

## Benchmarking big data



Cisco and Intel® partnering in innovation



The Transaction Processing Performance Council (TPC) launched TPCx-HS in 2014. The first—and to date, the only—vendor-neutral, industry standard big data benchmark, TPCx-HS measures the price, performance, and energy consumption of big data systems like Hadoop.

“Benchmarking provides an objective assessment of system performance and other variables,” says Raghunath Nambiar, distinguished engineer and chief architect of big data solutions for Cisco. “This is not only important for comparing and selecting technology options, but also for driving faster, less expensive, and more energy efficient systems.”

TPCx-HS is designed to measure both hardware and software, including:

- Hadoop Runtime
- Hadoop FileSystem API compatible systems
- MapReduce layers

It can also be used to assess a broad range of system topologies and implementation methodologies, in a technically rigorous and directly comparable, vendor-neutral manner.

### ESTABLISHING THE BIG DATA BENCHMARK

Cisco and MapR recently announced the first TPCx-HS results, setting the high performance benchmark for big data systems.

“We published not one but three results,” says Nambiar, “at 1TB, 3TB, and 10TB scale factors.”

The benchmark configuration includes the Intel® Xeon® processor-based Cisco Unified Computing System™ (Cisco UCS®) Integrated Infrastructure for Big Data, two redundant active-active Cisco UCS 6296 Fabric Interconnects running Cisco UCS Manager version 2.2, 16 Cisco UCS C240 M3 Servers running Red Hat Enterprise Linux Server 6.4, and MapR Distribution with Apache Hadoop.

“TPC benchmarks are highly regarded in the industry and we are pleased to establish these high-performance standards for Hadoop,” says Jon Posnik, vice president of business development at MapR. “These results validate the scale, performance, and availability of the MapR Distribution. We’re already seeing tremendous demand for the combined solution with Cisco to enable utility-grade Hadoop environments.”

### BIG DATA AND ANALYTICS VIRTUAL CONFERENCE

To learn more about how you can leverage your data to track down new opportunities, identify revenue streams, uncover hidden inefficiencies, and more, register at [UnleashingIT.com/virtualconference](http://UnleashingIT.com/virtualconference).

## Benchmarking big data



Cisco and Intel® partnering in innovation

### **BIG DATA REFERENCE ARCHITECTURES AND INFRASTRUCTURE BUNDLES NOW AVAILABLE**

Hadoop clusters, analytics software, and supplementary big data technologies all have one thing in common: A fundamental need for a rock-solid, tightly integrated infrastructure on which to operate.

“Big data is all about pulling a wide variety of data sources into a common environment,” says Raghunath Nambiar, distinguished engineer and chief architect of big data solutions for Cisco. “That requires alignment and orchestration spanning compute, storage, networking, and application environments.”

Cisco and its big data partners offer just that. Through the Intel Xeon processor-based Cisco Unified Computing System (Cisco UCS) Integrated Infrastructure for Big Data, organizations have access to a set of reference architectures and solution bundles that are designed for big data deployments. Validated designs are available for leading Hadoop platforms and analytics software.

“The architectures can be used as is or customized to meet specific business requirements,” says Nambiar. “We also have a number of single SKU bundles that make it easy to order and deploy big data environments.”

The new Cisco UCS Director Express for Big Data enables end-to-end cluster management and monitoring, he adds. And Cisco Nexus® 7000 and 9000 switches in tandem with Cisco® Application Centric Infrastructure (Cisco ACI™) deliver best-in-class scalability.

“With these reference architectures, solution bundles, and [Cisco] UCS Director Express, organizations can focus squarely on innovation and business outcomes,” says Nambiar. “They don’t have to worry about piecing together the infrastructure or optimizing it for the Hadoop and analytics software they plan to use.”

Read Nambiar’s blog at [UnleashingIT.com/m](http://UnleashingIT.com/m) to learn more about Cisco UCS Integrated Infrastructure.

This article first appeared in *Unleashing IT* Volume 4, Issue 4, and online at [UnleashingIT.com](http://UnleashingIT.com), available after subscribing at [UnleashingIT.com/Login.aspx](http://UnleashingIT.com/Login.aspx).

© 2015 Cisco and/or its affiliates. All rights reserved. Cisco, the Cisco logo, Cisco ACI, Cisco Nexus, Cisco Unified Computing System, and Cisco UCS are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [cisco.com/go/trademarks](http://cisco.com/go/trademarks). Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries.