

WHITE PAPER

## **New Product Review: Self-Contained Computer Room in a Shipping Container from Sun Microsystems**

***By Robert N. Sullivan, Ph.D., Vince Renaud P.E., Jonathan G. Koomey, Ph.D.  
and Bruce A. Taylor***

*Authors' Note: This New Product Review is based on field notes resulting from four site visits to observe and inspect Sun Microsystems' Modular Datacenter S20 (formerly Project Blackbox). As always, the Uptime Institute welcomes critical review and comment on all its published papers. Comments of merit may be included and attributed in future updates, with the permission of the reviewer. Send all review comments to [opinion@uptimeinstitute.org](mailto:opinion@uptimeinstitute.org).*

## Background and Introduction

Just over a year ago, Sun Microsystems announced a priority developmental concept under its Eco-Responsibility initiative code-named Project Blackbox (PBB). The first product release is now ready for sale as the Sun Modular Datacenter S20 (Sun MD). The Sun MD entails the design and assembly of a high-performance, high-density computer room housed within an industry-standard twenty foot long shipping container that has been structurally upgraded for Sun's design-intended use. While Sun uses the term "datacenter" to describe this product concept, the Uptime Institute feels "computer room" is more accurate because the customer must supply power and cooling from external site infrastructure systems (Uninterruptible Power Systems, Chillers, Engine Generators, switchgear, etc.) for the Sun MD to operate. Whether described as a data center or as a computer room in a box, the Sun MD provides an integrated and pre-proven high-density solution for customers wanting to deploy a lot of processing power in a hurry.

The Uptime Institute has conducted four technical inspections of PBB before it became the Sun MD. On June 20, Robert "Dr. Bob" Sullivan and Dr. Jonathan Koomey visited the prototype on the Sun campus in Menlo Park, CA. The host for this initial tour was Subodh Bapat, Vice President and Distinguished Engineer, Eco-Responsibility Office, Sun Microsystems. On July 11, 2007, Jonathan Koomey assisted the Sun team in conducting energy measurements on the prototype Black Box at Menlo Park. On October 23, Dr. Bob Sullivan and Bruce Taylor visited a manufacturing facility in the Denver area to inspect the latest version of PBB to note the progress made over the summer. Finally, on December 13, 2007 Vince Renaud visited the manufacturing site to examine concurrent maintainability and fault tolerance. Mark Monroe, Sun Director of Sustainable Computing, and Tim Jolly, Sun Senior Staff Engineer, were the hosts.

The purpose of this New Product Review is to verify, authenticate, and render technical opinion on the design and manufacturing specifications and on the product marketing claims made for the Sun MD. Its purpose is also to help Sun and the Institute appropriately advise their prospective customers on what they need to know and do to ensure the proper installation and integration of the Sun MD with their critical computing environments to meet performance requirements for computing availability, reliability, and resiliency.

In keeping with full disclosure, a key aspect of the Institute's assignment will be to advise Sun on its developmental roadmap for the Sun MD, particularly as it relates to compliance with each level of the four Site Infrastructure Tier Classifications for determining fault tolerance and concurrent maintainability.

The Institute's first observation is that the Sun MD concept is dependent upon the customer providing external sources for both power and cooling. It is important that customers recognize that the Tier functionality of the Sun MD can never be greater than the external site infrastructure equipment and systems to which it is connected. Many potential customers currently do not have a clear understanding of the concurrent maintainability or fault tolerant capability of their site. As such, there may be an expectation that

the installation of the Sun MD may be expected to deliver a level of concurrent maintainability or fault tolerance that will simply not be attainable due to the capacity, topology, or configuration of the customer's own external power, cooling or other site infrastructure support systems. In other words, if the source of power and cooling is non-redundant, then the Sun MD can only be expected to operate at that functionality level. If the source of power or cooling must be taken down for maintenance, then so must the Sun MD.

The Institute's next two observations are interrelated. Today, the term 'out of capacity' typically means being out of power and cooling (kW) capacity. This is a change from the recent past when being out of capacity typically meant being out of computer room floor space. For sites without sufficient power or cooling capacity, the Sun MD concept by itself does not offer relief without the simultaneous addition of supplemental power and cooling capacity. However, many sites also lack the necessary engineering skills, learning/experimentation time, and the physical space (in vertical as well as horizontal dimensions) to plan for and handle high-density heat loads. This leads to the next observation. Even sites having adequate power and cooling and computer room space may be out of "high-density" capacity in their computer room. They may not want to develop the required technical engineering expertise or their remaining space may physically be unable to achieve the cooling or weight densities that the Sun MD achieves in 160 square feet. For this reason, the PBB concept and the Sun MD product does offer a unique, integrated, and proven high-density cooling solution for heat loads of up to 25 kW per rack. The Institute expects to find field installations where the wall of the existing computer room was opened and the entire Sun MD shipping container was swung into the existing building and the wall re-closed.

The Institute undertakes this work with the full knowledge that the Sun MD is not the first or only effort at using a standard shipping container configuration for a modular, portable data center. APC was probably the first to develop such a transportable, server-room concept with full internal power and cooling generation, but stopped short of fully bringing it to market as a product offering. More recently, high-performance computing server developer, Rackable Systems, introduced their Integrated Concentro Environment (ICE) Cube. Undoubtedly, more such efforts will come in the future from other developers. The approaches taken by Sun and Rackable are markedly different: each deserves examination and consideration on its own merits for the intended purpose of the customer. It must be stressed that this New Product Review is not a product comparison. Instead, it reviews design and developmental progress on the PBB concept only for those prospective customer/user organizations who are now, or soon will be, evaluating the suitability and applicability of the Sun MD for their critical computing environments.

## Design Realization of the Sun Modular Datacenter S20

Sun has developed and taken great strides in the full realization of the concept of a computer room in a box—the box being, in fact, an ISO-standard 20-foot shipping container. Sun’s intended uses for this device include:

- Additional growth for an existing data center requiring more high-density floor space capacity in a hurry
- Relocation of non-critical, but high density applications from a constricted computer room
- Emergency and disaster recovery computing capacity brought to the site (when long-distance communications may be difficult)
- Modular growth for new applications or incremental expansion of existing applications
- Temporary computing capacity to bridge growth needs between an existing facility and the commissioning/provisioning of a new or expanded data center
- Rapid deployment of data center capacity in a remote location
- Location of a server computing facility within a regional grid offering less expensive, more reliable or available power

Eight 19 X 32-inch standard computer-room equipment racks with 40 rack units (RU) per rack have been placed within this standard container. Seven of the racks provide 280 RU for housing the customer’s IT equipment such as servers, storage, and switches. The eighth rack (called the infrastructure rack) reserves 15 RU for Sun MD monitoring, management, and control equipment, and provides 25 RU for customer networking equipment.

Four equipment racks are lined up along each side of the container with an access aisle down the middle of the unit (for a total of eight racks, four on each side). Power and communications cabling run in separate overhead trays. A design feature of the Sun MD is that it’s use of standard equipment racks means that it can be configured with any Sun or third-party servers, storage, or network equipment that conforms to the rack specifications. As a device, the Sun MD unit can be delivered to the customer site without being provisioned with the IT equipment payload. It also can be fully provisioned and configured by Sun, at the customer’s discretion.

Fully configured communications network interface panels with up to eight LC fiber (single or multi-mode) or RJ45 copper wire communications ports are mounted on each side of the Sun MD. Optionally, the ports on a network panel can be replaced with a bulk cable pass-through (accommodating up to 18 cables) for direct hook-up to the internal equipment.

The Sun MD is designed to operate in environments where the temperatures at the outside surface skin of the container range from -20°F to as high as 130°F, and where the relative humidity is up to 100 percent. The box is designed for inside operating temperatures ranging from 50°F to 95°F, with a relative humidity between 20

and 80 percent. A filter bank at one end of the unit maintains particulate cleanliness within the Sun MD.

A dehumidifier unit is mounted in the infrastructure rack to handle moist air introduced into the unit when the doors are opened in a humid environment. At this time, no capability exists in the design to introduce moisture into the air in dry environments to maintain a fixed relative humidity level. This problem is compounded by a personnel safety requirement to have a container door open whenever anyone is inside the container. (In future product releases, the container doors may be modified so they can be opened from the inside. This change will allow service personnel to work inside while the container is kept sealed.) For the current product release, Sun recommends that customers install a vestibule at the container doors to seal the container and prevent the introduction or loss of moisture. The vestibule will also provide protection for personnel and equipment during inclement weather and would provide space for additional humidification equipment.

## Power and Cooling Infrastructure

**Power:** The Sun MD is capable of supporting up to 178 kW of customer equipment load in its eight cabinets. While the device is designed to cool up to 200 kW of load, 22 kW are reserved for the Sun MD’s own infrastructure and external solar loading in high-temperature environments. The Sun MD internally distributes conditioned or unconditioned power depending upon what is supplied by the customer. Fault-tolerant power was available in the assembly-unit observed during the October 23 visit. Two customer power feeds can be utilized to supply redundant power to the equipment racks. Within the Sun MD, dual power distribution panels are located on opposite sides of the unit. Each power distribution panel supplies two branch circuits to each general purpose rack and one circuit to the infrastructure rack. Thus, each of the general purpose racks is supplied by four such circuits and can deliver up to 25 kW of dual redundant capacity. The monitoring and network infrastructure rack is supplied by two such circuits and can deliver up to 12.5 kW of dual redundant capacity.

According to a Sun technical note (confirmed by visual inspection of two different units during assembly on October 23), the Sun MD operates on three-phase power and comes in two versions: a low-input voltage version (190 to 220 VAC at 50/60 hertz) for North America, Taiwan, and Japan, and a high-input voltage version (380 to 415 VAC at 50/60 hertz) for the rest of the world. Amperage needed for a Sun MD implementation will vary depending on the payload in the Sun MD and the external voltage that is supplied to the Sun MD. For example, the full 178 kW of customer IT payload, along with the Sun MD infrastructure, will require 600A service at 208V, but will need less than 300A service at 415V. The output voltage (power distributed to the racks) is the same as the input voltage. In the low-voltage Sun MD, the branch circuits delivering that power are 60A. They are 32A for the high-voltage Sun MD. Regardless of whether the Sun MD is a low-voltage or high-voltage unit, the power distribution units (PDUs) in the racks ensure that the voltage delivered to the equipment in the racks always ranges between 190 to 240 VAC, depending on the input voltage.

At the demo PBB installation observed in June in Menlo Park, a single 480V power feed coming from the building electrical system was transformed down to 208V. The Sun MD is not designed to house a UPS, battery, or engine generator as part of its own internal infrastructure. Third parties may offer these separately.

**Cooling:** Cooling is provided by an external, customer-supplied, chilled water source. This can come from either a stand-alone air-cooled chiller package or from the existing building mechanical plant.

Two redundant chilled-water attachments are provided, each capable of (but not limited to) receiving the 60 tons of sensible cooling required at maximum 200 kW load. The PBB installation that was visited in December had only a single chiller attached and a single cooling loop.

Cooling is distributed via heat exchangers (HXs) placed between each pair of racks. Each HX contains a cooling coil and 10 variable-speed fan units. Each fan in a HX has a maximum capacity of 1,100 CFM, but under normal cooling situations the airflow volume is scaled back to 100 to 250 CFM. Each fan is always running. The fan speed is controlled by feedback from a series of resistive temperature devices (RTDs) located at the input of each heat exchanger. In many cases, the internal fans within the servers themselves are providing more than sufficient air flow through the HX cooling coil unless the air path results in significant backpressure. The fact that 20 percent fan operation is sufficient to maintain cooling is an indication that the server fans are doing the real work. This is a major energy savings for the customer because the HX fans run at only 20 percent. In other high-density computer room cooling solutions, the cooling unit fans always run and always consume energy. Beside the RTDs on the HX, an ambient temperature sensor is located in the infrastructure rack.

**Other Environmental Systems:** A fire-and-smoke detection system backed up by a clean-agent gas-fire suppression system is included in Sun MD units now in production. The first line of defense is a Very Early Smoke Detection Apparatus (VESDA) air-sampling system, backed up by two conventional smoke sensors that are tied into the gas suppression system.

## Fault Tolerance and Concurrent Maintainability

The Sun MD is a Tier II solution with some features of Tier III and Tier IV (Tier ratings refer to the Institute's Tier Rating System – see *Tier Classifications Define Site Infrastructure Performance* by W. Pitt Turner IV, P.E., John H. Seader, P.E., and Kenneth G. Brill for further information about the Institute's Site Infrastructure Tier Classification). Sun is considering alternatives for achieving a higher Tier rating in future product releases.

The Institute has some concerns regarding the ability of the Sun MD to allow customers to do things they would normally expect to do in a Tier II infrastructure. Customers will periodically need access to the IT hardware within a rack for maintenance or changes. During this time, proper environmental cooling must be maintained without the customer needing to shut off the IT hardware. Accessing the hardware currently requires separation of the rack

from its direct coupled heat exchanger and moving the rack away from the heat exchanger and into the central aisle. Once separated from the dedicated heat exchanger, rack cooling comes from the overall room. At low load densities, this should not be a problem. As densities rise, some racks or locations within racks during maintenance may be more sensitive to inadequate airflow and cooling than others. Assuring adequate cooling redundancy during IT hardware maintenance or heat exchanger or cooling loop maintenance is something Sun needs to validate with worst-case testing.

## Energy Efficiency

The Institute calculated the energy efficiency of the Sun MD by comparing the fan energy required for a computer room using 30 ton computer room air handling units with 10 horsepower motors driving blowers. These large motors are required to push cold air around within the computer room (a task which the Sun MD allows the existing internal fans within the servers themselves to perform at no or low incremental energy cost). Three 30 ton units would be required, 2 for capacity and one for redundancy for a total of 30 horsepower running. Assuming the motors are only loaded to 75 percent, the kW consumed per hour is 22.5. At full load, the fans in the Sun MD are estimated (with this estimate to be confirmed during February 2008 product testing) to consume 1.2 kW for a hourly saving of 21 kW. This equates to an annual energy savings of 184,000 kWh or 130 metric tons of CO<sub>2</sub> emissions.

## Recommendations and Questions for Sun Microsystems

Dr. Bob Sullivan suggests that Sun or Sun's customers may unintentionally create problems with local code enforcement personnel where the unit is intended to be installed by referring to the Sun MD as a "shipping container". By contrast, referring to the unit as something like a "self-contained server computing enclosure" may ease the process of obtaining necessary permitting to install the unit in areas other than those zoned as industrial. In some remote locations, local codes are not a factor, while obtaining communications or power and cooling connections to existing facility infrastructure may require siting the Sun MD very close to other corporate facilities which are likely to have strict zoning.

State and local code jurisdictions within the United States are quirky, and internationally they can be confounding. Customers seeking to site the Sun MD should research what nomenclature on a permit application will raise the fewest red flags to speedy approvals. Sun provides customers with planning documentation that helps them work with local code officials. The planning documentation addresses building, electrical, fire, and plumbing issues and codes.

Vince Renaud suggests that Sun might consider providing segmentation of the cooling loop piping at each HX or using two cooling loops for concurrent maintenance capability. Vince was also concerned about the redundancy, concurrent maintainability, and fault tolerance of the HX fan controllers, how these controllers relate to the Environmental Monitoring System (EMS), and what happens when maintenance is performed on the HX or the EMS.

Questions asked by Dr. Bob Sullivan and Vince Renaud and answered by the Sun MD technical team:

**Q:** Will Sun measure the airflow in each rack when it is separated from its dedicated heat exchanger to confirm concurrent maintainability of IT equipment, fan fault sequences and maximum operating temperatures under maximum heat load conditions?

**A:** Sun's measurement of airflow rates while running diverse server workloads indicates that the fans run at about 20 percent capacity most of the time. When an IT rack is separated from its dedicated HX, the HX remains in place and continues to provide cooling although the dedicated, direct, closely coupled cooling air connection between the rack and HX is broken. The servers in the rack being maintained will continue to pull the air they need from within the overall container and discharge the heated air back into the container. If the ambient temperature in the container becomes too hot, the Environmental Management System will speed up all fans in all HXs until the temperature drops. By calculation and by observation, there is more than sufficient fan and cooling capacity to maintain room ambient temperature when a rack is in the central aisle. However, Sun will confirm this by extensive testing at full load for each rack and RU location within a rack.

**Q:** At maximum load density, what happens to IT hardware cooling in the rack coupled to a HX if the HX or its associated fan controller fails? How does the Environmental Management System work in conjunction with the HX fan controllers? What happens if the EMS fails or needs to be maintained? How does the Sun MD deal with individual fan failures?

**A:** An HX will lose cooling capability only if it develops a leak. If that happens, the EMS will shut down both the chiller water supply and the electrical system in the entire container. It should be noted that the pipes and cooling coil in an HX are stationary plumbing with no moving parts and are extraordinarily reliable. Sun estimates their failure rate to be less than once every 80 years.

A separate fan controller is associated with each HX. The fan controller itself is not redundant. However, it is powered by redundant power sources, so a single power failure will not cause all of its fans to stop. If the controller itself fails, the default failure mode is to keep the fans at the last operating point, i.e. still spinning. Unless, of course the failure is one which causes the controller to turn the fans off. It's possible, just not very likely. This is a mature industrial controller system that is designed to fail in a "static" operating condition.

There are five airflow circuits at different RU elevations within the Sun MD. Each circuit is 8U in height and makes a full loop through the racks and heat exchangers at its specific elevation. Each airflow circuit has two fans and one temperature sensor per HX for a total of 16 fans and 8 temperature sensors in the circuit. The EMS continuously monitors the air temperature of the airflow circuits as they pass through each HX. If the EMS senses an increase in temperature at one of the sensors, it will instruct all HX fan controllers to increase the speed of their fans in the corresponding airflow circuit.

If a single HX fan controller fails in a way that causes it to turn off its fans, the temperature of the air in some air circuits may increase. If that happens, the EMS will sense the temperature rise in the affected air circuits and will compensate by instructing the other HX fan controllers to speed up the fans that participate in those air circuits. Air will continue to flow through the HX with the stopped fans, and thus will continue to be cooled by the coil in that HX. There will be no loss of cooling in any of the air circuits as the result of a fan controller failure.

If the EMS itself fails, the fan controllers for each HX will continue to run their fans at the last fan speed set by the EMS. The Integrated Management Server in the Sun MD will notice that the EMS has failed and send SNMP alerts and Emails to notify the need for service. Under the current design, if the EMS must be powered off for maintenance or replacement, communication will be lost to the HX fan controllers, and the controllers will stop the fans. Sun will review this design in a future release of the Sun MD.

The Sun MD does not monitor the operation of the individual fans. If a single fan in an HX should fail, the EMS may sense this as an increase in the temperature of its particular air flow circuit and increase the speed of the other fans in the circuit. However, due to the extensive fan redundancy in the Sun MD, the temperature may not rise significantly until several fans fail. Because the fans typically operate at less than 20 percent capacity, more than half the fans in the same circuit would have to fail before Sun MD would be unable to move enough air to cool the equipment in that circuit.

*[Institute Observations Relative to the HX, fan controller, and EMS]:*

- The Institute does not consider reliability calculations (a piping failure once every 80 years) to be a substitute for having concurrent maintenance capability or fault tolerance. If a leak occurs, the system must be shut down. This is one of the major reasons why the current product design did not achieve a rating above Tier II.
- The HX fans are not easily accessible to determine if they are actually turning. Replacement of an HX fan currently requires separating the equipment rack from the HX. Because of cable slackness issues, the equipment rack may be required to be powered off in order to get to the HX unit. It is likely that Sun will review this issue in a future release of the product.
- Loss of an HX controller or even the EMS will not affect the internal fans running within each server. These internal fans will continue to force air through the HX, but the discharge temperature may rise if the air flow is not sufficient depending on heat load.

**Q:** With the container doors required to be open for personnel safety while someone is inside, the Sun MD must provide cooling for the computer load and in addition provide cooling and dehumidification for air infiltration through the open doors. How are early customers handling this problem?  
**A:** Customers have addressed serviceability in different ways. Some customers are using an external vestibule with a passage door so the

container doors can be kept open while still isolating the Sun MD internal environment from the outside environment. Other customers have chosen to implement a deferred service model in which they don't open the doors for service work until the outside climate is suitable.

*[Institute Observations Regarding Service]:*

- Future product releases are likely to have the container doors openable from the inside thereby allowing personnel to easily enter and exit the container. With this change, the container doors can be kept shut during service and maintenance.

**Q:** How much cooling capacity is required?

**A:** Only at a full-power IT payload (178 kW) is the full 60 tons of cooling capacity required. Smaller payloads require smaller amounts of cooling capacity. When there are no people inside, the product has no source of latent heat load (e.g., no people and very limited air infiltration). The number of tons of required capacity needs to be adjusted up or down based on local conditions and whether a vestibule is installed. The 60-ton capacity covers the heat generated by the full-power IT payload, plus the heat generated by other Sun MD internal equipment, and the heat generated by solar loading. It does not include capacity adjustments that should be made for site-specific conditions such as altitude, ambient temperature, operating margin for the cooling equipment, and outside air leakage when the doors are open. Sun recommends that customers use 60 tons only as a rough estimate for budgetary planning, and that they work with an HVAC engineer and contractor to determine the actual chiller capacity that will be needed for their specific implementation of the Sun MD.

**Q:** To perform this investigation, the Institute worked more from verbal descriptions and engineering documentation than product marketing materials. How will the final specifications and product marketing materials for the January, 2008 product release vary from what the Institute has seen and observed?

**A:** The final specifications, product marketing materials, and user documentation for initial product release reflect the Sun MD seen by Dr. Bob Sullivan and Vince Renaud in the October and December timeframe.

## Conclusion

In the opinion of the Institute, Sun has materially achieved its stated primary and early-stage design objectives for the Sun MD. If a customer has the necessary power and cooling capacity to support the Sun MD, the Institute believes that released product represents a new and significant contribution to the realization of a flexible, scalable, lockable, modular, easily transported and installed computer room floor space which is capable of cooling extremely high-density IT hardware loads. As such, the Sun MD potentially offers a strong economic and operational benefits proposition for a variety of site installations and business needs. Through its four site visits over a six-month period, the Institute has observed considerable progress in the development of this device and an open willingness on Sun's part to further perfect the design. Data center operators facing density or capacity issues or new applications and

projects requiring new server capacity owe it to themselves to investigate Sun's Modular Data Center for its suitability and applicability to their present and near-term business needs. The Institute is prepared to answer any questions or advise interested data center operators as to the suitability and applicability of Sun Microsystems's Modular Data Center as a solution (whether temporary or long-term) to their specific requirements and how their existing site infrastructure systems would rate in the Institute's Tier Classification system.

## About the Authors

**Robert "Dr. Bob" F. Sullivan, Ph.D.**, had a 32-year engineering career at IBM. He was part of the IBM team that first identified zinc whiskers. Dr. Bob originated the Hot Aisle/Cold Aisle concept in 1993. For the Institute he is a staff scientist and has performed detailed research on how computer room cooling actually works and teaches the Institute's Seminar on High Density Cooling. For Upsite Technologies he conducts computer room KoldCheck diagnostics and KoldTune remediation services. For ComputerSite Engineering Dr. Bob delivers raised-floor engineering services, helping clients exceed gross computer room densities of 50 Watts/ft<sup>2</sup>.

**Jonathan Koomey, Ph.D.** is a Project Scientist at Lawrence Berkeley National Laboratory and a Consulting Professor at Stanford University, as well as an Institute Senior Fellow. Dr. Koomey is one of the leading international experts on electricity used by computers, office equipment, and data centers. His study, *Estimating Total Power Consumption By Servers in the U.S. and the World*, published in February, 2007 is already a standard reference work for anyone working in the field of data center energy. He has also published extensively on critical thinking skills. He holds M.S. and Ph.D. degrees from the Energy and Resources Group at the University of California at Berkeley, and an A.B. in History of Science from Harvard University.

**Vince Renaud, P.E.** is a Principal of ComputerSite Engineering and a Distinguished Fellow and Certification Authority for the Uptime Institute. As a consultant, Mr. Renaud has provided planning, design, construction, operation, and maintenance of mission critical facilities for the Department of Defense on a world-wide basis. While managing this key national facility, he implemented a Computerized Maintenance Management System resulting in demonstrated increased infrastructure availability. He is also a co-author of the Institute's *Operational Sustainability and Its Impact on Uptime Performance, Investment Value, and Resiliency* White Paper, among others, and a faculty member of the Institute's Site Uptime Network® and Seminar Series.

**Bruce Taylor**, Chief Strategy Officer of the Institute, is a 20-year executive of Information Technology media as a writer, editor, publisher, conference producer, speaker, and consultant specializing in IT publishing and conferencing. He is a founder of the technology, design, and futurist conference PopTech, Knowledge Management World and ImagingWorld. In the early 1990s, Mr. Taylor served as publisher of International Data Group's (IDG) Canadian IT publishing franchise. For ITWorld.com, Mr. Taylor is a regular interview host for IT thought-leaders.

## About the Uptime Institute, Inc.

Since 1993, the Uptime Institute has been a respected provider of educational and consulting services for Facilities and Information Technology organizations interested in maximizing data center uptime. The Institute has pioneered numerous industry innovations, such as the Tier Classification System for data center availability, which serve as industry standards today. The Institute's 100 global members of the Site Uptime Network® represent mostly Fortune 100 companies for whom site infrastructure availability is a serious concern. They collectively and interactively learn from one another as well as from Institute-facilitated conferences, site tours, benchmarking, best practices, and abnormal incident collection and analysis. For the industry as a whole, the Institute publishes white papers, offers a Site Uptime Seminar Series and a Site Uptime Symposium and Data Center Design Charrette Series on critical uptime-related topics. The Institute also conducts sponsored research and product certifications for industry manufacturers. For users, the Institute certifies data center Tier level and site resiliency.

---

The Uptime Institute's publications are protected by international copyright law. The Institute requires written requests whenever the Institute's literature or portions of the Institute's literature are reproduced or used. The Institute copyright extends to all media—paper, electronic, and video content—and includes use in other publications, internal company distribution, company websites, and handouts for seminars and courses. For more information, please visit [www.uptimeinstitute.org/resources](http://www.uptimeinstitute.org/resources) to download a Copyright Reprint Permission Request form.



Building 100, 2904 Rodeo Park Drive East, Santa Fe, NM 87505  
505.986.3900 · Fax 505.982.8484 · [uptimeinstitute.org](http://uptimeinstitute.org)