WHITE PAPER

Re-defining the role of the DBA

Moving to an open architecture approach

Published September 2013 © Ovum

EXECUTIVE SUMMARY

CATALYST

The role of IT in an organization is undergoing significant change, driven by the need for businesses to become more agile and have greater control over technology cost levers. IT transformation is all about changing to meet the current and future demands of the business. Central to this modernization is adopting new approaches to the roles and responsibilities of IT employees. A key player is the database administrator (DBA), who needs to understand how this critical role will change as organizations look to implement new flexible architectures to align with business demands.

SUMMARY

Most IT organizations run a mix of databases and workloads. Selecting the correct infrastructure for a database environment requires the organization to take a wider perspective on the issues and benefits, looking beyond the performance claims of a single-vendor solution. An open architecture approach should be leveraged to provide configuration flexibility from a storage, compute, and network perspective, while also delivering optimization at the database level. The argument for a single-vendor engineered solution is often based on performance and greater control for DBAs. However, with the right open infrastructure, IT organizations can both optimize databases and DBA efficiency and deliver these benefits across multi-database versions and workloads for greater flexibility and ROI.

There are several key areas that organizations should take into account when choosing an infrastructure for their databases, including availability, total cost of ownership (TCO), and open support of virtualization technologies. Higher availability reduces the impact of taking applications offline to add storage and perform updates and maintenance; any downtime costs money as these systems are typically mission-critical. Organizations can improve TCO by reducing licensing and maintenance costs, which can be achieved without impacting performance by consolidating databases to fewer servers with fewer CPUs. Support for different vendor virtualization technologies means organizations can select the technology they are most comfortable with. Many engineered systems fall short in these categories, exposing the business to greater complexity and management risk and a higher total cost of ownership when viewed holistically over a multi-year period.

KEY MESSAGES

- Moving beyond performance considerations to the holistic value of open architecture across all of IT
- Availability and the cost of application unavailability are important considerations for the business.
- Identifying how the DBA's role is changing and redefining the DBA's function with help from EMC
- Allowing open choice of database virtualization approaches.

© Ovum. This White Paper is a licensed product and is not to be photocopied

MOVING BEYOND PERFORMANCE CONSIDERATIONS TO THE VALUE OF OPEN ARCHITECTURE

By focusing on delivering database performance, some organizations have adopted a vendor-centric infrastructure approach. The major database vendors, with one exception, are also infrastructure hardware vendors, and customer choice is limited. Both the major hardware-based database vendors have engineered systems specifically designed for database workloads, but these solutions are optimized for their particular database offering. Thus customer choice is limited because such systems are tied to a single vendor's database product and do not allow the system to be used for other types of databases and/or applications. This limitation is particularly noticeable when organizations are using a different database for their online transaction processing (OLTP) applications, compared to their data warehouse/business intelligence applications. Using vendor-centric converged infrastructure requires at least two different systems. Ovum advocates caution in using such engineered systems as many are optimized for just one type of application, despite claiming the ability to support both. Specifically, IT organizations may be surprised to find that many of the features they would leverage in a data warehouse (DW) environment, for example, may not apply when using OLTP workloads.

To determine if a database and its infrastructure are tuned specifically for performance, key target metrics are needed, such as the database's I/O pattern. However, database I/O activity patterns vary depending upon workload activity. For example, the workload activity for a data warehouse is significantly different to that of an OLTP application. Companies often have a mix of both environments, and organizations struggle with the question of whether to use a separate, specially tuned, vendor-specific solution or an open architecture approach.

The engineered systems approach focuses on database performance, but this focus ignores the wider aspects of executing mission-critical applications. Although the database is probably the most critical part of these applications, the emphasis should not just be on performance. Organizations should also consider the total cost of ownership (TCO), availability, flexibility, and virtualization support.

For most CIOs, TCO is the benchmark by which they make decisions, and an open architecture approach provides a more holistic perspective on TCO when compared to a dedicated engineered systems approach. Typically, databases will be owned for many years, so the TCO should be calculated over more than the typical one- to three-year horizon that some vendors propose. Ovum advocates the use of a three- to-five year view when calculating TCO, because years four and five make a significant difference. When considering the TCO of database applications, particularly Oracle, the majority of the cost is due to the licensing. For Oracle, licensing is core-based, so consolidating Oracle databases on to fewer CPUs can represent significant savings. Many systems are not CPU bound, with CPU utilization remaining low at circa 20%. Therefore, by putting the active data set on fast flash storage, performance can be improved with fewer CPUs. This improvement is possible because the active data set is constantly performing I/O operations reading and writing the data to disk. Placing the active data set on fast flash storage can significantly improve I/O latency, which is not a function of the CPU. Consequently, additional CPU capacity has no effect.

Availability

Although TCO is the benchmark for the CIO, availability is the benchmark for business users – measured in terms of how much money is lost when systems are unavailable. The claim that engineered systems offer "five nines" (99.999%) availability is touted by many vendors at the marketing brochure level. However, once implemented in real-world application environments, specific IT management tasks can impact these marketing claims and have a dramatic impact on the IT organization.

© Ovum. This White Paper is a licensed product and is not to be photocopied

For example, applying quarterly patch updates can require multiple copies of a database environment to be created, and can involve system downtime. Other examples include downtime associated with swapping failed hardware components such as disk drives, and the ability to scale capacity on demand with databases online. It is not uncommon to see this reduce the availability of many engineered systems to "three nines" (99.9%). While this may seem academic, it translates to eight hours downtime in a year for "three nines", compared to five minutes a year for "five nines". By comparison, an open architecture does not bind a specific set of hardware resources to a specific database. As such, it does not suffer the disruption of a quarterly update. It also offers other non-disruptive benefits, such as the ability to swap failing disk drives and the ability to add additional storage for an application while a database remains online. Both of these improve availability, which translates to organizations losing less money through system outages. Furthermore, technologies that deliver application and data availability within the data center and over distance can be utilized to provide full infrastructure utilization and zero downtime.

TCO

An often-overlooked aspect of an engineered system is its inherent lack of flexibility. As previously mentioned, CPU utilization of systems is low, and consolidation therefore provides an opportunity to reduce costs. However, an engineered system must be able to support multiple varied workloads for consolidation to be an option. An open architecture approach provides organizations with the flexibility to use systems for a heterogeneous mix of workloads, and thus supports consolidation. This approach also enables organizations to provide granular management down to the application level, providing IT departments with the workload management flexibility to define differential SLA-based business needs.

Virtualization

A system’s ability to support virtualization is also a consideration. Virtualization helps organizations to reduce TCO while still delivering the required level of service and performance. Putting a database in a virtual machine can practically guarantee that it will receive the resources it needs, which in turn can deliver an expected performance level. Support for virtualization brings a number of other benefits, such as the ability to perform active workload balancing while providing high availability. An open architecture approach provides this capability across different virtualization technologies, such as VMware ESXi, Microsoft Hyper-V, and Oracle VM.

© Ovum. This White Paper is a licensed product and is not to be photocopied

Figure 1: Objectives of enterprise IT investment strategies

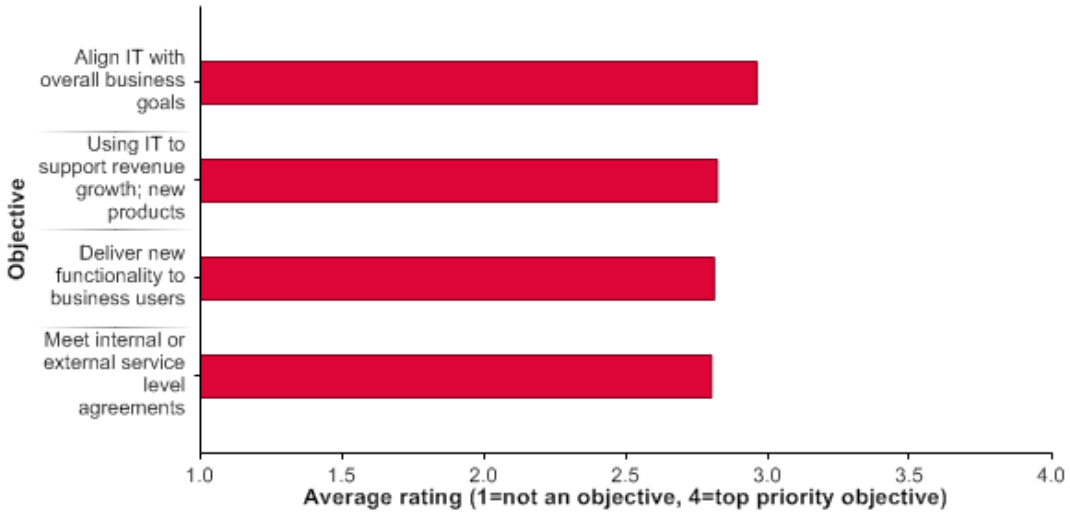

Ovum Technology Trends Survey 2012 n=17,422

Source: Ovum

A key requirement for any organization in today's economy is to deliver services when they are needed, and to reduce the operational spend on maintaining service delivery. The results of Ovum's 2012 Technology Trends research (see Figure 1) indicated that the primary objective for IT investment is to align IT with overall business goals. The value of an open architecture approach is its flexibility and the removal of the vendor lock-in aspect of any specific technology, enabling organizations to deploy the most appropriate database topology for a particular use case, which could change over time.

AVAILABILITY AND THE COST OF APPLICATION UNAVAILABILITY ARE IMPORTANT CONSIDERATIONS FOR THE BUSINESS

Most organizations that run Oracle databases and applications do so for mission-critical workloads. These are typically core OLTP applications or the business intelligence analytics systems that are used to run an organization's operations. The challenge for the IT department in these scenarios is to ensure that availability is not compromised by a software, hardware, or management problem.

Ovum believes that these three categories require different approaches to ensuring availability: software problems require the ability to quickly restore at a granular-level an application or data set to a known good point; hardware problems require the system to automatically detect issues and proactively move workloads off failing components; and management problems require the system to provision storage without taking applications offline, as well as the ability to patch systems while they are running and accessible.

IT organizations need to ensure their infrastructure is equipped to streamline protection and recovery from hardware, software, and management errors. Engineered systems rely almost exclusively on database software to provide all three levels of protection. This can place excess burden on the database and the DBA that manages it, and fails to take advantage of the significant advancements in storage infrastructure.

© Ovum. This White Paper is a licensed product and is not to be photocopied

An open architecture approach can combine the best of database availability features and augment them through added infrastructure capabilities to improve availability and data protection across all the databases it supports.

Data protection has never been more difficult and more critical to business survival. Organizations need to protect increasing amounts of data, but can tolerate increasingly less disruption to applications during the backup process. In the event that data is lost through software, hardware, or a management error, organizations must be able to recover it in ever-shorter periods of time. A key capability in this area is any-point-in-time recovery.

An open architecture approach can use a number of different techniques to significantly improve availability from a data protection perspective. Logical replication can be used to ensure that data is backed up and available when needed. Snapshots are point-in-time copies of changed bits or volumes on a storage block. They are quick to make and, because they are not full copies, are space-efficient. Snapshots can be read-only or read-write copies of files, LUNs, file systems, or any other type of container. Another approach is clones; these look a lot like snapshots, but are bit-identical copies of full volumes. They require the same storage capacity as the primary volume.

Another technique used by some leading storage vendors, such as EMC, is active-active database replication. Active-active database replication eliminates downtime not only for unplanned interruptions such as disasters or site failures, but also for planned outages such as migrations, upgrades, and systems maintenance. These solutions allow organizations to have different hardware setups, operating systems, databases and database versions, service packs, and so on. In many migration and consolidation projects, a bi-directional realtime replication capability is needed to perform active-active replication between old and new environments in order to completely eliminate downtime and minimize risks with phased migration.

When considering the level of availability and the recovery time objective (RTO), an open architecture approach can provide significantly improved results compared to a tightly coupled vendor-specific solution. For example, a large US bank reported that it was able to reduce recovery point times from 45 minutes to less than 5 minutes by using network-based open replication technology instead of a proprietary-based solution. Evidence such as this indicates that an open architecture approach can match or improve the RTO reported by engineered system vendors.

IDENTIFYING HOW THE DBA'S ROLE IS CHANGING

DBAs have an important and unique role within an enterprise. Traditionally, the DBA has been seen as someone focused on the maintenance of databases. The DBA understands the data better than most in the organization and appreciates that its value is critical to business success. The DBA maintains close links with the business and IT, and is able to demonstrate to business users that IT is transforming from a technical function to a more business-focused service-oriented operation.

One of the obvious changes that DBAs will see is that their role is moving from an operational process-centric administrative role to more of a business usage consultant. In this new role, there will be a shift from performing standard processes to becoming more of a data expert that can help end users understand how to access and utilize an organization’s data.

To enable this transformation, DBAs should seek to have infrastructure deployed that adds greater automation and simplicity to their routine daily IT management tasks.

© Ovum. This White Paper is a licensed product and is not to be photocopied

DBAs currently spend a significant amount of time reacting to problems and performing storage and capacity planning activities. Much of this activity is performed using an individual DBA's library of scripts that they have created over time. This reliance on manual scripts works well for single vendor database technologies and is one of the biggest changes many DBAs will experience as they move to a more open architecture approach. The benefit for the DBA will be in using open tools to create workflows that automate activities across different technologies, enabling improvements in productivity.

Performance

One of the most common, complex, and time-consuming tasks that a DBA performs is managing the performance and capacity planning of the database. Auto-tiering within storage, for example, uses automated techniques to ensure that "hot" data is placed on the fastest storage available, while "cold" data is moved to slower and cheaper storage devices. The advantages of this are amplified when the DBA is able to dynamically provision storage to run applications. These two capabilities enable the DBA to automate data storage placement and provision storage to an application using virtual provisioning. This technique allows an application to think it has more storage than is physically allocated, which allows the DBA to add more storage while an application is running. The business advantages are that these technologies will free up much of the DBA's time, allowing them to focus on other business-related activities, and reduce the downtime of applications from a management perspective.

Backup and Recovery

Backup and recovery remains important, even if it is the least visible job of the DBA. Such activities involve meticulous attention to detail when it comes to checking transaction log files and event logs. Both sources of information are critical to ensuring that a backup image is consistent and complete; if an inconsistent or incomplete backup image is used for database recovery, it will corrupt the database. Many database vendors have developed tools that eliminate this risk, but an open architecture approach provides the DBA with a greater set of options. The DBA's attention-to-detail skills are not lost; they can be re-directed to other activities that automation cannot perform, such as database strategy and architecture.

Test and Development

An open architecture environment can also help to streamline test and development operations.

First, an open architecture approach simplifies management by allowing production, test, and development to co-exist within the same infrastructure environment, and enabling data to move quickly and easily throughout the lifecycle using common infrastructure and best practices. Secondly, it improves the quality of application development, as test/development environments can be hosted on the same hardware as production environments. This eliminates the differences that could exist between two different systems, such as microcode, patch-level, and device driver. Finally, it enables development and test databases to be provisioned faster, accelerating application release cycles. With an open architecture approach, database provisioning can be performed rapidly and simply for a wide range of different database types. An added benefit is that this can be performed across multiple different environments, which is a distinct advantage over a single-vendor solution that is restricted to one database on one environment.

DBAs need to embrace automation technology to release them from manual processes. As the DBA moves into a new design-based role, they will need to consider adjacent activities, such as ensuring the capacity planning process is integrated into lifecycle management and adding compliance and auditing to governance.

© Ovum. This White Paper is a licensed product and is not to be photocopied

REDEFINING THE DBA'S FUNCTION WITH HELP FROM EMC

EMC and Oracle have a long history of working together on joint best practices and product integration. These engagements have all leveraged an open architecture approach that supports multiple Oracle database workloads and environments. Ovum has identified six key aspects of EMC's open architecture approach that represent changes to how DBAs work:

1. **Database management optimization** – EMC has developed a solution that provides the DBA with access to information that is normally obscured or that requires separate monitoring tools. The EMC VMAX Storage Plug-in for Oracle Enterprise Manager 12c enables the DBA to monitor performance through deep visibility into the storage system. This proactive management approach will allow the DBA to manage the database more effectively.

2. **Performance and capacity planning** – Automation capabilities, such as auto-tiering, manage the storage placement of data to optimize performance, freeing up the DBA to focus on more value-added business activities. Virtual provisioning capabilities enable storage to initially be over-provisioned to an application, so that it thinks it has more physical storage than it actually has. This allows the DBA to dynamically add storage to a running application, thereby minimizing downtime and ensuring database availability.

3. **Backup and recovery** are the most important tasks that a DBA performs, as they ensure data integrity. EMC, through its Oracle Recover Manager (RMAN) integration and purpose-built backup appliances, can quicken the process of backing up Oracle databases, thereby saving the DBA time. EMC claims these appliances increase the speed of backups by 50% and require 40% less effort from IT teams. Ovum believes these savings are a result of EMC's ability to provide the control needed to back up DBA operations; the choice is up to the customer and the advantage is that a single team is responsible for the entire backup operation.

4. **Data protection** represents another area where DBA's normally take a lead because of the specific set of skills needed. EMC RecoverPoint can recover an Oracle database and any applications it supports from a specific point in time, with more recovery granularity than that of competitive solutions. While Ovum has not tested this claim, the fact that EMC has database-specific solutions means it can justifiably claim to be more than a generalist in the backup and recovery space.

5. **EMC VMAX technology** removes the need for systems to be taken offline to perform maintenance updates. This capability ensures that "five nines" (99.999%) availability is truly obtainable, particularly when compared to engineered systems that require quarterly maintenance updates where systems must be taken offline to perform this work.

6. **Flexibility and the ability to support multi-hypervisor technologies** enable a DBA to migrate virtual machines over distance using EMC VPLEX for Oracle RAC (Real Application Cluster) environments. Ovum considers this to be an advantage for a DBA because it means that this technology can become more pervasive in the enterprise. It will also enable DBAs to work across locations on the same data at the same time, and make things like migrations or tech refreshes faster and simpler to perform. Additionally, VPLEX can provide continuous availability and make disaster recovery (DR) testing far simpler, as a single consistent read/writable database exists in both locations simultaneously.

All six of these capabilities are compatible with Oracle, making the DBA's life easier. EMC has developed these from its experience running one of the largest Oracle environments in the world with over 50 database

© Ovum. This White Paper is a licensed product and is not to be photocopied

grids. However, the most influential aspect is the online Community at EMC that brings DBAs and EMC infrastructure experts together in a free online portal.

ALLOWING OPEN CHOICE OF DATABASE VIRTUALIZATION APPROACHES

With regard to virtualization, the industry has seen most aspects of the infrastructure move toward a virtual mode of operating. However, the database layer has largely remained in the physical world. Fortunately, solutions have been developed that clone production databases and share them among many different users in a way that does not require a significant investment in additional storage. This has changed many DBAs' perceptions. The original concept behind database virtualization was to address three key areas that other virtualization technologies have missed.

The first area was to improve the ROI of database ownership. Traditionally, databases were implemented on a one-to-one basis with an associated physical server and storage array. Database virtualization broke this relationship and enabled IT departments to obtain cost savings through consolidation. However, the consolidation was based on the number of physical servers rather than the storage arrays. This was because the direct one-to-one relationship between database and server meant that these servers were typically running at 20% utilization rates. The associated benefits from consolidation were intended to improve operational efficiencies and provide business users with a more responsive and flexible service. The bottom-line for IT departments is that database virtualizations provide both opex and capex savings. Therefore, CIOs can report improved use of IT investments in hardware and reduced operational spending.

The second was to improve developer throughput and remove the testing bottleneck that most organizations experience. The concept of agile development is to develop and release incremental changes as fast as practical. However, for most development environments, the many different forms of testing required means that multiple copies of test data are needed, which take time to refresh between testing cycles. The virtualization of the database enables these virtual database instances to be made available within minutes, thereby speeding up the entire development lifecycle. It also enables an organization to make the process of requesting and deploying a test or development database a user self-service-driven process. This aspect provides increased flexibility and allows more user choice and control.

The third was to improve the quality of all software product releases. Traditionally, testing against a full production database was costly and time consuming. Database virtualization allows developers and testers to use identical copies of the production database, which can help to identify problems that testing on reduced sub-sets of data fail to expose. A database virtualization approach enables live production data to be masked and re-ordered to protect customer confidentiality once, and then multiple versions can be created and used in parallel. Data in test databases can also be kept up to date with live production data to ensure that testing is performed against actual "live" scenarios. Another benefit for testers is that database virtualization allows users to create virtual database instances so that comparisons can be made between data before and after tests are run. Ovum considers this an excellent aid for developers when troubleshooting problems.

© Ovum. This White Paper is a licensed product and is not to be photocopied

RECOMMENDATIONS

The decision about which infrastructure is best suited for database environments is not simple. Some vendors will push the performance aspect as being the most important, but Ovum recommends that organizations consider other aspects such as availability, TCO, flexibility, and support for virtualization, as well as performance. When all these aspects are considered, the open architecture approach provides a compelling alternative to some single-vendor engineered solutions.

COMPANY PROFILE

Table 1: Contact details

EMC Corporation
World Headquarters
176 South Street
Hopkinton, MA 01748
USA

www.emc.com

Source: EMC OVUM

Ovum's Knowledge Centers are new premium services offering the entire suite of Ovum information in fully interactive formats.
To find out more about Knowledge Centers and our research, contact us:

Ovum Europe
Mortimer House
37-41 Mortimer Street
London W1T 3JH, UK
+44 20 7551 9000
crmgroup@ovum.com

Ovum Australia
Level 5, 459 Little Collins Street
Melbourne, VIC 3000
Australia
+61 3 9601 6700

Ovum New York
100 Wall Street, 9th Floor
New York
NY 10005, USA
+1 212 686 7400

All Rights Reserved

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of the publisher, Ovum. Whilst every care is taken to ensure the accuracy of the information contained in this material, the facts, estimates, and opinions stated are based on information and sources which, while we believe them to be reliable, are not guaranteed. In particular, it should not be relied upon as the sole source of reference in relation to the subject matter. No liability can be accepted by Ovum, its directors, or employees for any loss occasioned to any person or entity acting or failing to act as a result of anything contained in or omitted from the content of this material, or our conclusions as stated. The findings are Ovum's current opinions; they are subject to change without notice. Ovum has no obligation to update or amend the research or to let anyone know if our opinions change materially.

© Ovum. Unauthorised reproduction prohibited

This report is a licensed product and is not to be reproduced without prior permission.