



Author: Marc Staimer

W H I T E P A P E R

Oracle Exadata X8: Why it Spells **Doom** for DIY Database Systems

Oracle Exadata X8: Why it Spells Doom for DIY Database Systems

Executive Summary

Business is all about time. Time is the construct that separates events and moments. Time is a non-renewable resource. It is perhaps the most precious resource an organization has. Once used, it can never be reclaimed. Time only marches in one direction...forward. Once spent it cannot be reclaimed or rewound. For IT organizations, there is never enough time. Time is not a friend. And as every business knows, time is money.

Time is the fundamental factor when it comes to databases. It is tied to everything in each and every database. There is the time required to set it up; index it; operate it; manage it; tune it (ongoing); protect it from outages of all kinds; create, patch, and modify SQL statements; troubleshoot and fix when something breaks. Plus, the additional time required to train, become knowledgeable, experienced, and skillful in each type of database. Make no mistake, most organizations have more than one type of database.

Then there is the time most IT pros think of when referring to databases and that's the amount of time consumed in performance. Time per query. Amount of database operations and processes completed in a given period of time. Time per analytical process. Time per batch process. Database throughput in a given period of time. Database backup in a given period of time. Database restores from outages in a given period of time (more commonly referred to as recovery time objectives or RTO.) Database response time to the application. End-to-end latency a.k.a. delay or the amount of time it takes for a process to complete end-to-end. Time to move data between databases (extract transfer and load or ETL) for different types of analytics. All are critically important.

And yet, many database vendors claim that their database is the fastest in the market. What they actually mean is that they are applying a very narrow specific database use case in a specific configuration. They are not claiming all of the other database factors consuming gobs of time are the fastest. They are in fact ignoring them in their claims. Installing, implementing, managing, operating, tuning, troubleshooting, patching, protecting, etc., the database is never mentioned in their claims.

And none of that takes into consideration the hardware. Every kind of database requires hardware to run on. And yes, databases can run in a VM or container, as long as the use is relatively light. Databases are resource-intensive. Performance tuning, problem solving, patching, and scalability can be exceedingly difficult to troubleshoot when implemented in virtual machines or containers. It is recommended for mission-critical databases to run on their own hardware. The time required to install, configure, troubleshoot, correct errors, manage, operate, tune, patch, tech refresh, and maintain the hardware dwarfs the time required for the database. Expertise is required for the server hardware, operating system, storage networking, networking and storage. Finding an administrator with all those skills is the exception. It is more typical to have multiple administrators for the different disciplines. Coordination between those administrators is a massive challenge that consumes more time.

Reducing or saving time is the central theme that underpins the Oracle Exadata architecture, software, features, functions, and advancements. Time savings as directly resulting from: significantly higher performance; automation of manually labor-intensive tasks; artificial intelligence machine learning (AI/ML) that enables highly skilled, knowledgeable, and experienced database administrators to focus on more strategic projects; moving compute executables closer to the data instead of moving the much larger data to the compute executables; unmatched scalability simplifying implementation, operations, and management; incorporation of multiple database types that eradicate time consuming ETLs (extract transfer and load processes) for different types of analytics; and database consolidation with pluggable databases (PDB) and container databases (CDB).

Everything in Oracle Exadata from soup to nuts is designed to **SAVE TIME AND COST**. It starts with delivering more actionable information quicker (saves time, reduces risk, and cost, while increasing revenues); accelerates time-to-market (saves time, increases revenues and market share); enhances agility and flexibility faster (saves time and reduces costs); reduces wasted non-productive time (saves time and reduces costs); and increases organizational profitability.

Saved time equals reduced costs, increased revenues, and increased EBITDA – earnings before interest, taxes, depreciation, and amortization (profits). This paper examines how the 9th generation of Oracle Exadata is able to recapture a lot more time than any other database system, and do far more in less time.

Table of Contents

Executive Summary.....2

A Closer Look at the Latest Oracle Exadata X8 Upgraded Hardware5

 Engineered Integration6

 Performance.....6

 Scalability.....6

Oracle Exadata X8 – Exclusive **Unique Software *i.e. Not Available in DIY DB Systems*.....7**

 Computational Storage Servers7

 Automatic I/O Reduction7

 Analytics Optimized Columnar Format7

 In-Memory Analytics to Storage7

 OLTP Bottleneck Elimination8

 Mixed Workloads and Unmatched Consolidation Performance Optimization8

 Simpler Analytics and Data Warehousing8

 Only Autonomous Database9

 Automatic Statistics & Indexing.....9

 Unsurpassed Availability.....9

 Superior Database Consolidation10

 Enhanced Cybersecurity.....10

All-inclusive Oracle Database 19c Running on Exadata X8..... 10

Dispelling DIY Database System Mythology 11

 Multi-national Multi-\$ Billion USD Retailer.....11

 Multi-national Multi-\$ Billion USD Financial Services Firm11

 Multi-\$ Billion USD Healthcare Organization11

Conclusion 12

For More Information on Oracle Exadata 12

A Closer Look at the Latest Oracle Exadata X8 Upgraded Hardware

There is a lot to like in the new Oracle Exadata X8-2 and X8-8. Starting with the hardware. Oracle continues their tradition of leveraging production proven state-of-the-art (SOTA) hardware and lowering costs.

Oracle Exadata X8 Upgraded Hardware	
New Database Servers (DBS)	<p>2 & 8-socket Intel 24-core Cascade Lake CPUs</p> <p>15% Faster clock rate</p> <p>“Spectre” & “Meltdown” mitigated in silicon</p> <p>No SW overhead to patch those known vulnerabilities</p> <p>Increased security</p>
New Storage Servers (SS)	<p>2-socket Intel 16-core Cascade Lake CPUs</p> <p>60% more cores available</p> <p>Offloads Oracle Exadata Database Servers for:</p> <ul style="list-style-type: none"> SQL XML JSON Encrypt/Decrypt RMAN Backup Filtering (BCT) Fast Data File Creation Many in-database Analytics & ML algorithms <p>Extreme All Flash (EF) Storage Servers</p> <p>High-Capacity (HC) Storage Servers</p> <p>40% higher capacity 14TB Helium HDDs</p>
New Extended Storage Server (XT)	<p>For infrequently accessed, regulatory/compliance, cool, or cold data</p> <p>Oracle Databases have visibility & seamless access to XT stored data</p> <p>High raw capacities</p> <ul style="list-style-type: none"> 10-15x usable vs raw capacity via Hybrid Columnar Compression (HCC) <p>Oracle Database processing without requiring HCC rehydration</p> <p>Same consistent security model & encryption as all Exadata storage tiers</p> <p>Neatly slides into the Exadata rack</p> <p>Same high performance RDMA Infiniband scale-out architecture</p> <p>Hadoop or object storage costs with Database Storage capabilities</p> <p>SS disk license can be forsworn to further reduce storage cost</p>

The question becomes how does the upgraded hardware save additional time, reduce costs, increase revenues, and put all DIY and other database systems to shame? Take a look.

Engineered Integration

It starts with engineered integration. Oracle Exadata X8 comes pre-packaged, pre-integrated, pre-configured, pre-tested, optimized, automated, and custom software implemented that makes the hardware and software aware of each other. Typical implementations take 2 days or less including cleanup. Although it can take as long as 3 days if there is a lot of testing, 2 days or less is most common. Compare that to a typical DIY database system implementation that are commonly measured in weeks, and just as often months. Time that the Oracle Exadata X8 is already in production speeding transactions and analytics. Time that can never be recovered in the DIY database system regardless of the database.



And there are some things that just are not available in a DIY database system regardless of the database. Things such as smart flash caching, smart scanning, smart flash logging, scalable in-database analytics and AI/ML.

Performance

Oracle Exadata X7 already has delivered the best database performance anywhere before the release of Oracle Exadata X8:

- ~ 350 GB/s throughput per rack
- Up to ~ 6.3 TB/s throughput per maxed out 18 racks
- ~ 4.8 million OLTP (8K) read IOPS per rack
- Up to ~ 86.4 million OLTP (8K) read IOPS per maxed out 18 racks
- Latency at ~.25 ms (250 microseconds) or less

Oracle Exadata X8 raises the bar even higher:


- **60%** more throughput at ~ **560 GB/s** per rack
- Up to ~ **10.08 TB/s** throughput per maxed out 18 racks
- **37%** more OLTP (8K) read IOPS per rack at ~ **6.57 million**
- Up to ~ **118.26 million** OLTP (8K) read IOPS per maxed out 18 racks
- Latency at ~.25 ms (250 microseconds) or less
 - For up to 3.5 million IOPS per rack
 - Or 63 million IOPS for a maxed-out system



That unparalleled higher Exadata performance – multiple orders of magnitude greater than any other DIY database system, even ones based on the Oracle Database – translates into more transactions and more throughput in less time. That in turn converts into faster response times, faster analytics, faster actionable information in less time, reduced time-to-decisions, and faster time-to-market. Those faster times mean higher productivity, lower costs, higher revenues, and higher profitability.

Scalability

Performance and performance scalability are a good start; however, capacity scalability is essential in keeping complexity, skills, knowledge, expertise, experience, and costs low while keeping admin manageability, flexibility, and productivity high. Oracle Exadata X8 does not disappoint.



Oracle Exadata X8 scales from 1/8th rack to 18 racks. And subsequently scales over 100X in both compute and storage capacity. Each rack has up to ~ **3 PB** raw HDD capacity and ~ **920 TB** of raw NVMe flash capacity. This is before HCC increases the raw capacities from ~ **10 - 15X**. At maximum scale, an Oracle Exadata X8 has up to ~ **48 PB** raw HDD capacity and up to ~ **16.56 PB** of raw NVMe flash capacity. Again, HCC increases the usable capacity by 10-15x.

Scalability is not a problem.


High scalability means fewer systems to manage. Fewer systems to manage requires significantly fewer pedantic tasks and thus takes much less time required for managing, operating, troubleshooting, patching, learning, training, etc. That saved time can be utilized more productively in application development leading to more revenues and reduce personnel costs.

This is the time savings just from the Exadata hardware engineering alone. The Exadata software has an equal or greater impact on time savings.

Oracle Exadata X8 – Exclusive Unique Software *i.e. Not Available in DIY DB Systems*

Oracle Exadata is the only database machine that is database-aware. The Oracle Database is the only database that is hardware-aware. All of which enables a level of hardware and software synergy with unique software capabilities simply not available anywhere else.

Computational Storage Servers




This exclusive capability enables extensive co-processing between the Oracle Exadata Database servers and the Oracle Exadata Storage servers. Oracle Exadata Smart Scan technology automatically offloads data-intensive and CPU-intensive SQL, XML, JSON, encryption/decryption, RMAN backup filter a.k.a. BCT, fast data file creation, in-database analytics and AI/ML to the Oracle Exadata Storage servers. Smart Scan delivers over 560 GB/s SQL scan throughput per rack while offloading database server CPUs.

By moving the compute executables closer to the data, processing is much faster since the data does not have to be moved. Only the results have to be moved to the database servers, which are a fraction of the size of the data being processed. It also frees up the database servers for more transactions and more complicated transactions. That reduced latency equates into much greater performance.

Automatic I/O Reduction

Oracle Exadata database-aware **Flash Caching** delivers exceptional NVMe flash performance at HDD capacities and lower cost. Oracle Exadata **Storage Indexes** eliminate I/O not relevant to a particular query, accelerating query response times. Faster response times turns into faster results, faster processes, faster actionable information, faster-time-to-market, lower costs and higher profits. In fact, Oracle Exadata is the only database system that can automatically generate and optimize storage indexes – more on that later.

Analytics Optimized Columnar Format




Oracle Exadata **Hybrid Columnar Compression (HCC)** radically reduces storage capacity requirements and costs. It also speeds up analytics by more than an order of magnitude (10-15x) because the Oracle Database analyzes the HCC data without rehydration. That massive reduction in data being analyzed provides results faster. Faster results lead to faster actionable information, decisions, time-to-market, lower costs, and more profits.

In-Memory Analytics to Storage

NVMe flash has pushed database performance bottlenecks into the database server CPU. Exadata solves this problem by first offloading SQL from the Oracle Exadata Database Server to the Storage Server. The Oracle Exadata Storage then automatically transforms the table data into in-memory database columnar formats within the Oracle Exadata Flash cache. This enables much faster vector processing for storage server queries.

Oracle Exadata utilizes NVMe flash as persistent non-volatile memory. This Oracle Exadata X8 exclusive capability works with both row format OLTP databases and HCC analytics databases concurrently. And it’s architecturally set up to take advantage of future non-volatile memory technologies generally referred to as Storage Class Memory or SCM.


Oracle Exadata In-Memory Analytics automatically accelerates both transactional and batch processes. The results are faster processing, faster analytics, shorter response



times, ability to handle much higher volumes, faster actionable information, faster time-to-market, lower costs, and greater profits.

OLTP Bottleneck Elimination

Oracle Exadata eliminates the common random I/O OLTP bottleneck. It does this by leveraging very low latency NVMe flash and RDMA Infiniband with Smart Flash Logging that optimizes OLTP logging to flash.



Oracle Exadata additionally eliminates OLTP stalls from failed or sick components. It does this with two exclusive processes: First, it quickly detects database and storage server failures in typically less than 3 seconds. This avoids the long timeouts

required for non-hardware aware databases that causes the system to hang. Oracle Exadata then employs sub-second redirection of I/Os around the sick devices, thus avoiding database timeout hangs. Fewer hangs results in considerably less non-productive downtime, lost revenues, and lost reputation. It also saves administrators time, increasing their productivity.


Another OLTP bottleneck Oracle Exadata eliminates is inter-node coordination. Direct-to-Wire Protocol gives 3x faster inter-node OLTP messaging. Smart Fusion Block Transfer eliminates log write on inter-node block moves. RDMA coordinates transactions between nodes. All of this equals faster OLTP. As previously noted, faster OLTP enables more of them in less time, faster results, faster revenues, lower costs, higher profits.

Other ways Oracle Exadata accelerates OLTP is with specialized Cache Fusion, RDMA Algorithms for communicating transaction information, memory optimized key-value style access for OLTP data, and optimized IoT ingest and loading. All of this translates into radically reduced Global Cache waits and much fast OLTP. That faster OLTP translates into faster results, revenues, and ultimately, profits.

Mixed Workloads and Unmatched Consolidation Performance Optimization

Oracle Exadata automatically prioritizes latency-sensitive operations such as critical network messages for locks, cache fusion, logging, and more. It also prioritizes OLTP I/O over analytic or batch I/O. Oracle Exadata additionally prioritizes important workloads based on user policies, prioritizing CPU and I/O by job, user, service, pluggable database (PDB), session, SQL, and more. For multiple databases and/or PDBs, Oracle Exadata automatically provides prioritization, isolation, and separation between multiple tenants.

This inimitable prioritization and isolation ensure the database performance required is the database performance delivered in a multi-workload and/or multi-tenant environment. That simplifies database consolidation without any loss of performance while saving management/operational time, hardware and supporting infrastructure costs.



Simpler Analytics and Data Warehousing

Oracle Exadata can maintain a Data Guard or more automated Active Data Guard standby databases even if non-logged direct loads are used such as data warehouse loads. Active Data Guard delivers rapid and automated NOLOGGING loads avoiding redo overhead. It automates data population into In-Memory columnar tables and In-Memory for external tables speeds up analytics on external data.

That automation and simplicity accelerates analytics and data mining producing faster results, faster actionable information, faster time-to-market, no time spent tuning (saves time), and faster profits.

Only Autonomous Database



Oracle Exadata is the only hardware system running the **Oracle Autonomous Database** in the Oracle Cloud. The Oracle Autonomous Database is currently the only level 4 autonomy database on the market today¹ and handles the vast majority of database operations automatically. This frees up software engineers to simply build and deploy their applications. They no longer have to be concerned, knowledgeable, skilled, or experienced in operational tasks such as monitoring, backup, or system tuning.

The Oracle Autonomous Database is based on built-in Oracle artificial intelligence and machine learning (AI/ML). The AI/ML tools are now part of the Oracle Database 19c Enterprise Edition. However, the level 4 autonomy is only available in the Oracle Cloud. In other words, no DIY database based on Oracle or anyone else comes with this autonomous capability.

Database autonomy saves enormous amounts of time in database implementations, operations, management, indexing, tuning, debugging, troubleshooting, patching, backing up, learning, and just about everything to do with the database. It's self-managing, self-indexing, self-error detecting, self-patching, self-healing, and self-learning. The Oracle Autonomous Database is the ideal location for Oracle to develop advanced capabilities, many of which later become available over time in on-premises Oracle Exadata environments. Many capabilities that started out as part of the Oracle Autonomous Database and are now available with the Oracle Exadata on-prem. Some examples include:

- Automatic detection and recovery from sick or failed servers, storage, switches, or links.
- Real-time hang management, anomaly detection, maintenance slot identification, bug identification and prioritization in database operations.
- Query optimization, real-time statistics and automatic indexing (more on these in the next section) in workload operations.

The time saved by autonomy is better spent developing applications, working on and completing projects that there's never enough time to get done, getting to market faster, reducing costs, increasing revenues, and profits.

It is important to note that bright industrious DBAs can leverage the extensive AI/ML tools in the Oracle Database 19c to duplicate, with substantial time and effort, much of what comes automatically with the Oracle Autonomous Database.

Automatic Statistics & Indexing

Oracle Exadata optimizer statistics are gathered in real-time as DML executes. There are no stale statistics.

Oracle AI/ML implements indexes automatically, adjusting quickly to changing workloads. Reinforcement learning endows Oracle Exadata the ability to learn from its own actions. Candidate indexes are validated before being implemented.

Oracle Exadata automatic indexing is able to create better, faster, more efficient, and more usable indexes in seconds. These indexes have proven to be more efficient than those created by experts and extremely knowledgeable database administrators with years of experience.




Unsurpassed Availability

It's often said that the best ability is availability. Oracle Exadata is the only database system to achieve IDC's AL4 fault tolerance. AL4 fault tolerance means faults and failures are not noticeable by users because there

¹ **Level 0:** No Automation. Zero autonomy, DBA performs all tasks. **Level 1:** DBA Assistance. DB is controlled by DBA, but some DB assist features may be included in the DB. **Level 2:** Partial Automation: DB has combined automation functions, but the DBA must remain engaged and monitor the environment at all times. **Level 3:** Conditional Automation. DBA is a necessity, but not required to monitor the environment. DBA must be ready to take control of the DB at all times without notice. **Level 4:** High Automation. DB is capable of performing most operations automatically without DBA assistance including tuning and indexing. DBA has option to take control of the DB. **Level 5:** Full Automation: DB is capable of performing ALL operations under all conditions without DBA assistance. DBA has option to take control of the DB.

is no loss of state or data. The database keeps running as if nothing at all happened. Only the administrator is aware of the faults or failures. AL4 fault tolerance is automatically configurable in the Oracle Exadata X8 when implemented in an Oracle Maximum Availability Architecture (MAA) Gold or higher environment.



Oracle Exadata utilizes and automates the intrinsic backup, recovery, and disaster recovery tools built in the Oracle Database including: RMAN; ability to recover a database, PDB, or CDB to any point-in-time; Flashback Database time enabling views or returns of database objects past states without point-in-time media recoveries; Zero Data Loss Recovery Appliance (ZDLRA) reducing RMAN backup time by up to 95%; and Oracle Data Guard synchronization of one or more database copies. It's this implementation of the Oracle Gold or Platinum MAA that delivers the AL4 fault tolerance.

AL4 in an Oracle Database DIY database system is not easy to accomplish at all. It is not automatically built-in. But it is doable with extensive time, effort, and cost, especially hardware costs as long as the Oracle MAA guidelines are followed.

Superior Database Consolidation

Oracle Exadata supports up to 4,000 pluggable databases (PDBs) with refreshable PDB copies per PDB store. The multi-tenant container database (CDB) of PDBs dynamically grows CDB memory (Oracle Database 19c), without disruption to the databases. Reduced downtime increases productivity, saves time, and increases profits.

Enhanced Cybersecurity

Advanced Intrusion Detection Environment (AIDE) detects and alerts when unknown changes to system software are made. Oracle Recovery Manager (RMAN) flags and does not back up un-recognizable database blocks such as those that are part of malware.

This enhanced cybersecurity is in addition to the best-in-class Oracle Database 19c automated database security including core database security for users, roles, authentication, etc.; Network Communication Encryption; Transparent Data Encryption (TDE); Column-Level Encryption; Oracle Key Vault; Database Auditing; Oracle Audit Vault; Oracle Database Firewall; Oracle Database Vault; Label Security; Real Application Security; Virtual Private Database; and Data Masking and Subsetting. More detailed Oracle security information can be found here: [Oracle Database Security Guide](#).

A crucial underreported aspect of database encryption is what happens when it's implemented. All database cybersecurity measures require processing, memory, and I/O resources to be effective. Encryption/decryption is one of the most intense and can reduce database performance by as much as 20%. But once again Oracle Exadata is unique. It takes advantage of the in-built computational storage to offload those resource intensive encrypt/decrypt functions to ensure there is a nominal impact to the Oracle Database performance at approximately 2% or less². This is something no DIY database system can do.

All-inclusive Oracle Database 19c Running on Exadata X8

The Oracle Database is the world's 1st and currently only multi-purpose all-inclusive database. What makes the Oracle Database multi-purpose and all-inclusive can be found in the extensively detailed Chris Craft blog: [What Makes Oracle Database Number One](#). The most important aspect of the All-inclusive Oracle Database is that it's the only database that covers the full range of data types and analytics with unparalleled native data protection and resilience, period.

Although the Oracle Database 19c can be implemented on standard servers from HPE, DELL, Lenovo, IBM, SuperMicro, Samina, Quanta, NEC, Fujitsu, and others, customers will get net incremental capabilities when running 19c on Exadata. Below are just some of the specific Exadata features that 19c users will benefit from:

² Some vendors assert self-encrypting drives (SED) eliminates encrypt/decrypt overhead more than Oracle Exadata TDE offload. For data encrypt/decrypt at rest, it is equivalent. However, SEDs does not encrypt data in-flight or backed up. That still requires TDE and its overhead.

Automatic Indexing	Database Aware Flash Cache	Pure Columnar Flash Cache
Smart Scan	Smart Flash Logging	Hybrid Columnar Compression (HCC)
Automatic In-Memory	Smart Write Back	Backup Change Tracking Filtering in Storage
In-Memory Fault Tolerance	Reverse Offload	Cross Node Parallel Query
In-Memory on Active Data Guard	Fast Node Death Detection	ExaFusion Direct-to-Write Interconnect
In-Memory External Tables	Fast Network Fail Detection	Faster Smart Scans Using Column-level Checksum
In-Memory Row Store	Reduced Brownouts for Instance Failures	Faster Row Store Ingest
Bloom Filtering	IO Resource Manager	Smart Fusion Block Transfer
Cloud-Scale	Automatic Monitoring	

Dispelling DIY Database System Mythology

Far too many IT professionals suffer from DIY database system illusions. They believe they can match both the performance, scalability, and functionality of a purpose-built engineered system. They also believe that utilizing commodity-off-the-shelf servers (COTS) and storage will cost much less than an Oracle Exadata X8 purpose-built engineered system. The trade press continually reinforces these illusions. And yet, both illusions are completely false, incorrect, and just plain wrong.

This document has just shown categorically there is no performance, scalability, or functionality equivalence between what the Oracle X8 brings to the table versus any DIY database system. The DIY database system simply cannot match the Oracle Exadata X8. But what about cost?

The cost illusion is so deeply embedded that many IT pros will not believe the empirical observable facts when presented. It's called adaptive reasoning. The facts are the facts. Here are real world actual DIY comparisons that led to very large Oracle Exadata sales.

Multi-national Multi-\$ Billion USD Retailer

The choice was between 5 different vendors, 144 COTS servers, 10 storage arrays, 20 SAN switches, and 28 racks running the Oracle Database and Oracle with the Oracle Exadata in 4 racks.

The retailer chose Oracle seeing an immediate capital expenditure (CapEx) savings of \$3.5 Million USD. They saw their license expenditures decline by 1,000 cores because Oracle does not charge Oracle Database licensing for computational storage cores in the Oracle Exadata storage servers. Annual power costs declined by nearly 60% from \$105,000 USD to \$42,000 USD. They saw a 99% reduction in patch operations alone going from 1,392 patches per year to just 16 Oracle managed patches.



All this while performance increased between 2-5x.

Multi-national Multi-\$ Billion USD Financial Services Firm



The choice was between DIY private cloud using Oracle Database, RHEL, VMware, and COTS server hardware and Oracle Engineered Systems (ODA, Exadata, ZFS Storage Appliance, and ZDLRA).

The bank chose Oracle because it saved them \$100 M USD over 7 years and was 1/3rd the cost of their alternative. Oracle Engineered systems were much faster to deploy and get into production, were significantly easier to manage and operate, handled multi-PB databases with ease, and ran noticeably faster than expectations.

Multi-\$ Billion USD Healthcare Organization

Replaced 600 IBM POWER Series systems and Terradata Database with 19 full Exadata racks and a clear roadmap to public cloud. The Healthcare provider sees a 10x performance improvement and \$ Millions USD in operational savings from license cost maintenance, rack space savings, power and cooling savings, management savings, personnel savings, and much more.



Conclusion

Oracle Exadata X8 performance, operational simplicity, manageability, time-to-market, total costs, and price/performance is difficult to nigh impossible to match with any DIY database systems or any other database system...period. This is why the Oracle X8 spells doom for DIY database systems.

For More Information on Oracle Exadata

Go to: [Oracle Exadata](#)

Paper sponsored by Oracle. **About Dragon Slayer Consulting:** Marc Staimer, as President and CDS of the 21-year-old Dragon Slayer Consulting in Beaverton, OR, is well known for his in-depth and keen understanding of user problems, especially with storage, networking, applications, cloud services, data protection, and virtualization. Marc has published thousands of technology articles and tips from the user perspective for internationally renowned online trades including many of TechTarget's Searchxxx.com websites and Network Computing and GigaOM. Marc has additionally delivered hundreds of white papers, webinars, and seminars to many well-known industry giants such as: Brocade, Cisco, DELL, EMC, Emulex (Avago), HDS, HPE, LSI (Avago), Mellanox, NEC, NetApp, Oracle, QLogic, SanDisk, and Western Digital. He has additionally provided similar services to smaller, less well-known vendors/startups including: Asigra, Cloudtenna, Clustrix, Conduvix, DH2i, Diablo, FalconStor, Gridstore, ioFABRIC, Nexenta, Neuxpower, NetEx, NoviFlow, Pavilion Data, Permabit, Qumulo, SBDS, StorONE, Tegile, and many more. His speaking engagements are always well attended, often standing room only, because of the pragmatic, immediately useful information provided. Marc can be reached at marcstaimer@me.com, (503)-312-2167, in Beaverton OR, 97007.