.
- **Customers have a lower overall cost and only pay for what they use:** vSMP Foundation for Cloud can allocate systems needed and move R Systems' customers through the hardware queue more quickly. Because there's more flexibility in terms of what systems R Systems can pull into a virtual SMP at any given point in time, they are able to give a customer an SMP at the size they need for the project they want to run today, cutting down costs significantly. Customers can now pay for what they need and R Systems can run SMPs simultaneously rather than serially. In the past, customers sometimes paid up to four times more just to have access to more RAM than can fit in a single box.
- **Reduces complexity:** In most big data applications, the file system capability is quickly overrun and organizations must turn to more expensive file systems. Additionally, these file systems add a layer of complexity. Because of the way vSMP Foundation pulls everything into one operating system, all that complexity evaporates. In certain circumstances, the money customers save in file system support would pay for an SMP machine.