IBM Power Systems: a Better Choice than Oracle Hardware for Mission-Critical Oracle Databases

IBM Competitive Project Office
Rebecca Ballough
Executive IT Specialist

Contents

Figures and Tables 3
Executive Summary 4
Reliability 5
IBM Power Systems Offer Minimal Unplanned Downtime 5
IBM Power Systems Offer Minimal Planned Downtime 5
Does Oracle Real Application Clusters Fix Oracle Server Availability Issues? 6
IBM Power Systems Offer Reliability by Design 8
Customer Satisfaction 9
Satisfaction with Products, Service and Support 9
Provisioning Time 9
Implications for Cloud Deployments 10
Conclusion 11
References 12

Figures and Tables

Figure 1: 2015 ITIC Survey - Unplanned Downtime Hours per Year 5
Figure 2: 2015 ITIC Survey - Percentage of unplanned downtime lasting longer than 4 hours 5
Figure 3: 2015 ITIC Survey - Planned Downtime Hours per Month 6
Figure 4: 2015 ITIC Survey - Satisfaction with server hardware vendor products, service and support 9
Figure 5: 2015 ITIC Survey – Satisfaction with time to Configure & Provision Servers 9
Figure 6:2015 ITIC Survey – Which vendors provide the most comprehensive security for the cloud 10

Executive Summary

Oracle offers a line of server solutions called Engineered Systems for customers running its database software. These products include the Oracle SuperCluster, which hosts databases and applications running on Solaris on Oracle Sun T5-8, M6-32, or the newest M7-8 hardware. The Engineered Systems line also includes database-only x86-based offerings, namely, the Database Appliance and the Oracle Exadata Database Machine, both of which use Oracle Enterprise Linux, similar to Red Hat Linux, as the operating system.

Oracle indicates (and its customers believe) that running Oracle software on Oracle hardware provides better availability, better support, and requires less installation time.

This belief has been discredited by a recent survey conducted by Information Technology Intelligence Consulting (ITIC). For more than seven years, ITIC has conducted an annual hardware and operating system reliability study, polling more than six hundred businesses. This survey is independent; no vendors sponsor this study, nor are participants compensated for providing data.

According to this ITIC survey, IBM® Power Systems® have less planned and unplanned downtime than Oracle x86 and SPARC servers, and have higher customer satisfaction with products, service, support, and provisioning time. As this same study notes, minimizing downtime has a measurable impact on business results: *"**A 98% majority of respondents say that a single hour of downtime per year costs their company over $100,000".***[1] An older study that breaks down losses by industry notes that an hour of downtime in banking may look more like $2.1M in losses per hour.[2]

Even with the recent announcement of new hardware products, including the new SPARC M7 and T7 lines, Oracle server availability is not expected to improve. Oracle does not highlight any additional reliability features in its new products; Oracle continues to provide what it calls fault tolerance only by configuring two or more servers.

Fixing downtime and customer satisfaction issues with a cloud strategy also doesn't look very promising for Oracle. Security is a key component of cloud-based solutions, and customers surveyed by ITIC viewed IBM as providing the most comprehensive security strategy and products for the cloud, with Oracle lagging 2.5 times behind.

[1] ITIC 2015 Global Server Hardware, Server OS Reliability Survey http://www.lenovo.com/images/products/system-x/pdfs/white-papers/itic_2015_reliability_wp.pdf

[2] Andrew Hiles, Five Nines: Chasing the Dream? http://www.continuitycentral.com/feature0267.htm

Reliability

IBM Power Systems Offer Minimal Unplanned Downtime

According to the ITIC 2015 Survey, enterprises define mission-critical reliability as having less than 12 minutes of unplanned downtime per year. IBM Power Systems running AIX® or Linux meet this target, with only 10.2 minutes unplanned downtime reported. Oracle x86 and SPARC based systems have more than 18 minutes unplanned downtime reported, which does not meet the standard for mission-critical reliability, and is nearly twice as much downtime as Power Systems. These results are based on comparable server configurations, with similar workloads and server age.

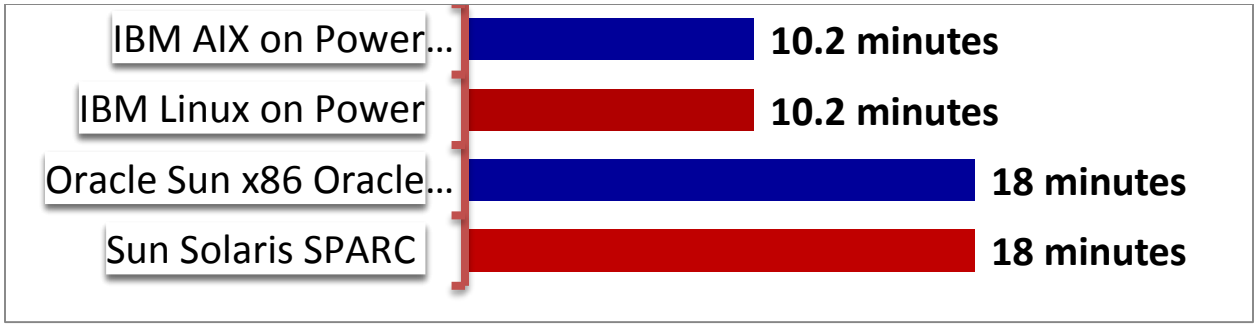

Figure 1: 2015 ITIC Survey - Unplanned Downtime Hours per Year

Even worse for Oracle customers, when the systems go down, they stay down much longer. The following survey results show that while very few outages on IBM Power Systems last longer than 4 hours, 10% of Oracle's x86 outages are at least that long. SPARC systems fare a little better, with only 4% of downtime incidences lasting 4 hours or longer.

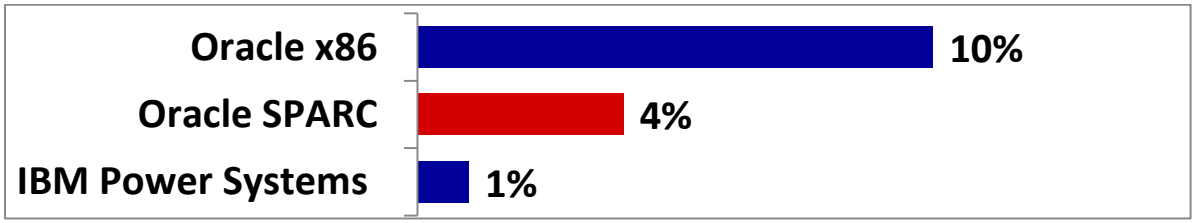

Figure 2: 2015 ITIC Survey - Percentage of unplanned downtime lasting longer than 4 hours

IBM Power Systems Offer Minimal Planned Downtime

Minimizing planned downtime is also a key element of system availability, particularly for worldwide organizations or those providing goods or services over the internet. Here, too, IBM AIX needs the least amount of time for applying patches, upgrades or other maintenance activities. In fact AIX needs less

than 10 times the amount of planned maintenance time as Solaris; Solaris 11 needs over 84 minutes per month of planned downtime, while AIX needs less than 8 minutes.

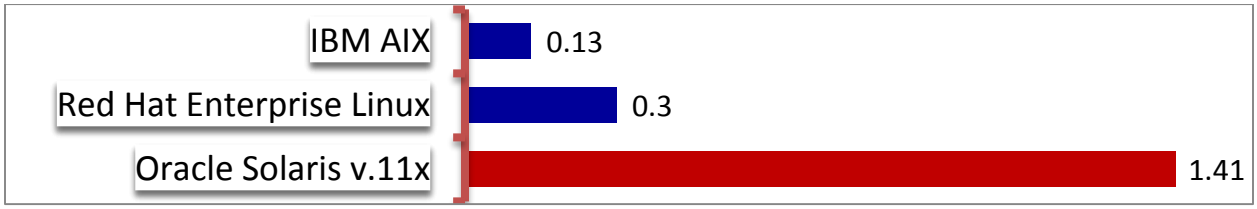

Figure 3: 2015 ITIC Survey - Planned Downtime Hours per Month

Does Oracle Real Application Clusters Fix Oracle Server Availability Issues?

Oracle's methodology for providing availability is to use two or more servers, rather than increase the reliability of those servers. Oracle Engineered Systems including SuperCluster, Exadata, and the Oracle Database Appliance are designed as two or more node clusters, using Oracle Grid Infrastructure as a clustering layer. Oracle Real Application Clusters (RAC) is a database feature (available for an extra license cost) to allow databases to leverage a multiple node Grid Infrastructure framework. Oracle RAC is an active-active design, where the database runs on multiple server nodes simultaneously.[3]

To see the difference between availability provided by Oracle RAC compared to availability provided by the server staying up and running, let's look at how an application behaves when a cluster node crashes.

When a node crashes in a RAC cluster, users who are connected to the surviving node or nodes remain connected and able to execute transactions. The degree of disruption the users on the failed node will experience varies based on Oracle features the application can utilize.

Transparent Application Failover (TAF) is an Oracle capability which allows sessions to automatically reconnect to another database instance when those sessions time out from their original host (such as when a node fails). Queries may continue on a surviving cluster node, while uncommitted transactions will be rolled back. TAF requires that sessions connect using a specific Oracle Call Interface (OCI) driver. OCI requires that the application use very specific drivers such as JDBC-OCI, PHP OCI8, or Ruby OCI8.[4] Notably, JDBC thin drivers cannot use this driver, so any application which connects using JDBC thin cannot even use this technology. This includes Oracle's own E-Business Suite product.

[3] A variation on this technology, RAC One-Node, is an active-passive design, where the database instance runs on only one node, although a second node may be brought online temporarily to assist with rolling upgrades. RAC One-Node is a licensed database feature. For pricing information, consult the Oracle Technology Global Price List, http://www.oracle.com/us/corporate/pricing/technology-price-list-070617.pdf

[4] Application Failover with Oracle Database 11g, http://www.oracle.com/technetwork/database/app-failover-oracle-database-11g-173323.pdf

If a session can't use TAF, that user or application will have to reconnect to the database some other way; users will have to log in again and manually resume activity. Stateless middleware may even need to be rebooted to reconnect to another database node.

Even when applications can use TAF, sessions are only failed over after a TCP Timeout is detected on the original host; this parameter is typically set to 5 minutes, so an individual session will hang for 5 minutes before it can be reconnected to another cluster node.

As previously mentioned, although queries may continue when TAF fails over a session[5], any uncommitted transactions will roll back. This can cause two big issues. First, long-running batch jobs may have many uncommitted transactions. Rolling back these transactions may take a long time and increase failover time overall. In addition, it may not be trivial to restart a batch job. Second, when a failover occurs in transactional environments there may be some confusion on the part of a user as to whether a transaction actually committed.[6] This can lead to unacceptable problems like multiple bank account transfers or purchase of duplicate items.

Due to the limitations of TAF, Oracle has made some advancements to this technology. First is Fast Application Notification (FAN), which proactively notifies the application when a database condition changes (such as an instance going down). This avoids the 5 minute TCP timeout delay. It can work in conjunction with Fast Connect Failover (FCF), which can cleanup and reconnect dead connections. However, in order to use FAN with FCF, the application must use Oracle integrated clients.[7]

Although the improvements with FAN and FCF make the reconnection time faster, they don't eliminate the disruption due to rolled back uncommitted transactions or uncertainty on the part of the user as to whether a recent transaction committed.

There are two new technologies in Oracle 12c to attempt to address the issue of questionable commits. However, these technologies are very restrictive, requiring either Oracle Weblogic Server or application code changes, so very few implementations are able to use them.[8]

---

[5] Assuming TAF is configured with FAILOVER_MODE=select as documented in Oracle 12.1 Database Net Services Administrator's Guide. This also increases overhead on the client server during normal select execution.

[6] As stated in the Oracle whitepaper Transaction Guard with Oracle Database 12c, "Transaction guard avoids the costs of clients receiving ambiguous errors …Without Transaction Guard, applications and users who attempt to retry operations following an error or timeout can cause logical corruption by committing duplicate transactions or committing transactions out of order." http://www.oracle.com/technetwork/database/database-cloud/private/transaction-guard-wp-12c-1966209.pdf

[7] FAN may be used on its own, but only if the application is specifically coded to handle the FAN events. Alternatively, FAN can call out scripts on the database tier so the server may take action when the database condition changes, but it may not be clear what action the server should take.

[8] Application Continuity for Java is an Oracle 12c feature which only works with Oracle Weblogic Server. Transaction Guard is a 12c feature which returns a message to the application to confirm whether or not the last transaction committed if a database server failed. Specific application changes are necessary to take action on the notifications provided by Transaction Guard.

There are no Oracle technologies which address the issue of rollbacks on a long-running batch job.

In short, RAC failover is very disruptive to long-running batch jobs, very disruptive to users or applications which cannot use TAF, and can result in a poor end user experience due to uncertainty with the last commit prior to an outage. The failover capabilities that Oracle provides are simply not the same as preventing unplanned server downtime.

IBM Power Systems Offer Reliability by Design

The fact that IBM Power Systems results show better reliability and less planned maintenance than Oracle hardware is a direct result of the focus IBM places on designing servers with reliability, availability, and serviceability in mind.

Central to Power's availability design is the inclusion of hardware checkers, which are built-in hardware probes that look for errors and record them in a reserved area on chip called fault isolation registers. A service processor then monitors that register data and takes any appropriate action (such as retry or sparing) proactively, avoiding downtime. This same capability can be used to determine what parts may need replacement, so the right part is replaced the first time, rather than needing lengthy diagnostic checks to attempt to identify or recreate a problem. This capability is referred to as First Failure Data Capture (FFDC).

But FFDC isn't just used to identify problems. It is also used during the design stage to simulate errors such that IBM can test for potential failure scenarios and work through fixes before the product is generally available. This greatly improves server availability.

Other key elements of IBM Power Systems availability are the features which protect memory. Power systems include specialized DIMMs with built in error checking and correcting (ECC) technology called Chipkill. Chipkill protects memory from single memory chip failure plus multi-bit errors from any part of a single chip. Dynamic bit-steering allows a spare memory chip to take over for one which exceeds an error threshold.

Some x86 servers offer a similar capability to Chipkill and dynamic bit-steering, called lockstep (also referred to as 'RAS' mode) and sparing, respectively. These features are available on the Xeon E5-2600 v3 processor family (the processor behind Oracle's X5-2 servers used in the Exadata and Database Appliance), but are not enabled by default. Enabling these features reduces available memory and decreases memory performance.

IBM Power Systems also offer Live Partition Mobility (LPM), a feature that allows transparent migration of a partition from one physical server to another. Session states remain active, so no special failover recovery is required. However, if desired, Oracle Grid Infrastructure, Oracle RAC, and RAC One Node are all available for IBM Power Systems running Oracle Databases. Application connectivity technologies like

TAF, FAN, Application Continuity for Java and Transaction Guard are also available on IBM Power Systems in versions and configurations where the Oracle database supports these features.

Customer Satisfaction

Satisfaction with Products, Service and Support

Survey respondents all indicated IBM products, service and support were at least satisfactory, with 73% of users rating their satisfaction at very good or higher.

Oracle hardware customers are much less satisfied, with only 47% of users giving a very good or higher rating. What's worse, 14% of Oracle hardware customers consider their service and support poor or even unsatisfactory.

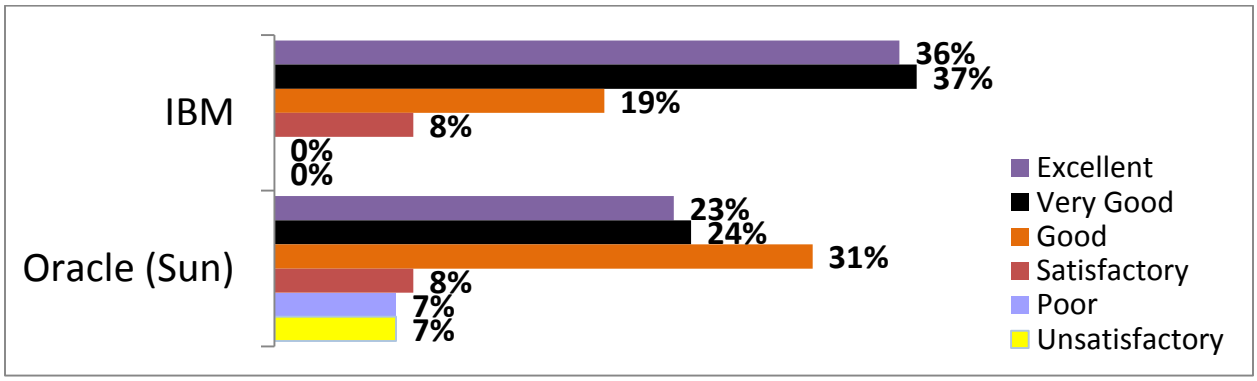

Figure 4: 2015 ITIC Survey - Satisfaction with server hardware vendor products, service and support

Provisioning Time

Fewer customers reported satisfaction with configuration and provisioning time with Oracle servers, with 18% of customers unsatisfied with the amount of time required. By comparison, 99% of IBM customers were satisfied with configuration and provisioning time.

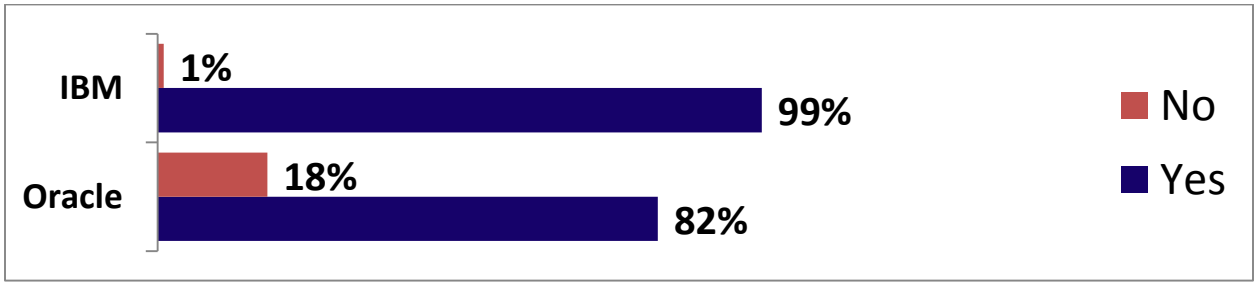

Figure 5: 2015 ITIC Survey – Satisfaction with time to Configure & Provision Servers

Oracle Enterprise Manager, a tool Oracle uses for managing and monitoring Oracle environments, can also be utilized to provision Oracle software on AIX on Power, if customers chose to do so.

Implications for Cloud Deployments

When customers are not satisfied with the availability levels of Oracle hardware on-premises, and are not satisfied with Oracle services and support, giving Oracle more control over their mission-critical systems by deploying them in an Oracle cloud seems like a poor choice.

If availability and support considerations are not a significant deterrent to Oracle cloud deployment, the viewpoint of surveyed users on Oracle cloud security should certainly make Oracle customers look toward other cloud options.

According to this same ITIC study, when asked the question "Which vendors provide the most comprehensive security products, services and security strategy for the cloud? (Select ALL that apply)", survey respondents gave IBM the highest rating, over 2.5 times higher than Oracle's cloud security offerings.

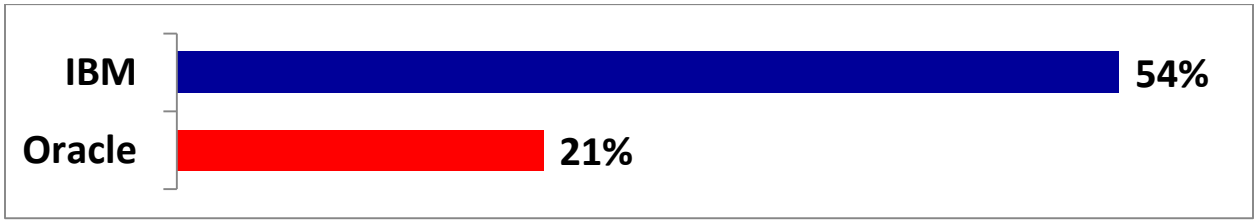

Figure 6:2015 ITIC Survey – Which vendors provide the most comprehensive security for the cloud

Oracle may claim to address security issues with its new M7 and T7 hardware lines, which it says offers "Security in Silicon". [9] Whether or not these claims are true remains to be proven; regardless, SPARC products including these T7 and M7 servers are not even available in Oracle's public cloud as an infrastructure as a service offering.[10]

[9] https://www.oracle.com/servers/sparc/index.html

[10] A trial VM configuration is available, but it for application validation with M7 processors only https://swisdev.oracle.com/

Conclusion

For customers who have requirements for mission critical enterprise databases, ITIC's independent studies show that IBM Power Servers offer the following advantages over Oracle servers:

- **1.7x less unplanned downtime**
- **10x fewer high impact outages (4 hours or more) than Oracle x86 systems, and 4x fewer high impact outages than Oracle SPARC servers**
- **10x less planned maintenance than Oracle Solaris**
- **Higher customer satisfaction with service and support**
    - 73% of Power users report Very Good or higher satisfaction, with no users reporting poor satisfaction levels
    - 47% of Oracle users report Very Good or higher satisfaction, with 14% reporting poor or unsatisfactory service
- **Better options for cloud security**

The ITIC survey results are clear; running Oracle databases on IBM Power Systems is a safer, more reliable choice than running them on Oracle hardware. And Oracle clustering technologies can't compensate for the availability advantage that IBM Power Systems provide.

References

1. ITIC 2015-2016 Global Server Hardware, Server OS Reliability Survey
2. Oracle Maximum Availability Architecture Client Failover Best Practices for Highly Available Oracle Databases
   http://www.oracle.com/technetwork/database/availability/client-failover-2280805.pdf
3. Application Failover with Oracle Database 11g
   http://www.oracle.com/technetwork/database/app-failover-oracle-database-11g-173323.pdf
4. Transaction Guard with Oracle Database 12c
   http://www.oracle.com/technetwork/database/database-cloud/private/transaction-guard-wp-12c-1966209.pdf
5. POWER8® Processor –Based Systems RAS
   http://public.dhe.ibm.com/common/ssi/ecm/po/en/pow03133usen/POW03133USEN.PDF
6. Oracle's SPARC T5-2, SPARC T5-4, SPARC T5-8, and SPARC T5-1B Server Architecture
   http://www.oracle.com/technetwork/server-storage/sun-sparc-enterprise/documentation/o13-024-sparc-t5-architecture-1920540.pdf
7. SPARC M6-32 Server Architecture
   http://www.oracle.com/technetwork/server-storage/sun-sparc-enterprise/documentation/o13-066-sparc-m6-32-architecture-2016053.pdf
8. Oracle Database Net Services Administrator's Guide
   https://docs.oracle.com/database/121/NETAG/advcfg.htm#NETAG013
9. Oracle SPARC T7, M7, and Supercluster M7 documentation
   https://www.oracle.com/systems/secured-computing-powered-by-sparc-m7.html
   https://www.oracle.com/engineered-systems/supercluster/supercluster-m7/features.html

© Copyright IBM Corporation 2015

IBM Corporation
Route 100
Somers, NY 10589
USA

Produced in the United States
August 2015
All Rights Reserved.

IBM, the IBM logo, AIX, and POWER8 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and or other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product or service names might be trademarks or service marks of others.

The information contained in this documentation is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this documentation, it is provided "as is" without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this documentation or any other documentation. Nothing contained in this documentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM (or its suppliers or licensors), or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

No other publication, distribution or use is permitted without the prior written consent of IBM. Customers who want a "deep drill down" on CPO Competitive Case Studies should be directed to the IBM Competitive Project Office under NDA. An NDA is required to explain CPO's methodologies, processes and competitive comparison. You can contact IBM Competitive Project Office by sending an email to ibmcpo@us.ibm.com

References in these materials to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in these materials may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way.

This document was developed with IBM funding. Although the document may utilize publicly available material from various vendors, including IBM, it does not necessarily reflect the positions of such vendors on the issues addressed in this document.