IT Central Station
Unbiased reviews from the tech community

User Reviews of Oracle Exadata and HPE Vertica

A PEEK INTO WHAT REAL USERS THINK

2016

IT Central Station helps tech professionals by providing...

A comprehensive list of enterprise level Data Warehouse vendors.

A sample of real user reviews from tech professionals.

Specific information to help you choose the best vendor for your needs.

To read more reviews about Data Warehouse, please visit:
http://www.itcentralstation.com/products/comparisons/oracle-exadata_vs_hpe-vertica

Disclaimer

IT Central Station Ltd. does not endorse or recommend any products or services. The views and opinions of reviewers quoted in this document, IT Central Station websites, and IT Central Station materials (i.e. Content) do not reflect the opinions of IT Central Station. We make no guarantees about the accuracy, currency, suitability, or quality of the content, and we assume no responsibility for unintended, objectionable, inaccurate, misleading, or unlawful content made available by Content authors or other third parties.

ABOUT IT CENTRAL STATION

User reviews, candid discussions, and more for enterprise technology professionals.

The Internet has completely changed the way we make buying decisions. We now use ratings and review sites to see what other real users think before we buy electronics, book a hotel, visit a doctor or choose a restaurant. But in the world of enterprise technology, most of the information online and in your inbox comes from vendors but what you really want is objective information from other users.

We created IT Central Station to provide technology professionals like you with a community platform to share information about enterprise software, applications, hardware and services.

We commit to offering user-contributed information that is valuable, objective and relevant. We protect your privacy by providing an environment where you can post anonymously and freely express your views. As a result, the community becomes a valuable resource, ensuring you get access to the right information and connect to the right people, whenever you need it.

Use IT Central Station to:

- Read and post reviews of vendors and products
- Request or share information about functionality, quality, and pricing
- Contact real users with relevant product experience
- Get immediate answers to questions
- Validate vendor claims
- Exchange tips for getting the best deals with vendors

244 5th Avenue, Suite R-230
New York, NY 10001

www.ITCentralStation.com
reports@ITCentralStation.com
+1 (646) 328-1944

ABOUT THIS REPORT

This report is comprised of a comprehensive list of enterprise level Data Warehouse vendors. We have also included several real user reviews posted on ITCentralStation.com. The reviewers of these products have been validated as real users based on their LinkedIn profiles to ensure that they provide reliable opinions and not those of product vendors.

IMPORTANT NOTICE

Did you find this whitepaper helpful? At IT Central Station, our philosophy is “Give to Get”. Our active community and unbiased reviews are made possible by your participation and as such, we ask that you share your expertise with us as well. Please email reviews@itcentralstation.com and one of our community managers will be in touch with you shortly. You can choose to review anonymously or not and your company name will not be included in the review.

If you found this report and/or the reviews on IT Central Station useful, we would greatly appreciate your participation in giving back to our community.

Oracle Exadata

HPE Vertica

Vendor:

Oracle

Hewlett Packard Enterprise

Overview:

Oracle Exadata Database Machine is a modern framework, engineered to run databases and to scale out database servers.

The main benefit of Exadata is speed. It hosts operating systems, CPU memory, and hard drives. It runs all types of databases, including online transaction systems, processors and data warehouses, while solving poor performances of old database architecture.

The Oracle Exadata Database Machine features a simple and fast database storage system that protects and backs up your critical data. It accelerates data warehouse performance for faster access to business information. It is ideal for companies looking to build up their infrastructure from scratch.

For more information on Oracle Exadata Database Machine, visit Oracle.com

HPE Vertica is the most advanced SQL database analytics portfolio built from the very first line of code to address the most demanding Big Data analytics initiatives. HPE Vertica delivers speed without compromise, scale without limits, and the broadest range of consumption models. Choose Vertica on premise, on demand, in the cloud, or on Hadoop. With support for all leading BI and visualization tools, open source technologies like Hadoop and R, and built-in analytical functions, Vertica helps you derive more value from your Enterprise Data Warehouse and data lakes and get to market faster with your analytics initiatives.

To learn more about HPE Vertica Advanced Analytics, visit our website .

Sample Customers:

PayPal, EBS, Organic Food Retailer, Garmin, University of Minnesota, Major Semiconductor Company, Deutsche Bank, Starwood, Ziraat Bank, SK Telecom, and P&G.

Cerner, Game Show Network Game, Guess by Marciano, Supercell, Etsy, Nascar, Empirix, adMarketplace, and Cardlytics.

Top Comparisons:*

Teradata vs. Oracle Exadata:
Compared 20% of the time.

HPE Vertica vs. Oracle Exadata:
Compared 19% of the time.

Microsoft Parallel Data Warehouse vs. Oracle Exadata:
Compared 13% of the time.

Oracle Exadata vs. HPE Vertica:
Compared 18% of the time.

Teradata vs. HPE Vertica:
Compared 15% of the time.

Netezza vs. HPE Vertica:
Compared 12% of the time.

Top Industries:*	Financial Services Firm	18%	Media Company	20%
	Media Company	15%	Financial Services Firm	17%
	Comms Service Provider	9%	Comms Service Provider	12%
	Energy/Utilities Company	7%	Marketing Services Firm	9%
Company Size:*	1-100 Employees	12%	1-100 Employees	30%
	100-1000 Employees	14%	100-1000 Employees	23%
	1000+ Employees	74%	1000+ Employees	47%

** Data is based on the aggregate profiles of IT Central Station Users researching this solution.*

Deep Dive: Valuable Features

Oracle Exadata

Charles Kim, Founder and President at Viscosity North America:

My top 4 most important features of Exadata are: 1. Smart Scan, the ability to offload intensive SQL workloads to the storage servers. Queries are offloaded to the storage layer and only the result sets of relevant data are returned to the database server thus significantly improving performance. 2. Exadata Hybrid Columnar Compression, where we can compress data from 10x to 50x. Deploying databases on the Exadata can significantly reduce the amount of storage that is needed. 3. Exadata Smart Flash Cache automatically moves data between DRAM, flash and spinning SAS or SATA disks to provide best performance. 4. Virtualization is a new feature introduced to the X5-2 family. Now Oracle's Database Machine can be catered to large enterprise mission critical databases and can house smaller databases that need isolation and now even application servers together with the database. We can connect data intensive applications to the database over low-latency, high throughput infiniband.

HPE Vertica

Joe Gonzalez, BI Manager, Vertica ASE Certified DBA at a marketing services firm with 1000+ employees:

I have found great use out of many features, most notably the Management Console and the Database Designer. Many people with lots of experience creating table projections can get frustrated trying to optimize some complex queries, however, in Vertica, the Database Designer is normally a big help in these situations. You can feed it your problem queries and it will make projection design suggestions for you. The ability to have multiple projections on a table to work with different queries is a big bonus. With the Management Console, I am able to easily monitor the system health without having to check individual system tables. The ability to view running queries and cancel problem ones from the Management Console is a very nice feature, and one that I requested when we first started using version 7.0 which had an early version of the Management Console included.

Nitin Vengurlekar, CTO/Architect at Viscosity North America:

The most valuable feature is definitely the performance. However, it's also a platform for both performance and consolidation. Once you start leveraging Exadata with Database In-Memory and Multitenant, you've got a really solid platform for database enterprise computing.

Chris Bradham, Technical Director at a consultancy with 1-100 employees:

We run a mixture of OLTP and Data Warehouse with the majority being OLTP. With that being said, the feature that provides us the most performance gains is the Smart Flash Cache for the OLTP databases. The "offloading" capabilities provide the biggest performance gains for DW such as smart scans and storage indexes. There is a new security feature which allows disabling ssh to the storage servers which will make my security folks very happy.

Manish Nashikkar, ITA - Oracle Apps DBA at Tata Consultancy Services:

Smart Scan Smart Flash Cache High Availability Consolidation Robust Storage Computing

Saurabh K. Gupta, Database Development Leader at a tech vendor with 1000+ employees:

The following features of Exadata are the most valuable to me: Cell offloading Smart Scan Flash Cache Infiniband Fabric IO Resource Manager Hybrid Columnar Compression Storage Indexes

SoftwareDataArch072, Software and Data Architect at a software R&D company with 1000+ employees:

The compute and processing engine returns the queries fast and let us use our analysis resources in a better utilization. The concurrency got better in this version and we are able to run more queries and load concurrently.

Matt Williams, Lead Data Scientist Machine Learning at a financial services firm with 1-100 employees:

Columnar storage makes 'hot data' available much faster than a traditional RDBMS solution. Also, Vertica scales up quickly and maintains good performance.

Rodrick Brown, Development Operations/SRE at a software R&D company with 100-1000 employees:

HA Clustering Speed / Performance

Davide Moraschi, Senior business Intelligence consultant at Asociación SevillaUP:

Speed in query in general and specifically in aggregate functions on multi-million rows tables.

DBA511, Database Administrator (DBA) at a software R&D company with 100-1000 employees:

The auto-distribution to all nodes for fault tolerance and query performance was pretty amazing.

Yoojong Bang, Data Scientist at a media company with 100-1000 employees:

The fact that it is a columnar database is valuable. Columnar storage has its own benefit with a large amount of data. It's superior to most traditional relational DB when dealing with a large amount of data. We believe that Vertica is one of the best players in this realm.

Mohammad Rafey, Senior Principal Consultant at Viscosity:

Some of the features which are specific to Exadata, like histories, offloading and accurate processing. One of the customers we work with had mostly thousands of batch jobs, which used to take, 30 to 40 hours, with the Exadata, and re-architecturing some of those solutions on the Exadata, moving them. It turned out to be running in one hour, or less than one hour. It's a huge, huge performance gain. And in the end, the company realized the return on the investment they have made into Exadata.

Suresh Muddaveerappa, Senior DBA and Architect at a tech services company with 100-1000 employees:

The system is engineered to run both databases and enterprise applications unlike other engineered systems, which are either tailored only for databases or only for applications. The ability to consolidate all of the databases is a big plus. The SuperCluster provides best of both worlds – Oracle Exadata functionality along with virtualization at the firmware and the OS kernel layers. On the Exadata side the features most valuable are – Hybrid Columnar Compression (HCC) for both Data Warehouse and OLTP workloads, Storage Indexes for Smart Scan, ability to use Flash Cache for DB storage, and more.

Rob Visser, CIO Group Services at a insurance company with 1000+ employees:

It is fast and integrated. It combines hardware with the software that we use from Oracle. It's a complete stack, not only for the hardware but the full implementation and delivery, so it’s really good for IT.

Saswata Sengupta, Solution Engineering and Arcitect - Big Data, Data Science and Cloud Computing at a tech company with 1000+ employees:

Speed and performance: Vertica stands top among its peers in the MPP world, delivering unparalleled speed and performance in query response time. Its distributed architecture and use of projection (materialized version of data) beats most of its competitors.

SeniorDBA728, Senior DBA at a local government with 1000+ employees:

Compression / speed with highly complex queries

Victor Daniel Gonzalez, Vertica DBA at a tech services company with 1-100 employees:

Analytical features are amazing, the integration is wonderful.

Sivaprasad Kunchakuri, Database Admin at a tech services company with 1000+ employees:

Replication

seniorda505386, Senior Database Administrator; HP Vertica, MySQL, MSSQL, Cloud Ops at a tech vendor with 100-1000 employees:

Speed of columnar data.

Asim Singh Mahat, Sr. Software Engineer (Database) at a tech services company with 1000+ employees:

Query Performance.

Vijay Lunawat, Director - OSP/Engineered Systems at a tech services company with 1000+ employees:

Oracle Exadata provides very good performance and reliability. It's a highly reliable system and provides maximum performance.

Architect6dc7, Unix/Linux Platform Architect at a financial services firm with 1000+ employees:

It provides performance for our highly critical applications. They seem to be doing real well with keeping up on the performance and everything, especially what I have seen with the new chipset that is being offered. I can't really think of anything that they are not doing that they should be doing.

Simone Traversari, Team Leader at a tech consulting company with 100-1000 employees:

Oracle Engineered Systems are the best solution for Oracle and non-Oracle applications: hardware and software are combined together to ensure high availability and maximum performance. Oracle Exadata contains all the advantages of standard Oracle Database, but on top of that there are exclusive features such as IORM (IO Resource Manager) and HCC (Hybrid Columnar Compression). IORM (in combination with Database Resource Manager (DBRM)) permits you to isolate I/O requests to avoid I/O saturation. HCC provides different levels of compression to reduce space allocation and to decrease I/O.

Chris Gilliam, President at DataMAPt, Inc.:

The place where I've used it are mainly on the data warehouses with larger implementations, and having everything in one enclosure is really a lot nicer. Having the disk, all the memory, and all the networking, makes not just the implementation easier, but as you're running it, a lot simpler to administer. Because we're asking it to handle everything, it's a simple solution to administer and support. It also has plugins that makes it easier to monitor.

Andrey Krasnov, High FrequencyTrading Systems and Strategy Architect/Quant Trader at a financial services firm with 1-100 employees:

Fast inserts, queries way faster than in SQL Server.

Matti Heimonen, DW Admin at a hospitality company with 1000+ employees:

Scalability, query speed.

Kjetil Halvorsen, Management Consultant at a software R&D company with 1-100 employees:

Scale-out, analytical functions, ML.

reviewer505263, Principal Data Architect (BI) at a media company with 1000+ employees:

Bulk loading data using copy Fast analytical functions Ability to extend functionality

Srikanth Iyer, Senior Data Architect at a media company with 1000+ employees:

Analytic functions.

Jfm, Owner at a tech vendor:

Geospatial

Karla Marisol Guajardo, Infrastructure, Data Center and PMO Coordinator at a comms service provider with 100-1000 employees:

Manage big data fast and easy.

DataSci505212, Data Scientist at a tech vendor with 100-1000 employees:

User Defined Extensions Analytic Functions

Michael Abbey, Manager, Principal Consultant at Pythian:

It's pretty simple in my opinion. It's called out the box, and it's called support. Sometimes with the engineered systems based on the support agreement that you have with the vendor, sometimes you hear about problems with your system from them rather than your applications experiencing some issues, and you're having to log a support ticket with Oracle. I think that's huge.

Director662, Director, IT at a tech services company with 1000+ employees:

The performance.

SrDBA3602, Lead Oracle Database Administrator at a insurance company with 1000+ employees:

The most valuable feature would be Smart Scan because it gives us the ability to perform a large amount of processing on the storage side, cutting down on I/O tremendously. That's one of the big features that really drew us into jumping on Exadata. We also like having the ability to have a lot of flash storage where a lot of our data can be in-memory, which speeds up transactions and processing.

Stefan Panek, Oracle Database Administrator & technical Project Manager at a tech services company with 1-100 employees:

The performance of Smart Scan, Storage Indexes, etc., and the High Availability features are all valuable to us.

André Gomes Lamas Otero, Data Scientist at a tech services company with 100-1000 employees:

Group by performance Analytic functions

SeniorDA505101, Senior Database Administrator at a financial services firm with 1-100 employees:

Recovery by table Copy cluster

Barry Searle, Sr. DevOps Engineer, Adometry at a tech vendor with 1000+ employees:

Super-fast aggregated results from massive data.

Filippo Onorato, CIO with 1000+ employees:

Massive data ingestion performance
Performance SQL standard query engine

reviewer504552, Sr. Developer, Big Data at a comms service provider with 1-100 employees:

Columnar data store

CIO384, CIO at a tech services company with 1000+ employees:

We use Vertica as our primary data warehouse. It works well, relatively, most of the time.

PM436379, Senior Project Manager at a comms service provider with 1000+ employees:

We're just now getting into Vertica, but it allows us to store and access big data very quickly. It comes down to being able to quickly identify where the root cause analysis is and where trends are, so you can actually try to almost predict where problems are before they really become a problem.

Javid Ur Rahaman, Practice Director at a tech services company with 100-1000 employees:

Exadata Smart Flash Cache - caches database objects with flash memory operations Integrated appliance with all components supported by one vendor, Oracle

James Palliparambil, Sr Enterprise Database Admin at Washington Metropolitan Area Transit Authority (WMATA):

The High Availability and rolling upgrade features are the most important for us. They help us achieve zero downtime and increased performance.

Robin Saikat Chatterjee, Head of Oracle Exadata Centre of Excellence at a tech consulting company with 1000+ employees:

We had excellent Flash Cache and Smart Scan usage. We also found that every new innovation in software that came out after our model was applicable to our model. So, we can run the exact same software as the latest x5-2 Exadatas and we can even install virtualization if we want to. We have also implemented the write-back Flash Cache.

Neil Rotherham, PL/SQL Developer at a marketing services firm with 1000+ employees:

It's a fast solution, though it does take some work to ensure that queries are offloading using bloom filters, etc.

Shivendra Narain Nirala, Senior Analyst (Senior Oracle DBA) at a tech services company with 1-100 employees:

Smart Flash Cache Storage Index Smart Scan Hybrid Columnar Compression Oracle ASM Security with encryption Oracle platinum support

Zohar Lerman, BI and Reporting Platform Teams and Tech leader at a software R&D company with 1000+ employees:

MPP Analytical functions HDFS Copy Resource management

DatabaseArch877, Vertica Database Architect at a tech consulting company with 1-100 employees:

Speed Parallelization SQL language High Availability

Zvika Gutkin, Big Data DBA & DevOps at a marketing services firm with 1-100 employees:

The speed of Vertica out of the box with the ability it has to perform complex analytics queries. In other databases, information will return in hours or even days while in Vertica it will be finished in minutes or even seconds. This is the best feature it has.

Andreas Scheel, Senior Data Warehouse Architect at a media company:

Storage abstraction through projections. It gives you the possibility to react to any kind of query with an optimal performance. The Workload Analyzer helps you easily to analyze your database workloads and recommends tuning opportunities to maximize the database performance. This in turn reduces your operational costs. I love the hybrid storage model and due to that the full control of load and query behavior. I also like the ability to read semistructured data with FlexTables for DataExploration.

Justiono Handoko Putro, Sr. Lead - Infrastructure Design Database at a retailer with 1000+ employees:

Most, if not all of the time, a database workload is a disk I/O-bound, not CPU nor memory. The most time spent in a database operation is reading a data block from or writing a data block to physical disks. This is true for any type of workloads -- small or large block read and write, sequential or random access. The Exadata technology minimizes this bottleneck by utilizing Oracle's proprietary intelligent storage software which understands Oracle database I/O characteristics and thus is able to minimize number of disk I/O, and High-speed Infiniband network for storage network (40 Gb/sec throughput, 5 times higher than the typical fibre channel used for SAN).

Gafur Ameti, Senior Oracle DBA - RAC and Exadata at a tech services company:

There are several features I've found valuable, including Exadata Hybrid Columnar Compression, Exadata Smart Scan Technology, and Smart Flash Cache.

SeniorSolArch232, Senior Solution Architect at a comms service provider with 1000+ employees:

The extreme performance and data compression that comes straight out of the box.

Ruthwik Bankar, Project Manager at a tech services company with 1000+ employees:

The intelligent features have been built to utilize the storage layer in a much better way.

Merian Croos, Lead Software Engineer - Theatrical Global at a marketing services firm with 1000+ employees:

Fast query processing for historical data analytics. Write Optimized Store (WOS) continuous data loading without drastically impacting performance of OLAP queries. It's one of the few columnar databases that has the capability to provide near real time data delivery for analytics with minimal delay sourcing data from traditional databases or NoSQL data stores or any unstructured data sources.

Eva Donaldson, BI Architect / Software Engineer at a tech services company with 1-100 employees:

Speed of query response time for complicated queries on tables with billions of rows including joins on varchar columns. There is no limitation on which columns can be queried or joined on and we see query times in the milliseconds for a lot of queries that just won't return at all from other products. Ease of administration. The Management Console we thought was a nice to have turns out to give us insight on what is happening behind the scenes so easily it has sped up query tuning, insight as to what jobs are running, and resource use on the boxes the product sits on. Style of deployment. We were able to build out a server farm exactly as we are accustomed to. We did not have to buy fancy hardware. Our first cluster was deployed on servers we had sitting around from other migrations and replaced products. As we grow also the growth is native to how we do business.

Carlos Anibal, Oracle Database Specialist at a comms service provider with 1000+ employees:

SmartScan and Cell Offloading, which can drastically reduce I/O from queries, improving response time.

Kathy Duret, Oracle DBA at a retailer with 1000+ employees:

There are so many it's hard to just name a few but - Cell off-loading Reliability ExaCheck InfiniBand

Undrakh-Orgil Odonkhuu, Software Engineer at a tech services company with 100-1000 employees:

Exadata Storage Server software Smart Scan

Jason Arneil, Infrastructure Principal Director at a tech company with 1-100 employees:

One of the less well appreciated features of engineered systems in general, and Exadata in particular, is the speed with which a customer can get up and running. Having all components, both hardware and software, fully configured, tested, and optimised on delivery is a major customer win. To fully appreciate this you have to consider the alternative, which is traditionally a best of breed approach, and how long this takes customers, and the effort involved in wiring everything and ensuring full interoperability between both the hardware and software components. Having experienced the issues with this approach many times in the past, the fully converged approach of engineered systems certainly saves a lot of effort.

Trevor Watkins, Managing Partner at Thorium Data Science:

Vertica's analytic capabilities are its key strength. It can aggregate and analyze data at massive scale and neatly bring the calculation logic to the data with external procedures in C, Java and R. The Vertica architecture means it can process/ingest data in parallel to reporting and analyzing because of its in-memory Write-Optimized Storage sitting alongside the analytics optimized Read-Optimized Storage. Which brings us to projections and the DB designer which intelligently structures how data is actually stored on disk to improve the queries you actually run against it. So tables are a logical construct which are operated on as per other DBMS systems, but there's a whole next level of intelligence in optimization for querying that puts Vertica in another league.

TechTeamLead470, Technical Team Lead, Business Intelligence at a tech company with 100-1000 employees:

The most valuable feature is the merge function, which is essentially the upsert function. It's become our ELT pattern. Previously, when we used the ETL tool to manage upserts, the load time was significantly longer. The merge function load time is pretty much flat relative to the volume of records processed.

Kristopher Landon, Associate at a tech services company with 100-1000 employees:

The biggest, most valuable feature for us is the clustering aspect with a share-nothing mentality. Most clusters usually require their own shared storage, shared subnet, etc. and this becomes a pain and a nightmare to maintain. The second most valuable feature is that it's very easy to maintain. It's a breeze once you know how to handle it with your scenario in mind.

Arup Nanda, Principal Global Database Architect at a hospitality company with 1000+ employees:

The most valuable feature is the integration of the various components that make the performance soar. More specifically the presence of storage indexes along with the Exadata Storage Server software, which delivers process offloading, allows us to deliver better performance without indexes--saving space and CPU. This innovative feature is not something we would be able to replicate with other solutions. In addition, we consider the features smart scan, smart flashcache, smart flash log and hybrid columnar compression very useful.

Julian Dontcheff, Managing Director at a tech services company with 1000+ employees:

Stability, high ROI and low TCO. Plus performance, of course. Being able to compress the data, especially in big size databases, has been extremely useful as HCC can deliver quite high compression rates. The business benefits of OES are as many as the pure technical benefits/features. Exadata with the In-Memory option is levels about SAP HANA.

Evgeny Benediktov, Big Data, Analytics and Hadoop Expert, Vertica DBA (Technical Leader), Architecture Group at a tech vendor with 1000+ employees:

Ability to get top performance for in-advance known aggregative SQL queries.

Architect281, Architect at a hospitality company with 100-1000 employees:

It provides very fast query performance after good designs of projections. It's easy to implement for 24/7 data load and usage because you don’t have to worry about “load time slots” since you can load data into reporting tables at all times without worrying about their query load. It just keeps up and running all the time.

SoftwareEng525, Software Engineer at a marketing services firm with 1-100 employees:

Having projections as a parallel for indexes in a simple MySQL helped keep our data access fast and optimized.

Deep Dive: Room for Improvement

Oracle Exadata

Charles Kim, Founder and President at Viscosity North America:

For the initial instantiation process, the Oracle Exadata Deployment Assistant (OEDA) could have better error checking and pre-check validation as you navigate through the tool. When executing the OneCommand utility, generated logs are decent but the logs are not detailed enough to pinpoint to where the error occurred in the stack. Oracle can do a better job with error isolation. After the OneCommand, other one-off commands have to happen (i.e. we have to login to the infiniband and cisco switch). Ideally, OEM could be leveraged to configure the remaining components of the Exadata after the OneCommand. The idea would be to reduce the number of people required to support the stack. If we leverage OEM, we can leverage a single resource that minimally understands the stack to support the workflow.

Nitin Vengurlekar, CTO/Architect at Viscosity North America:

I think Oracle Enterprise Manager has come a long way with monitoring and managing Exadata, but getting the patching down would go a long way. I think the patching is still lacking quite a bit. If we can get the patching and upgrade in place, that would make a huge impact in overall supportability. I realize there is new OEM functionality recently introduced to support this, but it needs to be flushed out and tested.

HPE Vertica

Joe Gonzalez, BI Manager, Vertica ASE Certified DBA at a marketing services firm with 1000+ employees:

Most of the changes I need are in 7.2. We have not yet been able to upgrade to it, but are planning on doing so by the end of the year. Some of the things that are supposed to have been improved/added to Vertica in 7.2 are - Better support for Live Aggregate Projections, including support for the DISTINCT function The ability to restore a single object from a full backup instead of having to do an object level backup first The ability to more easily restore a backup to an alternate cluster The addition of numerous hints to customize your queries The ability to enforce check constraints on table load Major upgrades to the Query Optimizer

SoftwareDataArch072, Software and Data Architect at a software R&D company with 1000+ employees:

Loading times for “real time” sources - for example, loading from Spark creates a load on the DB at high scale Connectors to other sources such as Kafka or AWS Kinesis Better monitoring tools Better integration with cloud providers - we were missing some documentation regarding running Vertica on AWS

Chris Bradham, Technical Director at a consultancy with 1-100 employees:

The patching process is a continually evolving. It's changed drastically over the years and foresee some continued refinement on the horizon, especially when comparing the process to something like ODA. We have one one-off patch which requires a SR with Oracle Support every time the quarterly patch comes out. This is a bit of a pain but I do know sometimes these one-offs get rolled into the quarterly patch or in our case the next version (11.2.0.4).

Manish Nashikkar, ITA - Oracle Apps DBA at Tata Consultancy Services:

Exadata Linux systems have Intel CPUs inside. I would suggest that if Oracle could work together with Intel to have some more intelligence at the CPU level, then there would be nothing like it.

Saurabh K. Gupta, Database Development Leader at a tech vendor with 1000+ employees:

Starting from Exadata X5, virtualization is supported. The feature, though, for some reason, is not much talked about. If Oracle starts promoting OVM and VMWare on Exadata through white papers and case studies, I'm sure it will enable new IT setups in practice. For instance, virtualization support can enable ETL nodes to be resident with Exadata.

Mohammad Rafey, Senior Principal Consultant at Viscosity:

The key features with the Exadata is offloading the query processing, some of them, at the storage level. That's where Oracle has to make it smarter. I think it's already smarter, the Exadata storage. I can't think of a specific feature that can put in, but that's where I think there's room for more improvement. Adaptability is a little bit challenging for the customer because of the licensing and the pricing. That's where I think they can make a big difference.

Matt Williams, Lead Data Scientist Machine Learning at a financial services firm with 1-100 employees:

I'm concerned that HP Enterprise has sold their software business, and worry about future investment to enhance predictive/machine-learning capabilities.

Rodrick Brown, Development Operations/SRE at a software R&D company with 100-1000 employees:

More frequent updates.

Davide Moraschi, Senior business Intelligence consultant at Asociación SevillaUP:

Sadly, it does not support stored procedures in the way we are used to thinking of them. There is the possibility to code plug-in in C++, but that's out of our reach. Correlated sub-queries are another point where we'd love to see enhancements, plus the overall choice of functions available. ETL with SSIS was not as easy as one we had expected (must remember to COMMIT and we had some issues with datetime + timezone, but that's was probably our fault). OleDB and .NET providers need some touches; and another great improvement would be support for Entity Framework, which so far I haven't seen. There is no serious graphical IDE for HPE Vertica, that's frustrating. One free option available is DbVisualizer for Vertica, but it's a bit basic.

Suresh Muddaveerappa, Senior DBA and Architect at a tech services company with 100-1000 employees:

In the area of Solaris zone-level virtualization, it would be good to have memory capping as a tool for memory management. Currently for 11g databases running on Exadata with smart scan against hybrid columnar compressed (HCC), tables could result in errors. It would be nice to have a patch rather than the current solution of upgrading the databases to 12c.

Rob Visser, CIO Group Services at a insurance company with 1000+ employees:

We are using it now predominantly as a hardware stack managed by our company. We would love to see it more as a service; a database-as-a-service offering.

Vijay Lunawat, Director - OSP/Engineered Systems at a tech services company with 1000+ employees:

A major concern from customers is that it's very costly, but if they think about the total cost of ownership over 3-5 years, I think the total cost attests to the business readiness and it doesn't come out to be very costly.

Architect6dc7, Unix/Linux Platform Architect at a financial services firm with 1000+ employees:

It can be difficult to patch and maintain because there are so many databases running on it.

Chris Gilliam, President at DataMAPt, Inc.:

Sometimes pricing can be a bit of an issue, especially if customers don't know exactly what they're getting. Maybe they've gone through the whole process and they didn't get the right size or the right amount. I've worked with several customers who had implementations that were too small and then had to upgrade, as well with others who were overpowered with their implementations.

DBA511, Database Administrator (DBA) at a software R&D company with 100-1000 employees:

In the versions I worked with, if a majority of the nodes were being loaded under heavy, sustained rates the nodes would see some dramatic decreases in performance due to the data shuffling that needed to occur between all the nodes. To work around that we ended up doing most of the loading in one or two nodes and that helped significantly. The synchronizations problems occurred when loading about 10 billion events, at a rate of about 100k tuples/second/node across 5 nodes. One of the suggestions from Vertica engineering was to increase the number of nodes to offset how much data was being sync'd per node.

Yoojong Bang, Data Scientist at a media company with 100-1000 employees:

Speed: It's already doing what it is supposed to do in terms of speed but still, as a user, I hope it gets even faster. Specific to our company, we do store the data both in AWS S3 and Vertica. For some batch jobs, we decided to create a Spark job rather than Vertica operations for speed and/or scalability concerns. Maybe this is just due to the computation efficiency between SQL operations vs. a programmatic approach. Even with some optimization (adding projections for merge joins and grouped by pipelined), it's still taking a longer time than a Spark job in some cases.

Saswata Sengupta, Solution Engineering and Arcitect - Big Data, Data Science and Cloud Computing at a tech company with 1000+ employees:

Projections take up a lot of space and hence, compression can be improved. Installation can be more intuitive via a simple, lightweight web client instead of the command line.

Michael Abbey, Manager, Principal Consultant at Pythian:

As is always the case with Oracle, when some late-breaking fancy technology gets under their radar, self-invented or perhaps something that they're picking up that they'd like to compete with from another vendor, they're all over it. I can't specifically think of anything myself. Outside I guess of speed and maybe the other two things I could think of are speed and speed, but I'm not suggesting that speed is an enhancement because there's anything wrong with the speed of the system now, but of course we always like to do things in four nanoseconds rather than seven.

Director662, Director, IT at a tech services company with 1000+ employees:

They need to reduce the cost. It needs to be a true appliance so you don’t have to manage individual components inside it. It needs to provide a better maintenance process, which means no downtime, and true rolling upgrade.

SrDBA3602, Lead Oracle Database Administrator at a insurance company with 1000+ employees:

There's a lot with Exadata that can be done on a black-box level which depends entirely on what the vendor is able to do for you, which is good from a certain perspective. But we hope, as this technology matures, Oracle allows customers to have a little more hands-on customization of some certain elements to better suit their environments. So we're looking forward to the ability to do greater customizations to suit our business needs.

SeniorDBA728, Senior DBA at a local government with 1000+ employees:

Performance tuning Not much by way of any documentation: The explain plans are very difficult to read / understand. I tried to diagnose some specific queries using the DBD Vertica utility, etc. For one example of using the explain plan, the query was complex with lots of joins and so on (the query took up about three A4 pages), but the explain plan I printed out took up in excess of 32 A4 pages. How on earth would you read that? No visual tools were available that I could find. Very little if any training available in the UK: Our company wasn't able to find any on the topic. We found very little if any documentation (from the vendor) that was of much use. Cloning / export was not well documented; poor examples.

Victor Daniel Gonzalez, Vertica DBA at a tech services company with 1-100 employees:

Documentation, there are functions that are not documented. UDF SDK, I'd like to see a step by step simulator example in a manual. The read-me code is good, however, an example would be great for starters.

Sivaprasad Kunchakuri, Database Admin at a tech services company with 1000+ employees:

vbr.py needs to be improve to support diff no of nodes source to target.

seniorda505386, Senior Database Administrator; HP Vertica, MySQL, MSSQL, Cloud Ops at a tech vendor with 100-1000 employees:

Performance tuning; user community is needed.

Stefan Panek, Oracle Database Administrator & technical Project Manager at a tech services company with 1-100 employees:

The Lifecycle Management over the whole stack could be improved over what it already provides. The rolling upgrade feature on the database itself should be improved. It needs some Java updates as well as there is too much downtime related to Java issues.

Javid Ur Rahaman, Practice Director at a tech services company with 100-1000 employees:

It needs built-in big data features. I'd like to see smart auto-healing features with machine-learning libraries. Less power consumption would be nice. It needs integrated cloud software to enable cloud connectivity.

Robin Saikat Chatterjee, Head of Oracle Exadata Centre of Excellence at a tech consulting company with 1000+ employees:

Oracle has actually made leaps and bounds in terms of improvement since our model. Our main gripes were with the storage capacity (ours are 600GB disks, now you can get 8TB disks) and flexibility of configuration (now you have the ability to opt for elastic configurations, capacity on demand and virtualization), so our gripes have all been addressed. We would like an option of a hardware-only support solution, but Oracle currently does not provide this.

Neil Rotherham, PL/SQL Developer at a marketing services firm with 1000+ employees:

As soon as your joins get complicated, things start to go wrong. You lose much of the offloading and start spilling into temp space. We were reporting on billions of rows and it was a constant problem so a custom reporting framework had to be created. Parallel queries need to be carefully managed and the DBAs need to be right on top of the resource manager, as if you aren't managed properly as a user, you can end up waiting behind others.

Asim Singh Mahat, Sr. Software Engineer (Database) at a tech services company with 1000+ employees:

More analytical functions. Optimization around DML operations such that we be able to use it more. I understand that Vertica is not meant for that, but I would really love it if it has capability to support procedural processing. Optimization around DML refers to fast DELETE and UPDATE statement such that we can leverage Vertica around those operations. I do understand that Vertica is not meant as an OLTP system, neither I'm asking to have it similar but if DML operations can be optimized, that would be admirable.

Regarding, my comment on the capability to support procedural processing - Vertica as of now is mainly used via SQL only. If we have to use any procedure based operation, we do it via User Defined Functions. If within Vertica itself, we have the capability to create & execute procedure similar to that for functions, it would be a plus. Again, I do understand that this may be against the architecture of Vertica but if there is anything that can be revised to get these supported, that would be preferred.

Andrey Krasnov, High FrequencyTrading Systems and Strategy Architect/Quant Trader at a financial services firm with 1-100 employees:

Some GUI Tools out of the box, better python integration. I would love to see some nice query engine, tooled specifically to Vertica extensions of SQL (with IntelliSense). We currently use TOAD, but it has a lot of bugs and does not provide full support of all Vertica features. For example with the "copy to" command it is buggy - extremely hard to debug. Other issues would be better query plan display and better management tools like command line admin tools. (MC would probably solve a bit here, I saw a demo on the conference). We have several avid Python users in the company, but this how maybe lower priority after I reviewed my conference materials and found that Vertica now has native Python driver.

Shivendra Narain Nirala, Senior Analyst (Senior Oracle DBA) at a tech services company with 1-100 employees:

Shared RAM for multiple instances Hardware update in terms of storage Ability to migrate a database to Exadata from a normal environment

Justiono Handoko Putro, Sr. Lead - Infrastructure Design Database at a retailer with 1000+ employees:

With its value proposition, Exadata is being used to run mission-critical, 24x7 applications. Unfortunately, not all hardware in the Exadata rack are hot-swappable (such as memory, processor and battery maintenance in older models), and therefore business application downtime may be required for those hardware replacements. Infiniband cabling work may need a complete downtime as well, for example, when we need to connect multiple racks.

Gafur Ameti, Senior Oracle DBA - RAC and Exadata at a tech services company:

There were some minor problems moving databases to Exadata, but it was more a problem with the internal processes with my current employer.

SeniorSolArch232, Senior Solution Architect at a comms service provider with 1000+ employees:

X5 supports virtualization, but I think the performance on virtualization is not all that great. I have heard it from a customer, though I admit I haven't tested it out myself.

Matti Heimonen, DW Admin at a hospitality company with 1000+ employees:

Data velocity and manageability.

Kjetil Halvorsen, Management Consultant at a software R&D company with 1-100 employees:

More ML, both data prep, models, evaluation and workflow. Improved support for deep analytics/ predictive modelling with machine learning algorithms. This area of analytics need a stack of functionality in order to support the scenario. The needed functionality includes: Data preparation. Scaling, centering, removing skewness, gap filling, pivoting, feature selection and feature generation Algorithms/models. Non-linear models in general. More specifically, penalized models, tree/rule-based models (incl.

ensambles), SVM, MARS, Neural networks, K-nearest neighbours, Naïve bayes, etc. Support the concept of a “data processing pipeline” with data prep. + model. One would typically use “a pipeline” as the overall logical unit used to produce predictions/scoring. Automatic model evaluation/tuning. With algorithms requiring tuning, support for automated testing of different settings/tuning parameters is very useful. Should include (k fold) cross validation and bootstrap for model evaluation Some sort of hooks to use external models in a pipeline i.e.

Ruthwik Bankar, Project Manager at a tech services company with 1000+ employees:

More control for database admins is needed for storage indexes and pricing.

Carlos Anibal, Oracle Database Specialist at a comms service provider with 1000+ employees:

Massive updates statements should not disable HCC.

Kathy Duret, Oracle DBA at a retailer with 1000+ employees:

Most of my suggestions for Exadata itself have been addressed in X5. Oracle keeps improving the reliability and adding more hot swappable parts for the hardware and software. The only improvements left are not with the product itself but with MOS support, ASR and patching. For MOS software support, there are some very talented people there, but not enough of them. They need more software engineers with in depth training on how to use the ASR system and the jump gate. When the jump gate is updated or patched, it seems to take a bit to get things working properly again so patching and ASRs can be done. Platinum patching continues to improve but it is not a hands off solution.There is quite a bit of manual effort involved with each patch cycle to co-ordinate the patching. With some of the patching there, is downtime, and most of it is rolling but there are a few cases where you do need have downtime.

Undrakh-Orgil Odonkhuu, Software Engineer at a tech services company with 100-1000 employees:

The ASR needs to be improved.

Kjetil Halvorsen, Management Consultant at a software R&D company with 1-100 employees:

data prep in Vertica + model from Spark/R. Parity functionality for the Java SDK compared to C++. Today the C++ SDK is the most feature rich. The request is to bring (and keep) the Java SDK up to feature parity with C++. Streaming data and notifications/alerts. Streaming data is starting to get well supported with the Kafka integration. Now we just need a hook to issue notifications on streaming data. That is, running some sort of evaluation on incoming records (as they arrive to the Vertica tables) and possibly raising a notification.

reviewer505263, Principal Data Architect (BI) at a media company with 1000+ employees:

Faster deletes!

Srikanth Iyer, Senior Data Architect at a media company with 1000+ employees:

Ability to use analytic functions in where clauses, being able to use aliases in the where and order by clauses will make query writing/reading a lot easier.

Jfm, Owner at a tech vendor:

Profiling, query optimize, management.

Karla Marisol Guajardo, Infrastructure, Data Center and PMO Coordinator at a comms service provider with 100-1000 employees:

The time that the mediation process takes and historical information that I can store.

Jason Arneil, Infrastructure Principal Director at a tech company with 1-100 employees:

Several barriers to entry have been overcome with the latest X5 generation, in particular moving to more flexible sizing (elastic configurations) allows customers to choose the exact fit of compute and storage resources they require.

Arup Nanda, Principal Global Database Architect at a hospitality company with 1000+ employees:

There are three possible enhancements: Build a stronger, more responsive support team. Add a RAID-5 like storage layout for customers to save space with the full understanding that performance will be less (which could be fine for non-prod systems). Add storage level replication without the use of Data Guard, which is quite useful in case of organizations that rely on that technology for their Disaster Recovery effort.

Julian Dontcheff, Managing Director at a tech services company with 1000+ employees:

The Initial investment price could use improvement. It is often the stumbling block. Many organizations have limited or small budgets for the initial investment. This has somehow changed with X5 when the entire Exadata configuration no longer needs to be licensed when using OracleVM.

DataSci505212, Data Scientist at a tech vendor with 100-1000 employees:

More Machine Learning algorithms--Random Forest for sure!

André Gomes Lamas Otero, Data Scientist at a tech services company with 100-1000 employees:

Debug custom functions in r.

SeniorDA505101, Senior Database Administrator at a financial services firm with 1-100 employees:

Improve Vertica logging and messages to Vertica startup commands.

Barry Searle, Sr. DevOps Engineer, Adometry at a tech vendor with 1000+ employees:

Better feedback from installation. I would like to see more meaningful errors returned and more graceful handling of those. Thankfully, we don't often hit error conditions.

Filippo Onorato, CIO with 1000+ employees:

Machine learning implementations Support for Cloud based environments like Google Compute Engine

reviewer504552, Sr. Developer, Big Data at a comms service provider with 1-100 employees:

Add geospatial indexes (sounds like they have done it in version 8.0)

CIO384, CIO at a tech services company with 1000+ employees:

I just expect it to work and be serviceable. When we ran into issues, there seemed to be a lot of different opinions for how to resolve the issues and that was the feedback I gave to them. You talked to one tech, you talk to a different tech they had a much different approach. That was a big frustration point for us. The upgrade path and which way we should go. So at the end it created a lot of confusion for us, so I wouldn't upgrade it again lightly. We're going to remain on it for the next year, but we'll probably re-evaluate at that point if we want to continue with Vertica or something else.

PM436379, Senior Project Manager at a comms service provider with 1000+ employees:

I guess really the only thing there is if you get a server big enough to handle Vertica, it does just fine. If you're working in a small business, it will tend to overtake most of their budget from a cost perspective because you need so many servers, so much storage, to be able to handle all that stuff.

Zohar Lerman, BI and Reporting Platform Teams and Tech leader at a software R&D company with 1000+ employees:

Active-Active clusters with online replication. Greater query concurrency. Better documentation/white papers as there arte lots of undocumented issues.

DatabaseArch877, Vertica Database Architect at a tech consulting company with 1-100 employees:

While the documentation is very extensive and relatively complete, it’s poorly organized and there are way too few examples. It’s come a long way since the first version I saw, but it still has a long way to go. Plus, there is very little information on the internet. I can find a solution to nearly any MS SQL Server problem using Google. Not so for Vertica.

Zvika Gutkin, Big Data DBA & DevOps at a marketing services firm with 1-100 employees:

The internal documentation. As a DBA, I really want to understand how its internals works. It needs to handle high concurrency short queries better as Vertica is not handling these well, and we have had to develop our own tool to help us with our dashboard.

Andreas Scheel, Senior Data Warehouse Architect at a media company:

I would love to see direct connections to other DMSs. Something like a direct connector to Oracle, MySQL, MS SQL, MongoDB, etc. so that you can copy data between Vertica and other vendors directly and more easily without an ETL tool, dump, transport, or load data.

Merian Croos, Lead Software Engineer - Theatrical Global at a marketing services firm with 1000+ employees:

Stability, scalability (3 node Community Edition) and backup/restore all need to be worked on. Without proper work load management and resource pool allocation, any batch/ETL or streaming jobs which refreshes data frequently will impair OLAP query performance.

Eva Donaldson, BI Architect / Software Engineer at a tech services company with 1-100 employees:

Getting the Management Console up and running as expected was a bit of a challenge.

Trevor Watkins, Managing Partner at Thorium Data Science:

We've built a data ingestion tool to sit alongside Vertica for easy data loading, and I would personally like to see extended developer tooling suited to Vertica – think published PowerDesigner SQL dialect support, IDE with IntelliSense, and stored procedures which we've also had to build a work-around module for.

TechTeamLead470, Technical Team Lead, Business Intelligence at a tech company with 100-1000 employees:

We've had issues with query time taking longer than expected for our volume of data. However, this is due to not understanding the characteristics of the database and how to better tune its performance.

Kristopher Landon, Associate at a tech services company with 100-1000 employees:

Developer Tools - Vertica really needs some kind of IDE plugin for a system such as Eclipse or IntelliJ. Developing external functions in Vertica can kind of be like shooting in the dark sometimes. Also, an improved monitor or monitoring with alerting built-in that actually works would be a welcome addition. They truly need a Python or some script that can handle all of the low-level system changes for you and find out how the customer has heavily modified their nodes before the install. Some automation here would help a lot. The product overall is a great product, however management tools as well as monitoring tools are lacking. The product does, however, offer a lot of information in the form of system views and tables, but most of the data is hard to translate with out the help of their support team.

Evgeny Benediktov, Big Data, Analytics and Hadoop Expert, Vertica DBA (Technical Leader), Architecture Group at a tech vendor with 1000+ employees:

I really would like to see Vertica able to use heterogeneous storage (RAM, SSD, HDD). Another issue I have seen is that the SQL optimizer fails to make optimizations that competing products are able to do. That's something that should be improved as well.

Architect281, Architect at a hospitality company with 100-1000 employees:

We are running our data transformations as an ELT process inside Vertica; we have data at least on the landing area, temporary staging area, and final data model. Data transformations require lots of deletes and updates (which are actually delete/insets in Vertica). Delete in Vertica doesn't actually delete data from tables, it just marks them as deleted. For us to keep the performance up, purge procedures are needed and a good delete strategy needs to be designed and implemented. This can be time consuming and is a hard task to complete, so more 'out-of-the-box' delete strategies would be a nice improvement.

SoftwareEng525, Software Engineer at a marketing services firm with 1-100 employees:

We would like to see better documentation and examples as well as further simplicity in creating clusters, adding nodes, etc. I understand the GUI is very simple but sometimes more insight into what the product is doing and where errors are occurring would help debugging.

Deep Dive: Improvements to My Organization

Oracle Exadata

Charles Kim, Founder and President at Viscosity North America:

We are able to provide a complete solution to our customers from data center installation/configuration, firmware upgrades, ILOM upgrades, OS patching/upgrades, cluster patching/upgrades and database patching/upgrades. We can upgrade the entire stack in a single evening with minimal outage. Depending on the customer's tolerance for downtime, we can either perform the patching or upgrade in a rolling fashion. Patching and upgrade services have proven to be a value-add differentiator for SMB and mid-market customers where resources and budgets are often limited. A single consulting resource can patch/upgrade the entire Exadata stack because patching and maintenance has become reliable and simplified.

Nitin Vengurlekar, CTO/Architect at Viscosity North America:

If you have server footprint reduction or consolidation as a requirement, then you can consolidate several databases and configurations onto one platform. Combine that with performance, it becomes a great avenue for cost reduction in terms of OpEx and CapEx. Exadata cannot be one-dimensional solution, performance, consolidation, and scalability must be part of the overall solution.

HPE Vertica

Joe Gonzalez, BI Manager, Vertica ASE Certified DBA at a marketing services firm with 1000+ employees:

The fast columnar store database structure allows our query times to be at least 10x faster than on any other database. This enables us to get answers to data questions as well as numerous analytics on our data out to our internal and external clients quickly. We are able to integrate our Vertica data warehouse with Tableau to create numerous reports quickly and efficiently. What was once a two year backlog of report requests on our old data system has been virtually eliminated now that we are using Vertica to provide the solutions.

SoftwareDataArch072, Software and Data Architect at a software R&D company with 1000+ employees:

We built an internal dashboard using the MicroStrategy to increase visibility to our management and our employees. Also, we built tool to expose the data to our selected partners and users to create better engagement with our platform.

Chris Bradham, Technical Director at a consultancy with 1-100 employees:

We have a number of statistics collected before cutover on our legacy environment compared to Exadata. Without doing anything other than copying the data across, we saw significant performance gains for most key processes. We receive feedback from users stating how fast the performance is compared to other systems. Performance issues are few and far between. Our database environment is extremely stable compared to the legacy DB configuration. We upgraded from a X2-2 quarter rack to a X5-2 eighth rack and experienced significant performance gains.

Manish Nashikkar, ITA - Oracle Apps DBA at Tata Consultancy Services:

Because of Exadata Systems, we were able to consolidate all the applications databases into one, e.g, Oracle EBS, Siebel, Hyperion etc.

Saurabh K. Gupta, Database Development Leader at a tech vendor with 1000+ employees:

The Exadata database machine is a perfect solution to cater to the needs of large data warehouses. It aptly addresses the problems of database consolidation and marks an end to the argument of running mixed workloads together. With Exadata, we were not only able to consolidate our large data warehouses in a single machine, but we also stabilized the peak workloads. The throughput issues simply vanished with the features like Flash Cache and Infiniband Network Fabric. On the operational fronts, Exadata streamlines the tasks and responsibilities in a distinctive manner across the levels of an enterprise IT structure. Solution architects act as machine administrators and make sure that the machine is being used in optimal way. At the same time, database administrators, being the familiar heads, don't find much change in roles. Overall, Exadata justifies the total cost of ownership (TCO) by balancing the capital and operational expenses - thereby achieving a higher ROI.

Matt Williams, Lead Data Scientist Machine Learning at a financial services firm with 1-100 employees:

Performance management of high-traffic sites - Vertica's ease of scaling has been invaluable for one of our main customers.

Rodrick Brown, Development Operations/SRE at a software R&D company with 100-1000 employees:

We're able to retrieve queries nearly instantaneous for our custom analytical tools we built on top of Vertica.

Davide Moraschi, Senior business Intelligence consultant at Asociación SevillaUP:

Data Warehouse response times have decreased of one order of magnitude with respect to the previous solution (SQL Server + Oracle).

DBA511, Database Administrator (DBA) at a software R&D company with 100-1000 employees:

Our data warehouse at the time was a multi-terabyte PostgreSQL cluster. It worked really well, but we wanted to increase the size to many TB's and our due to our query and loading patterns we found greater performance from Vertica's multi-node warehouse.

Yoojong Bang, Data Scientist at a media company with 100-1000 employees:

Large-volume queries are executed in a relatively short amount of time, so that we could develop reports that consume data in Vertica.

Mohammad Rafey, Senior Principal Consultant at Viscosity:

Some of the key benefits with our customers moving to the Exadata, they can get it consolidated, all the databases into one platform, one environment. And other advantages, you get a lot of the new features out of the box of the Exadata itself. We don't have to shelve around a lot of different parts of the machines, the networking component, the storage components, you know, all together, one into one machine, into the Exadata. It's a huge, huge advantage for the customers, right? They don't have to look into the history, they don't have to look into the database machines, you know, networking part of it all income passes in one machine. And Oracle has engineered it specifically for the purpose of data. Optimized the database solutions. So once you move it, you get the benefits of your query optimization, your history optimization, your machine data with computer optimization. Everything is blended into one machine.

Suresh Muddaveerappa, Senior DBA and Architect at a tech services company with 100-1000 employees:

With SuperCluster and Exadata, all of the servers and storage are integrated within the same rack. This reduces the configuration and setup time, increases performance, and makes maintenance and patching easy. We are able to consolidate all of the databases and application stack on a single SuperCluster with Exadata.

Rob Visser, CIO Group Services at a insurance company with 1000+ employees:

We have seen performance and reliability benefits. It's also beneficial from a cost perspective.

Saswata Sengupta, Solution Engineering and Arcitect - Big Data, Data Science and Cloud Computing at a tech company with 1000+ employees:

This product is used for in-database analytics for reports and queries that require very fast response times. Complicated multi-table queries perform very well, and the company has improved on business operations looking at hot data from various dimensions.

SeniorDBA728, Senior DBA at a local government with 1000+ employees:

We use it for analytics (marketing).

Victor Daniel Gonzalez, Vertica DBA at a tech services company with 1-100 employees:

It has helped us escalate, we need information almost real-time.

Sivaprasad Kunchakuri, Database Admin at a tech services company with 1000+ employees:

Replication and Node recovery in 8.0.

Andrey Krasnov, High FrequencyTrading Systems and Strategy Architect/Quant Trader at a financial services firm with 1-100 employees:

Certain research which was unattainable beforehand, now is in reach.

Matti Heimonen, DW Admin at a hospitality company with 1000+ employees:

It is the foundation of our new Data Warehouse platform.

Vijay Lunawat, Director - OSP/Engineered Systems at a tech services company with 1000+ employees:

The business outcomes are much faster; end users are very happy so they can focus on their day-to-day work rather than waiting for the computer, the response from the system. In many customer cases, reports that used to take 10-12 hours get completed in 1-2 hours, a massive reduction. A business cycle that used to take about 3-5 days to turn around can be turned around in 1 day, so that's where we optimize our business utilization; much more efficiency.

Simone Traversari, Team Leader at a tech consulting company with 100-1000 employees:

Oracle Exadata and Oracle Exalogic are standard-based. Customers can easily migrate their existing applications based on Oracle Database and Oracle WebLogic: they can take advantages of Engineered Systems with minimal effort during the migration process. Customers can benefit from high performance without modifying Java or SQL code and they can decide to use specific features after migrating to Exadata/Exalogic. Also, Engineered Systems can be used to consolidate complex environments: for example, our customer was able to decommission about 20 blades, using a single 1/8 Exalogic x4-2. Finally, Exadata and Exalogic simplifies patching/upgrading process, unifying firmware upgrades, ILOM upgrades, OS patching, Database/WebLogic upgrade/patching in one "single-vendor" activity.

Kjetil Halvorsen, Management Consultant at a software R&D company with 1-100 employees:

We are an HP partner. A SQL-based compute platform like Vertica enables far less human overhead in operations and analytics.

reviewer505263, Principal Data Architect (BI) at a media company with 1000+ employees:

Great reporting.

Srikanth Iyer, Senior Data Architect at a media company with 1000+ employees:

We are trying to data mine customer event data. Having the ability invoke analytic functions without having write self join SQL statements ... just brilliant.

DataSci505212, Data Scientist at a tech vendor with 100-1000 employees:

We were able to implement new algorithms without having to move data out of Vertica into a compute cluster. This allowed us to offer Analytics for Cybersecurity to our customers.

André Gomes Lamas Otero, Data Scientist at a tech services company with 100-1000 employees:

We could run group by queries thousand of times faster, we are able to test more models and improve accuracy.

SeniorDA505101, Senior Database Administrator at a financial services firm with 1-100 employees:

VBR backup used to take more than one week to back up 70 TB of data. After upgrading to latest version, it is taking about 48 hours.

Chris Gilliam, President at DataMAPt, Inc.:

Definitely having the ability to have one machine do everything has been the biggest improvement for us organizationally. From the administration side, it helps. If one or two admins can handle it, they don't have to go all over the place to, such as to the storage team, the network team, or anywhere else because they can do a lot of things internally. Definitely, we've seen cost, time, and HR savings with several implementations.

Michael Abbey, Manager, Principal Consultant at Pythian:

The thing that attracts them to the two products but mainly the Exadata is the storage technology is revolutionary whether it's the only solution of that flavor on the market or not, I'm not familiar with. It's the storage solution, it's the speed, and it's the high availability. If somebody wants 5 or 6 9s worth of availability, a very good place for them to be is with one of Oracles engineered system solutions.

Director662, Director, IT at a tech services company with 1000+ employees:

It makes it possible to deliver high performance Oracle OBIEE applications to our customers.

Barry Searle, Sr. DevOps Engineer, Adometry at a tech vendor with 1000+ employees:

We can process vast amounts of data, fast and with a high degree of reliability.

Filippo Onorato, CIO with 1000+ employees:

DWH core platform is based on it

PM436379, Senior Project Manager at a comms service provider with 1000+ employees:

The ability to access in-store, big data, and be able to create keywords for faster resolution and look up an individual, hey we did this problem before. It'll show you all the steps and everything, along with different products. Vertica is pretty much the database behind it. It really does the performance aspect of it.

Zohar Lerman, BI and Reporting Platform Teams and Tech leader at a software R&D company with 1000+ employees:

It's enabled us to develop our new reporting system which is used as a SaaS by hundreds of users. We can also load massive amounts of data in seconds and query it with SLA for online dashboards.

DatabaseArch877, Vertica Database Architect at a tech consulting company with 1-100 employees:

I have seen queries that take over 24 hours on MS SQL Server to complete, complete in less than 10 minutes on Vertica. I have seen queries that take several minutes, up to an hour, on MS SQL Server, complete in less than 10 seconds, sometime less than one second on Vertica. That allows analysts to spend their time analyzing results instead of waiting for results. Certain types of analysis weren't even possible before, simply because it took too long.

SrDBA3602, Lead Oracle Database Administrator at a insurance company with 1000+ employees:

As with any organization, before we embark on anything, we have to look at the ROI as well as the ability of a particular tool. Exadata helps us to consolidate and manage our databases as one with much faster processing. This, of course, is the key in speeding up our ability to do work and provide customer service. We're a mutual fund company, so speed is really key in providing satisfactory customer service. Exadata really gives us the key factors from an IT perspective, which are performance and availability.

Stefan Panek, Oracle Database Administrator & technical Project Manager at a tech services company with 1-100 employees:

We have a complete HA environment with two Exadata machines across two datacenters. We've able to use Data Guard to perform switchovers, perform maintenance, and do other tasks with little or no downtime at all.

Javid Ur Rahaman, Practice Director at a tech services company with 100-1000 employees:

Our business cycle has benefited greatly with Exadata as our as operational efficiency has improved by 200%. Also, our business analytics platform performance has improved by 300%.

James Palliparambil, Sr Enterprise Database Admin at Washington Metropolitan Area Transit Authority (WMATA):

The biggest advantage to our implementation of Exadata is that because it's an Oracle-engineered system, our database administrations don't have to worry about performance tuning. It's already optimally tuned and, when compared to other systems, it's really provided us with better performance.

Zvika Gutkin, Big Data DBA & DevOps at a marketing services firm with 1-100 employees:

Vertica is in our core technology stack. We are serving reports and dashboard to clients from it. It's very important to us that it fulfills its function correctly and provides us with an advantage over our competitors.

Andreas Scheel, Senior Data Warehouse Architect at a media company:

We are now able to procde real-time insights into our tracking data, and with that show how our customers are using the products that we have. Furthermore, it is now possible for our Data Science department to easily, and quickly train their new data mining models and get answers faster than ever before. With the hybrid storage model along with well designed resource pools and storage abstraction through projections, we are now able to easily load new data constantly throughout the whole day. While doing this, we can still be available to perform data analytics on new and legacy data quickly, and even Microstrategy for enterprise reporting doesn’t need to cache data. Most reports can be generated with live queries and still finish within seconds. So in a nutshell: - Faster Information Insight (Data to Insight cycle) - Less complexity on data modeling- Less operational costs

Merian Croos, Lead Software Engineer - Theatrical Global at a marketing services firm with 1000+ employees:

With traditional RDBMS historical data analysis or any complex queries took minutes to complete. With the addition of Vertica to handle big data queries, these reports are now returned in under 15 seconds. The biggest performance improvements obviously are for queries that have to analyze a large amount of historical data.

Robin Saikat Chatterjee, Head of Oracle Exadata Centre of Excellence at a tech consulting company with 1000+ employees:

Our machine is used for our internal development of Oracle-based solutions, PoC benchmarking, and training/familiarization. We have gained a lot of information on the product and it has helped us master this technology. As we are a consulting firm, this is an invaluable resource for our clients as they get the benefits of our research.

Neil Rotherham, PL/SQL Developer at a marketing services firm with 1000+ employees:

It replaced a legacy database (SAND) and was part of a major overhaul. It didn't change the way the company worked per se, but it did force the creation of an Exadata compatible reporting framework, and it was that framework which enabled a leaner rewrite of the existing reports and provided a basis for new ones.

Shivendra Narain Nirala, Senior Analyst (Senior Oracle DBA) at a tech services company with 1-100 employees:

We were using a normal RAC on a Linux server for online applications and Oracle on the same platform for data warehousing. After moving to Oracle Exadata, we were able to add more users, and queries get executed very quickly for data warehousing, around one quarter to half the time when compared to our previous solution under normal conditions.

Eva Donaldson, BI Architect / Software Engineer at a tech services company with 1-100 employees:

We can have insight into data we never had before. We can provide that insight to internal users so we do not have to generate reports for them all the time. With response times like these there is no concern of having them wait for results to return and so they do not think things are broken.

Trevor Watkins, Managing Partner at Thorium Data Science:

Our consultancy has introduced Vertica to a number of clients, from small scale ones who benefit from the free tier and per TB pricing model to have a powerful analytics cluster fairly cheaply to large investment banks who have been able to handle data at a scale that wouldn’t otherwise be possible.

TechTeamLead470, Technical Team Lead, Business Intelligence at a tech company with 100-1000 employees:

HP Vertica has helped us democratize data, making it available to users across the organization.

Kristopher Landon, Associate at a tech services company with 100-1000 employees:

Loading raw data and leveraging column store to quickly aggregate the values as well as run a general analysis were the biggest improvements we found. Before, we had to scrub the data or reformat, load it, possibly scrub it some more, and then run the first set of analysis, and so on. With Vertica, we were able to combine some of these steps, such as loading gzip data directly into the table and leveraging R in Vertica to run all of the analysis.

Justiono Handoko Putro, Sr. Lead - Infrastructure Design Database at a retailer with 1000+ employees:

There are several scenarios in which this technology can help. Most package (COTS) software are not designed, and therefore are not optimized, for a specific database platform, e.g. Oracle. Performance starts becoming a problem when the database size and workload (e.g. transaction rate) are high. Even if you, as a customer, have reached out to vendor support, chances are they are not going to customize their code just for one database platform, Oracle. They may help you to do small tuning here and there, but without optimizing the application code, the result is limited.

Exadata comes as a quick fix in this situation. Oracle Exadata will sweep the problem under the carpet. Secondly, most applications have a mixed-workload, not purely OLTP and not purely batch/reporting. Unfortunately in many aspects, database tuning for the fastest response (as in OLTP workload) and for the highest throughput (as in reporting/DSS workload) are a contradiction. Again, Oracle Exadata can be a quick fix by minimizing the number of I/O on a fast and high-throughput storage network. It helps on any kind of workloads.

Gafur Ameti, Senior Oracle DBA - RAC and Exadata at a tech services company:

We have dramatically improved performance for the batch jobs, after we moved from the traditional Oracle 11G databases to the new Exadata servers. Performance has improved several-fold and moving (consolidating many small databases from smaller Linux/Unix servers into Exadata cluster) has resulted in lower production support cost and effort for BAU DBAs.

Evgeny Benediktov, Big Data, Analytics and Hadoop Expert, Vertica DBA (Technical Leader), Architecture Group at a tech vendor with 1000+ employees:

HP Vertica is an outstanding backend for Big Data-scale interactive dashboards/BI. Achieving top performance however requires a deep understanding of the product architecture and experience in fine tuning of Vertica.

Architect281, Architect at a hospitality company with 100-1000 employees:

We have been able to move from nightly batch loads to continuous data flow and usage. This hasn't happened just because of Vertica, we have renewed our data platform pretty thoroughly, but definitely Vertica is one major part of our new data platform.

SoftwareEng525, Software Engineer at a marketing services firm with 1-100 employees:

This product has enabled us to keep very large amounts of data at hand for fast querying. With enough hardware force behind it, we were able to use Vertica as our primary reporting database without having to aggregate data, thus enabling us to provide many reports without having duplicated data or large aggregation steps.

SeniorSolArch232, Senior Solution Architect at a comms service provider with 1000+ employees:

With Exadata, everything is simply better optimized, and our organization has been leveraging the fact that it can now manage larger volumes of data much more easily than it could previously. For example, we almost never hear our internal or external customers complain about performance issues these days. Performance problems were a recurring problem for us in the past.

Ruthwik Bankar, Project Manager at a tech services company with 1000+ employees:

It's one of the most stable, high-performing products I've used.

Carlos Anibal, Oracle Database Specialist at a comms service provider with 1000+ employees:

My company's data warehouse was consolidated on a single database with improved capability and performance. Maintenance tasks and jobs start to execute much faster on the new environment. Now our company is acquiring a new hardware to support our SAS application.

Kathy Duret, Oracle DBA at a retailer with 1000+ employees:

Many diverse applications were migrated to the Exadata platform that had either performance or stability issues. All have had very reliable and performance since. Their hardware engineers can’t be beat. They are the best people I have ever worked with and parts are installed very quickly and efficiently. They’re very easy to work with, schedule jobs with, and get things done. Also, they are very proactive as well.

Undrakh-Orgil Odonkhuu, Software Engineer at a tech services company with 100-1000 employees:

Sped up SQL response time Easy to get support from the vendor

Jason Arneil, Infrastructure Principal Director at a tech company with 1-100 employees:

In addition to the enhanced speed of operational readiness, the other major, but less well appreciated benefit, is one of standardisation. Many large organisations have a wide and varied oracle database, both in terms of hardware and software. A big advantage I have seen is standardisation onto the one hardware platform and a minimal number of software versions. This has improved operational effectiveness.

Arup Nanda, Principal Global Database Architect at a hospitality company with 1000+ employees:

Prior to Exadata, we were relying on daily ETLs from multiple source systems to load data into base tables, then we prepared a summary and materialized views from the base to data marts which were eventually exposed to the end users. This process used to take 18 hours. So, the data was always at least 18 hours late, and occasionally later, if the ETLs failed. This caused two issues: first, most managers simply were not able to make quick decisions because the data was stale and second, for most important functions, the users got the data from our OLTP system which taxed that system's capacity.

Arup Nanda, Principal Global Database Architect at a hospitality company with 1000+ employees:

After Exadata, we observed three immediate benefits: The ETLs finish in 3 hours (down from 18), making the data fresh and enabling the managers to make quick decisions. Campaigns used to take days but take a few hours now Users no longer go to the OLTP system, reducing the impact there and saving us from upgrades There is no need to build summaries; so some of the data is available almost immediately, allowing most decisions to be taken on near-realtime data which was impossible pre-Exadata

Julian Dontcheff, Managing Director at a tech services company with 1000+ employees:

In the past, Infrastructure units were mostly organized in teams like the Database team, the Storage team, the Backup team, the Unix team, the Monitoring team, etc. Now, with Exadata, they have become sort of obsolete. Only one unit is needed for administering the whole stack, which means a lot of savings on IT level.

ORACLE EXADATA REVIEW BY A REAL USER

Charles Kim *Verified by IT Central Station*
Founder and President at Viscosity North America

Valuable Features:

My top 4 most important features of Exadata are:

1. Smart Scan, the ability to offload intensive SQL workloads to the storage servers. Queries are offloaded to the storage layer and only the result sets of relevant data are returned to the database server thus significantly improving performance.

2. Exadata Hybrid Columnar Compression, where we can compress data from 10x to 50x. Deploying databases on the Exadata can significantly reduce the amount of storage that is needed.

3. Exadata Smart Flash Cache automatically moves data between DRAM, flash and spinning SAS or SATA disks to provide best performance.

4. Virtualization is a new feature introduced to the X5-2 family. Now Oracle's Database Machine can be catered to large enterprise mission critical databases and can house smaller databases that need isolation and now even application servers together with the database. We can connect data intensive applications to the database over low-latency, high throughput infiniband.

Improvements to My Organization:

We are able to provide a complete solution to our customers from data center installation/configuration, firmware upgrades, ILOM upgrades, OS patching/upgrades, cluster patching/upgrades and database patching/upgrades. We can upgrade the entire stack in a single evening with minimal outage. Depending on the customer's tolerance for downtime, we can either perform the patching or upgrade in a rolling fashion.

Patching and upgrade services have proven to be a value-add differentiator for SMB and mid-market customers where resources and budgets are often limited. A single consulting resource can patch/upgrade the entire Exadata stack because patching and maintenance has become reliable and simplified.

Room for Improvement:

For the initial instantiation process, the Oracle Exadata Deployment Assistant (OEDA) could have better error checking and pre-check validation as you navigate through the tool. When executing the OneCommand utility, generated logs are decent but the logs are not detailed enough to pinpoint to where the error occurred in the stack. Oracle can do a better job with error isolation. After the OneCommand, other one-off commands have to happen (i.e. we have to login to the infiniband and cisco switch). Ideally, OEM could be leveraged to configure the remaining components of the Exadata after the OneCommand. The idea would be to reduce the number of people required to support the stack. If we leverage OEM, we can leverage a single resource that minimally understands the stack to support the workflow.

WHAT REAL USERS ARE SAYING...

"A single consulting resource can patch/upgrade the entire stack because patching and maintenance has become reliable and simplified."

"Once you start leveraging it with Database In-Memory and Multitenant, you've got a solid platform for enterprise."

"Exadata can significantly improve performance but there's a learning curve in a few key areas."

"Features I Find Valuable Include: Smart Flash Cache, Smart Scan and the Robust Storage."

"It aptly addresses the problems of database consolidation. No more arguments about running mixed workloads together."

"Key benefits include consolidated databases into one platform and one environment."

Use of Solution:

I've been using Exadata since V2 in 2009.

Deployment Issues:

For new DBA's trying to understand the Exadata infrastructure, they struggle initially trying to understand all the components between ILOM, ASR, and OEM (what they manage, what they support, and what problem area that they detect) and which does what since they are overlapping. Consolidation management across the Exadata is another common issue on the Exadata. OEM is good at managing individual targets but not so great yet on identifying culprits across environments with heavily consolidated databases. When you look at the compute wholistically, it is difficult to identify what database is consuming most CPU and most I/O resources. Doing any kind of showback is difficult to do across instances.

Stability Issues:

I only see stability issues with Exadata when I see too much consolidation, and/or the Exadata is not sized properly. Often the customer tries to throw everything and anything on the Exadata, by over-parallelizing OLTP/batch processes without any resource management across any databases.

Scalability Issues:

With the Exadata, if I see an issue with scalability, it is typically goes back to being a sizing issue. The real question I have to ask is: did you get the right Exadata configuration for your database(s). If your Exadata configuration is sized properly, you should not have scalability issues.

If you let every database see every CPU on the Exadata compute node, you can potentially run into scalability issues. Customers who do not take advantage of database resource manager or IO Resource Manager (IORM) often run into performance issues in a consolidated environment. Likewise, if a customer tries to over-parallelize their application code, it can cause scalability issues. We tend to see more issues with improper management of parallel execution on the Exadata because it is perceived as something you can throw anything at.

Customer Service:

The Exadata stack is well known and has become a standard platform with Oracle customers; thus, triage to resolution has become much more streamlined. When you create a service ticket with Oracle Support, time to resolution is significantly reduced. On another note, Oracle Field Support engineers are remarkable. They have been quick to respond, flexible, knowledgeable, and willing to work with our schedule.

Technical Support:

Oracle Support Engineers are outstanding but you have to know how to navigate the system. Oracle Support provides all the technical support for the Exadata including the database software, OS, and hardware. Not knowing how to navigate through Oracle's support structure and escalation policies can leave you feeling stranded by the vendor. Logging a support ticket with any component of the Exadata is no different than logging a support ticket for a database issue. You need to know how to raise a severity for a service ticket and how to escalate a support issue with the duty manager when production issues occur. Having a seasoned DMA (database machine administrator) is crucial to a successful Exadata deployment.

Having said that, often with Exadata customers, they can create a service ticket with the wrong Oracle Support group. This can cause confusion and elongated response times at early stages of the Exadata deployment as the service ticket gets routed to different teams within Oracle Support.

Previous Solutions:

Lot of our customers come from best of breed technologies (UCS/Dell/HP, EMC/Hitachi) to choosing Exadata. Typically, our customers choose Exadata for pure performance in IOPs, throughput, and low latency for their database workloads; however, we have seen a trend of customers choose the Exadata platform because they are short staffed, have high rates of attrition, and thus, have inability to support the hardware and software technologies.

Initial Setup:

The initial setup is straightforward with the Exadata X5-2. We have to re-image the factory Exadata and leverage OneCommand for configuring the compute and storage nodes.

Implementation Team:

A lot of the initial setup is configured by the Oracle ACS (Advanced Customer Support) organization; however, Viscosity is a certified Exadata implementation specialist and often perform the initial installation and configuration at the customer data center. By the time the customer receives a fully configured Exadata, they are ready to deploy databases as a RAC or non-RAC database.

ROI:

For our customers, we are able to significantly reduce both CapEx and OpEx for customers for 3-5 years TCO. We are able to:

Accelerate implementation to meet the functionality demand from the business users

Deliver quick implementations to meet the market demands

Lower implementation costs

Cost and Licensing Advice:

Setup costs for the Exadata varies from customer to customer and depends on database size, number of databases, and number of applications. For our most recent customer in Dallas Texas, the cost rolled up to approximately 800k for 2 X Quarter racks, which include storage cell software, and 25k for setup services. To determine day-to-day cost, we estimated the cost to average about 7k per month for both QTR racks.

Other Solutions Considered:

A lot of our customers typically look at two other options. We have seen customers perform side-by-side evaluations by building their own high performance system with EMC/Pure Storage/Violin All Flash Array and UCS/HP/Dell blades or perform comparisons with a converged system such as VCE's Vblock. Customers typically ask for a proof-of-concept demo and run performance benchmarks with their own database and application to see the immediate impact and value-add for their organization. Our last 3 customers have compared Exadata with the Vblock.

Other Advice:

For new customers who are about to embark on the Exadata journey, they should consult with a vendor specializing in Exadata implementations for the first set of database migrations and technical direction. Customers should also do the proper sizing exercises either with Oracle or with the Exadata Specialty niche vendor to buy the suitable Exadata configuration what will suit their business needs for the next several years.

When purchasing Exadata, they should also look into purchasing either the ZFS Storage Appliance (ZFSSA) or the Zero Data Loss Recovery Appliance (ZDLRA) to offload their backups to leveraging Infiniband technology for maximum throughput.

OEM CC 12c provides a comprehensive monitoring and management of the Exadata platform. Not only can OEM monitor and maintain at the hardware level for compute, storage and network but also at the OS, cluster and database level. OEM Cloud Control can monitor all components of the Exadata.

ORACLE EXADATA REVIEW BY A REAL USER

Nitin Vengurlekar *Verified by IT Central Station*
CTO/Architect at Viscosity North America

Valuable Features:

The most valuable feature is definitely the performance. However, it's also a platform for both performance and consolidation. Once you start leveraging Exadata with Database In-Memory and Multitenant, you've got a really solid platform for database enterprise computing.

Improvements to My Organization:

If you have server footprint reduction or consolidation as a requirement, then you can consolidate several databases and configurations onto one platform. Combine that with performance, it becomes a great avenue for cost reduction in terms of OpEx and CapEx. Exadata cannot be one-dimensional solution, performance, consolidation, and scalability must be part of the overall solution.

Room for Improvement:

I think Oracle Enterprise Manager has come a long way with monitoring and managing Exadata, but getting the patching down would go a long way. I think the patching is still lacking quite a bit. If we can get the patching and upgrade in place, that would make a huge impact in overall supportability. I realize there is new OEM functionality recently introduced to support this, but it needs to be flushed out and tested.

Stability Issues:

It's been great. The stability has really, really improved a lot over the last couple of years.

Scalability Issues:

It's been fantastic. With all the additional flash, with the faster CPUs, the faster disks, it's really come a long ways. The introduction of the X6 is also going to be an interesting avenue.

Technical Support:

Tech support has been really good. That goes a long way because you have some comfort, stability in the fact that you've got somebody that you can lean on. You don't always have those skill sets. Since we're a services company, we do have those skill sets in that area, but most customers may not have that.

These guys are very well versed in the technology. They understand not only the storage software part of it, but the database as well as the hardware aspect of it -end to end. The most surprising aspect of Exadata is really the support.

Initial Setup:

It is quite difficult, but again since we're a services company, we're pretty well versed in it. We actually help customers do that end-to-end. Not only do we help bring Exadata in, we help do the installation and configuration and migration to the Exadata, as well as sustainability for future.

It takes a lot of planning depending on your downtime for conversion, depending where [platform] you're coming from. All those things play a big part in what tools you use and how you use them.

Cost and Licensing Advice:

It is an expensive stack, so there's those things to keep in mind. The upside is that once you've bought into the technology, you have to embrace it , if that's the direction you want to go.

Other Advice:

It is really good for what it's supposed to do. It's a great technology for Oracle-based solutions. There's pluses and minuses in everything. Any time you get an engineered system like that, you have to be geared for organizational [boundary] changes. Understand how the technology is supposed to work, how it's going to change the way you do your day-to-day business.

ORACLE EXADATA REVIEW BY A REAL USER

Chris Bradham *Verified by IT Central Station*
Technical Director at a consultancy with 1-100 employees

Valuable Features:

We run a mixture of OLTP and Data Warehouse with the majority being OLTP. With that being said, the feature that provides us the most performance gains is the Smart Flash Cache for the OLTP databases. The "offloading" capabilities provide the biggest performance gains for DW such as smart scans and storage indexes. There is a new security feature which allows disabling ssh to the storage servers which will make my security folks very happy.

Improvements to My Organization:

We have a number of statistics collected before cutover on our legacy environment compared to Exadata. Without doing anything other than copying the data across, we saw significant performance gains for most key processes. We receive feedback from users stating how fast the performance is compared to other systems. Performance issues are few and far between. Our database environment is extremely stable compared to the legacy DB configuration. We upgraded from a X2-2 quarter rack to a X5-2 eighth rack and experienced significant performance gains.

Room for Improvement:

The patching process is a continually evolving. It's changed drastically over the years and foresee some continued refinement on the horizon, especially when comparing the process to something like ODA. We have one one-off patch which requires a SR with Oracle Support every time the quarterly patch comes out. This is a bit of a pain but I do know sometimes these one-offs get rolled into the quarterly patch or in our case the next version (11.2.0.4).

Use of Solution:

Over 3 years

Deployment Issues:

The deployment went extremely smooth. I think this was in large part to not trying to modify anything, just doing an "AS IS" data pump copy of the data of our OLTP databases. We made some changes afterword to better leverage the Flash Cache and Hybrid Columnar Compression where it made sense to do this. For Data Warehouses, index strategies should be reevaluated to take advantage of smart scans. Indexes on primary keys only, would be a good starting point

Stability Issues:

There is redundancy built throughout the Exadata so even when we've experienced a disk failure, it's a very low stress situation. Early on we had some performance issues with DBFS and a node eviction problem. DBFS was resolved through a combination of settings changes and a quarterly patch. The node eviction was resolved through a one-off patch that eventually got rolled into a quarterly patch. I would chalk up these issues to being early adopters. We do have an occasional bug but I can't think of any that would be unique to Exadata with the database software. At least this provides some degree of comfort that Exadata is not the source of the issue.

Scalability Issues:

The key for scalability is selecting the appropriate disk configuration and the proper size rack configuration. The two options are High Capacity and High Performance. If ever in doubt, always go with High Capacity. The performance difference is negligible at best, however having the extra space allows for more consolidation. That's the entire point of Exadata, to consolidate databases. We've added a few databases to the Exadata since we originally started to use the environment and there has been no performance impact. In our case, a Quarter rack was appropriate but for larger environments, this may not be enough.

Customer Service:

In terms of overall Oracle customer service, we've had good experiences on this front. Oracle has provided us access to their experts and continually check to see how things are going. Whenever an issue comes up, they treat the problem seriously. Since we support a government customer, Oracle is extra motivated to ensuring we have a successful experience.

Technical Support:

On the hardware side, customer service is quite good. Any disk failures get replaced in a day and with triple redundancy for disk, it's not been a concern. Software customer service has improved over the years. Early on was a little rough as I will say the software wasn't fully mature. As the product has matured, so has the software support's capability to resolve issues more quickly. We can't take advantage of ASR, however this seems like a major improvement for customer service in terms of responsiveness.

Previous Solutions:

We didn't switch, we were doing a technology refresh and went with Exadata.

Initial Setup:

There is a definite learning curve initially. We had to learn about migration options, shared mount point options, how to integrate with Cloud Control, patching, health check, how to optimize, and how to harden the Exadata environment. Since we went live, many more folks use Exadata so there's more how to's and best practice documents available so the learning curve isn't nearly as steep. We learned a lot in the process and now have a tremendous amount of expertise in setting up, configuring, optimizing and maintaining the Exadata.

Implementation Team:

We implement Exadata in-house and have gone through several migration methodologies.

Cost and Licensing Advice:

We had ACS perform the initial Startup Pack, however there are companies that can do this much cheaper to lower the original setup cost, such as ours. Day-to-day cost is greatly reduced compared to our legacy environment as we no longer have to serve as "fire fighters." In terms of pricing, Exadata is probably not going to be the lowest cost option. There is a price to pay for performance and stability. With that being said, I have not heard of any customers who have regretted the purchase and/or looking to get off the technology. On the contrary, I can't imagine going to another solution at this point and trying to justify this with the user community in terms of why the system performance degraded. Can't imagine that would go over too well.

Other Solutions Considered:

We had a custom solution and evaluated Exadata versus the custom solution. Exadata was actually a cheaper solution due to the number of cores. Oracle software licenses are based on processor so if comparing a Quarter Rack versus a 4+ four node custom solution, Exadata may win out from this perspective. We were looking at a 5 node RAC which would have doubled the cost of our software licenses when compared to the equivalent with a Quarter rack of Exadata. Besides, the performance metrics indicated Exadata would easily outperform the custom solution which made our decision a no brainer.

Other Advice:

Exadata is a powerful solution. As I mentioned there is a learning curve. Working with a company that has experience with Exadata can help avoid potential pain points and maximize the ROI.

ORACLE EXADATA REVIEW BY A REAL USER

Manish Nashikkar *Verified by IT Central Station*
ITA - Oracle Apps DBA at Tata Consultancy Services

Valuable Features:

Smart Scan

Smart Flash Cache

High Availability

Consolidation

Robust Storage

Computing

Improvements to My Organization:

Because of Exadata Systems, we were able to consolidate all the applications databases into one, e.g, Oracle EBS, Siebel, Hyperion etc.

Room for Improvement:

Exadata Linux systems have Intel CPUs inside. I would suggest that if Oracle could work together with Intel to have some more intelligence at the CPU level, then there would be nothing like it.

Use of Solution:

I've been using it since June 2013.

Deployment Issues:

We didn't have any deployment issues.

Stability Issues:

After three to four months we may need to perform a rolling reboot of the DB/storage servers.

Scalability Issues:

I would say it's the best solution in terms of scalability.

Customer Service:

8/10.

Technical Support:

9/10.

Previous Solutions:

We used both EMC Storage and IBM Servers.

Initial Setup:

The setup was simple, however, the database migration was a bit challenging, but there was no show stopper anywhere.

Implementation Team:

We are a vendor TCS (Tata Consultancy Services) and have as a customer the Taj Group of Hotels (IHCL), where we deployed and implemented Exadata Systems.

ROI:

It's almost 100%.

Other Solutions Considered:

IBM

HP

Other Advice:

Go for it if you are looking for consolidation/scalability.

HPE VERTICA REVIEW BY A REAL USER

Joe Gonzalez *Verified by IT Central Station*
BI Manager, Vertica ASE Certified DBA at a marketing services firm with 1000+ employees

Valuable Features:

I have found great use out of many features, most notably the Management Console and the Database Designer. Many people with lots of experience creating table projections can get frustrated trying to optimize some complex queries, however, in Vertica, the Database Designer is normally a big help in these situations. You can feed it your problem queries and it will make projection design suggestions for you. The ability to have multiple projections on a table to work with different queries is a big bonus.

With the Management Console, I am able to easily monitor the system health without having to check individual system tables. The ability to view running queries and cancel problem ones from the Management Console is a very nice feature, and one that I requested when we first started using version 7.0 which had an early version of the Management Console included.

Improvements to My Organization:

The fast columnar store database structure allows our query times to be at least 10x faster than on any other database. This enables us to get answers to data questions as well as numerous analytics on our data out to our internal and external clients quickly. We are able to integrate our Vertica data warehouse with Tableau to create numerous reports quickly and efficiently. What was once a two year backlog of report requests on our old data system has been virtually eliminated now that we are using Vertica to provide the solutions.

Room for Improvement:

Most of the changes I need are in 7.2. We have not yet been able to upgrade to it, but are planning on doing so by the end of the year. Some of the things that are supposed to have been improved/added to Vertica in 7.2 are -

Better support for Live Aggregate Projections, including support for the DISTINCT function

The ability to restore a single object from a full backup instead of having to do an object level backup first

The ability to more easily restore a backup to an alternate cluster

The addition of numerous hints to customize your queries

The ability to enforce check constraints on table load

WHAT REAL USERS ARE SAYING...

"The ability to view running queries and cancel problem ones from the Management Console is a very nice feature."

"The concurrency got better in this version and we are able to run more queries and load concurrently."

"Columnar storage makes 'hot data' available much faster than a traditional RDBMS solution."

"We built a custom analytical tools on top of Vertica."

"Data Warehouse response times have decreased. It doesn't support stored procedures in the way we are used to thinking of them. "

"I liked the auto-distribution to all nodes for fault tolerance and query performance."

Major upgrades to the Query Optimizer

Use of Solution:

I have been using HPE Vertica for about three and a half years now.

Deployment Issues:

The installation and deployment was fairly straightforward, and was taken care of by our Linux admins.

Stability Issues:

Our system is very stable.

Scalability Issues:

Adding one node was very easy, as was adding memory to all nodes. We are currently in the process of setting up a Dev / DR environment which is going very smoothly.

Customer Service:

I have a great relationship with Vertica customer support. They are friendly, knowledgeable, and are quick to respond.

Technical Support:

HPE Professional Services have also been a huge help to us when needed. They are well worth the investment.

Previous Solutions:

At my last company, we initially used Aster Data (now owned by Teradata). Once our database grew too large, it was unable to handle the number of transactions we were completing per day. Many queries on our largest table were taking from 20 minutes to over an hour to complete. Right out of the box, our longest queries went down to under a minute, most completing in a matter of seconds.

Initial Setup:

The initial setup was straightforward. We used an HPE-affiliated vendor to purchase and properly set up the equipment, completed a POC, and then we had HPE Professional Services assist with the transition from our old system to Vertica.

Implementation Team:

We implemented through a vendor. I highly recommend using IIS, they are amazing.

Cost and Licensing Advice:

The pricing, based on raw TB of data stored, is fair and affordable. You can have multiple projections per table without incurring a cost beyond the initial data load. The fact that a Dev and a DR cluster are included in the license cost is a great value!

Other Solutions Considered:

We evaluated Vertica and Greenplum, and chose Vertica due to cost and a number of existing use cases that were nearly identical to ours.

Other Advice:

My only advice is to seriously consider using Vertica for your data warehouse needs. I have normally just gone with the flow and learned whatever tools our company chose. When we switched from Aster Data to Vertica, I made the initial recommendation to do so. I am so happy with this product that I am now an HP ASE Certified Vertica Administrator, and moved to a new job that is also using Vertica. I would not have changed jobs if I were not able to continue using this product. I am also recommending to management that we purchase HP IDOL for our upcoming audio and video analytics needs. HP Big Data Solutions is a great product suite, and I have bet my career on its future growth.

HPE VERTICA REVIEW BY A REAL USER

SoftwareDataArch072 *Verified by IT Central Station*
Software and Data Architect at a software R&D company
with 1000+ employees

Valuable Features:

The compute and processing engine returns the queries fast and let us use our analysis resources in a better utilization.

The concurrency got better in this version and we are able to run more queries and load concurrently.

Improvements to My Organization:

We built an internal dashboard using the MicroStrategy to increase visibility to our management and our employees. Also, we built tool to expose the data to our selected partners and users to create better engagement with our platform.

Room for Improvement:

Loading times for “real time” sources - for example, loading from Spark creates a load on the DB at high scale

Connectors to other sources such as Kafka or AWS Kinesis

Better monitoring tools

Better integration with cloud providers - we were missing some documentation regarding running Vertica on AWS

Use of Solution:

We've been using Vertica for a year.

Stability Issues:

In case of one HD failure in the cluster, the entire cluster got slower. We feel that it should be able to handle such issues.

Scalability Issues:

No.

Technical Support:

The support was slow and didn’t provide a solution in most cases. The community proved to be the better source for knowledge and problem solving.

Initial Setup:

Pretty straightforward, the installation was simple and we added more nodes easily as we grew.

Cost and Licensing Advice:

Vertica is pretty expensive, take into account the servers and network costs before committing.

Other Solutions Considered:

We evaluated both AWS Redshift and Google BigQuery.

Redshift didn’t fulfill our expectations regarding query latency at high scale (over 60 TB). Regarding BigQuery, we found the pricing structure pretty complex (payment per query and data processed) and harder to control.

Other Advice:

Don't plan a production usage on high-scale straight on Vertica, use caching or other buffers between the users and the DB. Get yourself familiar with the DB architecture before planing your model (specifically, make sure you know ROS/WOS and projections). Try to avoid LAP before your schema gets stabilized.

HPE VERTICA REVIEW BY A REAL USER

Matt Williams *Verified by IT Central Station*
Lead Data Scientist Machine Learning at a financial services firm with 1-100 employees

Valuable Features:

Columnar storage makes 'hot data' available much faster than a traditional RDBMS solution. Also, Vertica scales up quickly and maintains good performance.

Improvements to My Organization:

Performance management of high-traffic sites - Vertica's ease of scaling has been invaluable for one of our main customers.

Room for Improvement:

I'm concerned that HP Enterprise has sold their software business, and worry about future investment to enhance predictive/machine-learning capabilities.

Use of Solution:

3 years.

Stability Issues:

Not really.... Vertica shines on stability.

Scalability Issues:

No, scalability is also a strength of the solution.

Technical Support:

9 out of 10. HPE has some excellent engineers who are eager to help us make Vertica work well.

Previous Solutions:

I've been a 'full stack' data expert for years, started on Oracle and SQL Server, moved to Hadoop, Mongo, etc, but Vertica was the right fit for large enterprises with high performance demands and ease of scalability.

Initial Setup:

Initial setup is a bit clunky, like most complex, tunable products can be.

Cost and Licensing Advice:

Negotiate when their fiscal year is about to close :)

Other Advice:

It's a solid product that keeps its promises. I do worry about HP Enterprise's sale of Vertica to Micro-Focus

Rating: 8/10 - it works very well, but some customers worry about 'Vendor lock-in'.

HPE VERTICA REVIEW BY A REAL USER

Rodrick Brown *Verified by IT Central Station*
Development Operations/SRE at a software R&D company
with 100-1000 employees

Valuable Features:

HA Clustering

Speed / Performance

Improvements to My Organization:

We're able to retrieve queries nearly instantaneous for our custom analytical tools we built on top of Vertica.

Room for Improvement:

More frequent updates.

Use of Solution:

1 year

Stability Issues:

None.

Scalability Issues:

None.

Technical Support:

Very knowledgable team which has provided excellent documentation for every issue we've had to troubleshoot.

Previous Solutions:

MonetDB -- unstable, frequent crashes.

Initial Setup:

Straightforward, was able to get the database up fairly quickly with minimal effort.

Cost and Licensing Advice:

We're still using the Community Edition (CE).

Other Solutions Considered:

MonetDB, Cassandra, Amazon RedShift.

Other Advice:

Great product, very mature and robust. Vertica is able to scale to meet our demands as we scale our business 10x.

Oracle Exadata and HPE Vertica Alternatives

VENDOR (BY RANKING)	SOLUTION	
Hewlett Packard Enterprise	HPE Vertica	Rating: 8.06 (40 reviews)
Oracle	Oracle Exadata	Rating: 8.8 (32 reviews)
EMC	EMC Greenplum	Rating: 7.82 (20 reviews)
IBM	Netezza	Rating: 8.06 (20 reviews)
SAP	Sybase IQ	Rating: 8.0 (7 reviews)
Microsoft	Microsoft Parallel Data Warehouse	Rating: 7.0 (3 reviews)
Infobright	Infobright	Rating: 9.0 (2 reviews)
Teradata	Teradata	Rating: 8.0 (2 reviews)
Microsoft	Microsoft Analytics Platform System	(2 reviews)
Oracle	Oracle Database Appliance	Rating: 8.5 (2 reviews)
SAP	SAP NetWeaver Business Warehouse	Rating: 5.5 (2 reviews)
Actian	Actian ParAccel	Rating: 7.0 (1 review)
Apache	Apache Hadoop	Rating: 9.0 (1 review)
IBM	IBM InfoSphere Warehouse	Rating: 8.0 (1 review)
IBM	IBM Smart Analytics System	Rating: 8.0 (1 review)
Oracle	Oracle Big Data Appliance	Rating: 7.0 (1 review)
Actian	Actian Ingres	
Teradata	Aster Data Map Reduce	
Cazena	Cazena	
Dell	Dell Quickstart Data Warehouse Appliance	

To read more reviews about Data Warehouse, please visit:
http://www.itcentralstation.com/products/comparisons/oracle-exadata_vs_hpe-vertica

	InfoWorks	InfoWorks
	Kognitio	Kognitio WX2
	SAND Technology	SAND
	Snowflake Computing	Snowflake Computing
	TIBCO	TIBCO Live Datamart
	Treasure Data	Treasure Data

Top Data Warehouse Vendors

Over 176,460 professionals have used IT Central Station research on enterprise tech. Here are the top Data Warehouse vendors based on product reviews, ratings, and comparisons. All reviews and ratings are from real users, validated by our triple authentication process.

1. HPE Vertica

2. Oracle Exadata

3. Teradata

4. Netezza

5. EMC Greenplum

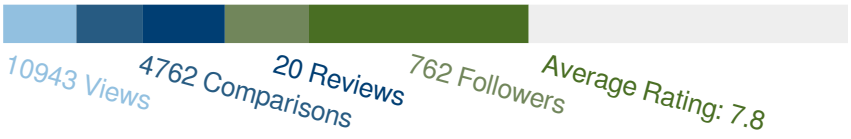

6. Sybase IQ

To read more reviews about Data Warehouse, please visit:
http://www.itcentralstation.com/products/comparisons/oracle-exadata_vs_hpe-vertica

7. Microsoft Parallel Data Warehouse

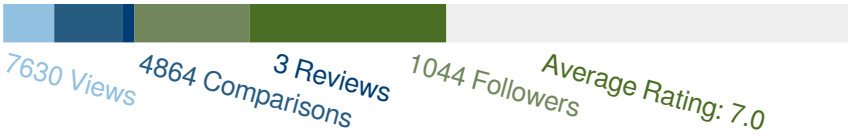

8. Infobright

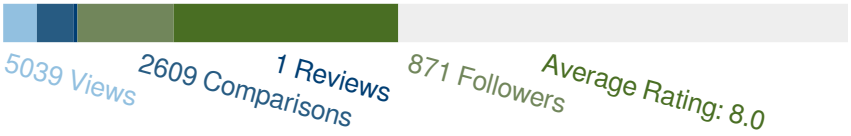

9. Apache Hadoop

10. IBM Smart Analytics System

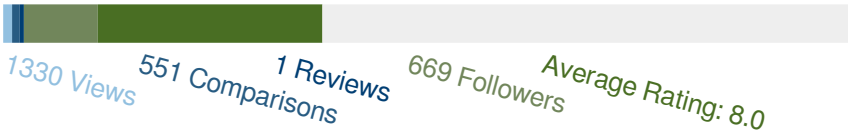

11. IBM InfoSphere Warehouse

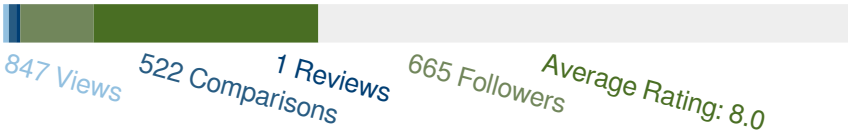

12. Actian ParAccel

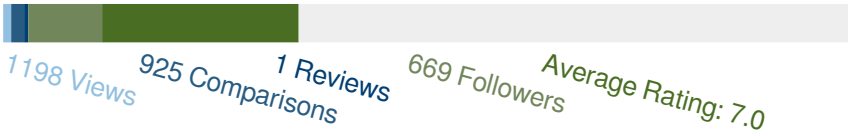

13. Oracle Database Appliance

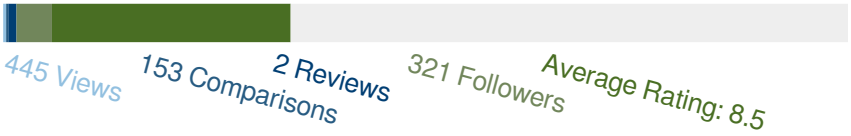

To read more reviews about Data Warehouse, please visit:
http://www.itcentralstation.com/products/comparisons/oracle-exadata_vs_hpe-vertica

14. Oracle Big Data Appliance

806 Views 228 Comparisons 1 Reviews 418 Followers Average Rating: 7.0

15. SAP NetWeaver Business Warehouse

1525 Views 545 Comparisons 1 Reviews 687 Followers Average Rating: 3.0

16. Snowflake Computing

1946 Views 1295 Comparisons 0 Reviews 690 Followers Average Rating: 0.0

17. Actian Ingres

672 Views 538 Comparisons 0 Reviews 752 Followers Average Rating: 0.0

18. Aster Data Map Reduce

625 Views 516 Comparisons 0 Reviews 669 Followers Average Rating: 0.0

19. Kognitio WX2

595 Views 456 Comparisons 0 Reviews 667 Followers Average Rating: 0.0

20. SAND

309 Views 222 Comparisons 0 Reviews 666 Followers Average Rating: 0.0

To read more reviews about Data Warehouse, please visit:
http://www.itcentralstation.com/products/comparisons/oracle-exadata_vs_hpe-vertica

21. Microsoft Analytics Platform System

1061 Views 599 Comparisons 2 Reviews 494 Followers Average Rating: 0.0

22. TIBCO Live Datamart

505 Views 380 Comparisons 0 Reviews 597 Followers Average Rating: 0.0

23. Cazena

274 Views 207 Comparisons 0 Reviews 479 Followers Average Rating: 0.0

24. Treasure Data

264 Views 196 Comparisons 0 Reviews 472 Followers Average Rating: 0.0

25. InfoWorks

142 Views 99 Comparisons 0 Reviews 400 Followers Average Rating: 0.0

26. Dell Quickstart Data Warehouse Appliance

138 Views 114 Comparisons 0 Reviews 316 Followers Average Rating: 0.0

Chart Key

Views: Number of total page views

Comparisons: Number of times compared to another product

Reviews: Total number of reviews on IT Central Station

Followers: Number of followers on IT Central Station

Average Rating: Average rating based on reviews

*The total ranking of a product (i.e. bar length) is based on a weighted aggregate ranking of that product's Views (weighting factor = 17.5%), Comparisons (17.5%), Reviews (17.5%), Followers (17.5%), and Average Rating (30%). Reviews and ratings that are more than 24 months old, as well as those by resellers, are excluded from the rankings. For each ranking factor, the score (i.e. bar segment length) is calculated as a product of the weighting factor and its position for that ranking factor. For example, if a product has 80% of the number of reviews compared to the product with the most reviews in its category, then the product's bar length for reviews would be 17.5% (weighting factor) * 80%.*

To read more reviews about Data Warehouse, please visit:
http://www.itcentralstation.com/products/comparisons/oracle-exadata_vs_hpe-vertica

TOP 5
Based on Views

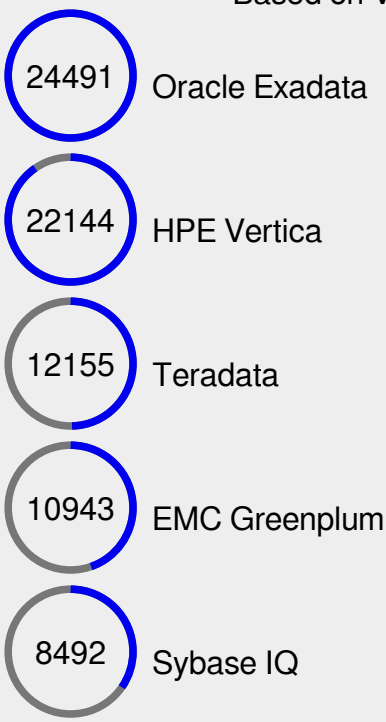

TOP 5
Based on Views of Comparisons

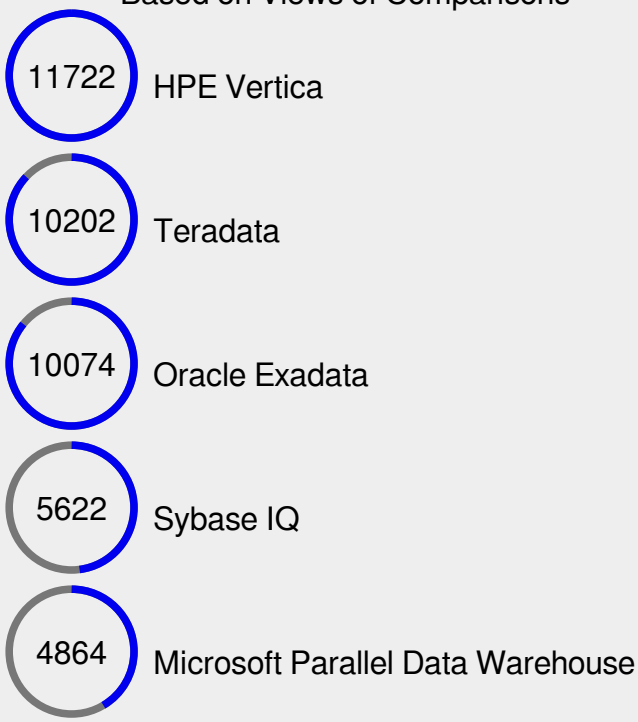

TOP 5
Based on Followers

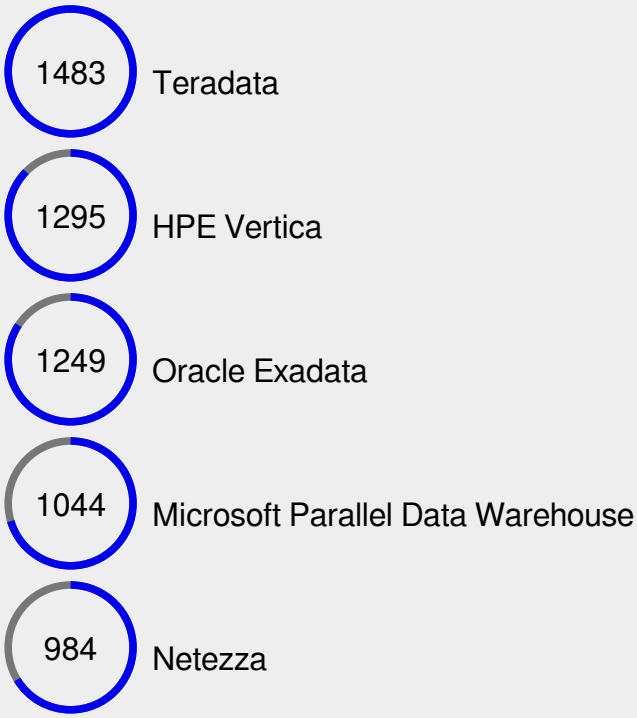

TOP 5
Based on Rating

To read more reviews about Data Warehouse, please visit:
http://www.itcentralstation.com/products/comparisons/oracle-exadata_vs_hpe-vertica

Join the IT Central Station Community

This report features just a few of the Data Warehouse reviews on IT Central Station. If you would like to read more about what real users are saying about the many Data Warehouse solutions on the market, visit the Data Warehouse page at IT Central Station:

http://www.itcentralstation.com/category/data-warehouse

There are a number of ways you can participate in the IT Central Station community. Write a review, read a comment, or just follow a product. Either way, we'll be sure to let you know when people are talking about the solutions you care about!

Give to Get

Did you find this whitepaper helpful? At IT Central Station, our philosophy is "Give to Get". Our active community and unbiased reviews are made possible by your participation and as such, we ask that you share your expertise with us as well. Please expect to be contacted shortly by an IT Central Station Community Manager. We will ask you for 10 minutes of your time to review a product that you use either via a 10 minute phone interview or questionnaire. You can choose to review anonymously or not and your company name will not be included in the review.

If you found this report and/or the reviews on IT Central Station useful, we would greatly appreciate your participation in giving back to our community.