

Introduction Minimizing Planned Downtime of SAP Systems with the Virtualization Technologies & systematic review and research challenges in cloud computing

Mrs.Priyanka Hariom Singh

In-charge Principal Of K.S.E.T'S COLLEGE OF COMPUTER APPLICATIONS

ABSTRACT

Cloud Computing has been used by different types of clients because it has many advantages, including the minimization of infrastructure resources costs, and its elasticity property, which allows services to be scaled up or down according to the current demand. From the Cloud provider point-of-view, there are many challenges to be overcome in order to deliver Cloud services that meet all requirements defined in Service Level Agreements (SLAs). High availability has been one of the biggest challenges for providers, and many services can be used to improve the availability of a service, such as SAP infrastructures on which hundreds of users depend for their work and especially when systems such as materials management, supply chain management (SCM), and customer relationship management (CRM) run concurrently and interact with each other. we can also find infrastructure and middleware solutions. This systematic review has as its main goal to present and discuss high available (HA) solutions and new technologies implementation which are introduced by oracle i.e Minimizing Planned Downtime of SAP Systems with the Virtualization Technologies for Cloud Computing, and to introduce some research challenges in this area.

Keywords *Cloud computing , biggest challenges for providers, SAP infrastructures, Service Level Agreements (SLAs) ,High availability, Systematic review*

I.INTRODUCTION

Cloud Computing emerged as a novel technology at the end of the last decade, and it has been a trending topic ever since. The Cloud can be seen as a conceptual layer on the Internet, which makes all available software and hardware resources transparent, rendering them accessible through a well-defined interface. Concepts like on-demand self-service, broad network access, resource pooling [1] and other trademarks of Cloud Computing services are the key components of its current popularity. Cloud Computing attracts users by minimizing infrastructure investments and resource management costs while presenting a flexible and

elastic service. Managing such infrastructure remains a great challenge, considering clients' requirements for zero outage [2, 3].

Service downtime not only negatively effects in user experience but directly translates into revenue loss. A report [4] from the International Working Group on Cloud Computing Resiliency (IWGCR)¹ gathers information regarding services downtime and associated revenue losses. It points out that Cloud Foundry² downtime results in \$336,000 less revenue per hour. PayPal, the online payment system, experiences in a revenue loss of \$225,000 per hour. To mitigate the outages, Cloud providers have been focusing on ways to enhance their infrastructure and management strategies to achieve high available (HA) services.

According to [5] availability is calculated as the percentage of time an application and its services are available, given a specific time interval. One achieves high availability (HA) when the service in question is unavailable less than 5.25 minutes per year, meaning at least 99.999 % availability ("five nines"). In [5], authors define that HA systems are fault tolerant systems with no single point of failure; in other words, when a system component fails, it does not necessarily cause the termination of the service provided by that component.

Delivering a higher level of availability has been one of the biggest challenges for Cloud providers. The primary goal of this work is to present a systematic review and discuss the state-of-the-art HA solutions for Cloud Computing. The authors hope that the observation of such solutions could be used as a good starting point to addressing with some of the problems present in the HA Cloud Computing area. Cloud computing has certainly benefited many enterprises by reducing costs and allowing them to concentrate on their core business competence rather than IT and infrastructure issues. But, for all the generally well-earned hype, there are still distinct **disadvantages of Cloud Computing** – especially relating to smaller operations – that you should consider before taking the leap. In this research paper , I'll try to offer some key concerns along with strategies and technologies for addressing them.

The six main disadvantages of Cloud Computing:

- 1) Downtime
- 2) Cloud Computing disadvantages: security and privacy
- 3) Cloud Computing disadvantages: vulnerability to attack
- 4) Limited control and flexibility
- 5) Cloud Computing platform dependencies
- 6) Cloud Computing costs

Even with all of the above disadvantages of Cloud Computing , the environment has immense potential for many business models. As platforms mature and the economies of scale continue to grow, costs will continue to fall and reliability and security standards will improve.

This work is structured as follows: "Cloud outages" section describes some Cloud outages that occurred in 2014 and 2015, and how administrators overcame these problems; "Systematic review " section presents the

methodology used to guide our systematic review for solving Downtime disadvantage; “Overview of high availability in Clouds” section presents an overview regarding HA Cloud solutions; “Results description” section describes works about HA services based on our systematic review result; “Discussions” section discusses some research challenges in this area; and “Final considerations” section delineates final considerations

II.SYSTEMATIC REVIEW

In this work, we adapted the systematic review proposed by [6], in order to find strategies that address HA Clouds. Next, we describe each activity (see Fig. 1) in detail and describe how we address it.

[Open image in new window](#)

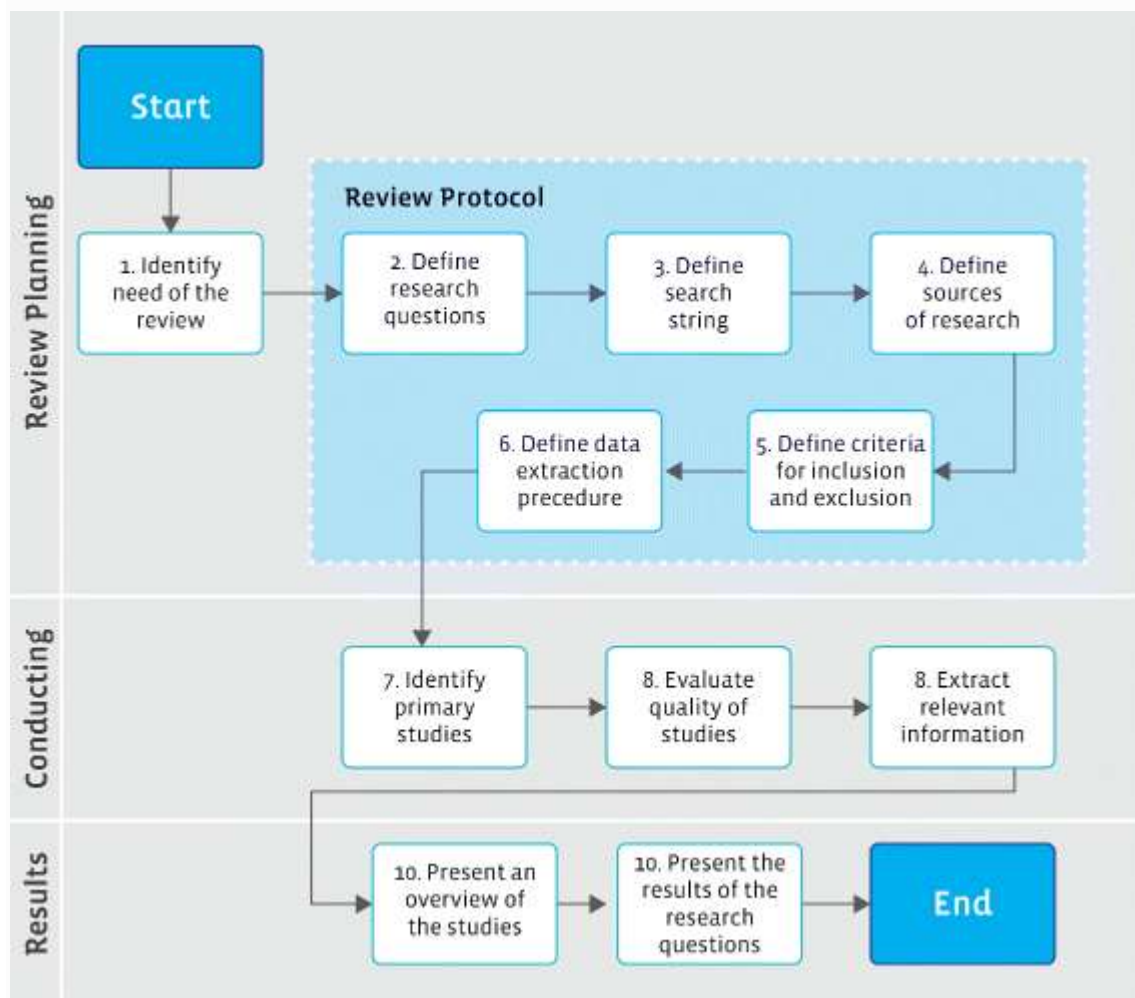


Fig. 1 Systematic review process

Activity 1: identify the need for the review

As stated previously, high availability in Clouds remains a big challenge for providers since Cloud infrastructure systems are very complex and must address different services with different requirements. In order to reach a certain level of high availability, a Cloud provider should monitor its resources and deployed services continuously. With information about resources and service behaviors available, a Cloud provider could make good management decisions in order to avoid outages or failures.

Activity 2: define research questions

In this activity, we need to define which questions we want to answer. The main goal of this work is to answer the following research questions (RQ):

- RQ.1: What is the current state-of-the-art in HA Clouds?
- RQ.2: What is the most common definition of HA?
- RQ.3: What are the HA services implemented by HA Cloud solutions?
- RQ.4: What are the most common approaches used to evaluate HA Cloud solutions?
- RQ.5: What are the research challenges in HA Clouds?

Activity 3: define search string

In this activity, we need to define which keywords we will use in selected search tools. For this work, we used the following expressions: “cloud computing” AND “high availability” AND “middleware”.

Activity 4: define sources of research

For this work, we chose the following databases: IEEE Xplore⁵, Science Direct⁶, and ACM Digital Library⁷.

Activity 5:

define criteria for inclusion and exclusion In order to limit the scope of this analysis, we considered only journals and conferences articles published between 2010 and 2015. The keywords “cloud computing” and “middleware” or “framework” were required to be in the article.

Activity 6: define data extraction procedure

Data extraction is based on a set of items to be filled for each article: keywords, proposal, and future works.

Activity 7: identify primary studies

The search returned 9, 63, and 145 articles in IEEE Xplore, Science Direct, and ACM Digital Library, respectively, totaling 217 works.

By reading all abstracts and using the criteria for inclusion or exclusion, we selected 19 papers for data extraction and quality evaluation. This number is justified because the keyword “high availability” is very common in Cloud Computing, especially in its own definition, and so most of articles had this keyword in them. However, in most cases high availability was not their research focus.

Activity 8: evaluate quality of studies

The quality evaluation was based on checking if the paper is related to some HA Cloud proposal for middleware or framework.

Activity 9: extract relevant information

This activity involves applying the data extraction procedure defined in Activity 6 to the primary studies selected in Activity 7.

Activity 10: present an overview of the studies

In this activity, we present an overview of all articles we selected in Activity 8, in order to classify and clarify them according to the research questions presented in Activity 2. The result of this activity is presented in “[Overview of high availability in Clouds](#)” section.

Activity 11: present the results of the research questions

After an overview about studies in HA Clouds, we had a discussion in order to answer the research questions stated in Activity 2. The results of this activity are presented in “Overview of high availability in Clouds” section.

Overview of high availability in Clouds

In this Section, we present an overview about Activity 10, presenting some characteristics of the selected articles in HA Cloud. Figure 2 shows the number of articles published per year from 2010 to 2015.

[Open image in new window](#)




Fig. 2 Number of articles per year

Concerning research source (Fig. 3), we can see that ACM has more articles published in HA Cloud area.

[Open image in new window](#)




Fig. 3 Number of articles per research source

Some articles define the term "high availability". For instance, authors in [7] say "the services provided by the applications are considered highly available if they are accessible 99.999 % of the time (also known as five 9's)". The Table 1 outlines the various definitions of "high availability" we identified through our research, as well as the source of each definition.

Table 1

High availability definitions

Reference	Definition
Achieving High Availability at the Application Level in the Cloud [7]	The services provided by the applications are considered highly available if they are accessible 99.999 % of the time (also known as five 9's)
Managing Application Level Elasticity and Availability [25]	High availability is achieved when the outage is less than 5.25 minutes per year
Scheduling highly available applications on cloud environments [35]	High availability systems are characterized by fewer failures and faster repair times
Are clouds ready for large distributed applications? [36]	High availability is defined in terms of downtime that is the total number of minutes the site is unavailable for events lasting longer than 5 minutes over a 1-year period
Software aging in the eucalyptus cloud	Availability is defined as the ability of a system to perform

Reference	Definition
computing infrastructure: Characterization and rejuvenation [37]	its slated function at a specific instant of time.

We also observed that many services are implemented in conjunction in order to offer a HA Cloud.

Figure 4 shows monitoring, replication, and failure detection are the most implemented services, identified in 50 % of studies in the research. Please, note that there are more services than published works because it is common to implement more than one service in a proposal.

[Open image in new window](#)

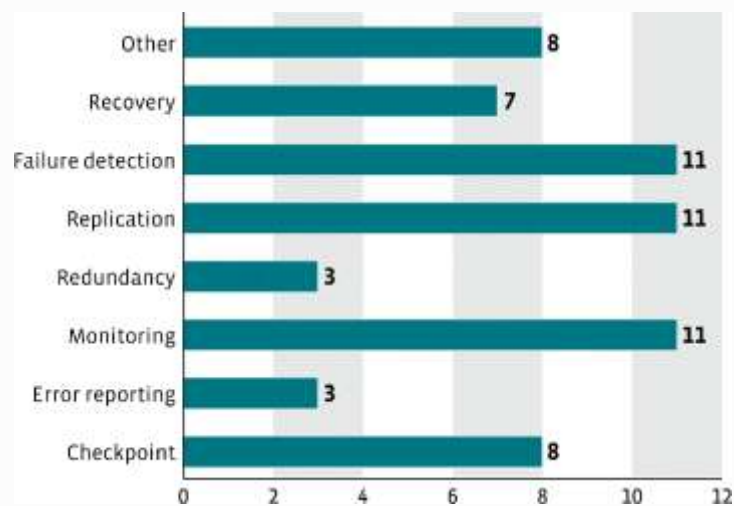


Fig. 4

HA services implemented by solutions

Figure 5 shows how solutions were evaluated in the studies we analyzed. We can see experimentation is the most popular technique used. These results indicate that research about this topic is working to derive proposals with fast application to the cloud computing industry.

[Open image in new window](#)




Fig. 5

Approaches used to evaluate HA solutions

The analysis should be performed based on comparison metrics. Work presented in [8] defines some metrics used to evaluate HA solutions, as shown in Table 2.

Table 2

Metrics for HA evaluation from [8]

Metric	Definition
Reaction time	Delay between the occurrence of the failure and the first reaction of the availability management solution.
Repair time	Duration from the first reaction until the faulty entity is repaired.
Recovery time	Duration from the first reaction until the service is provided again
Outage time	Time between the failure happening and the service recovery. In other words, outage time is the amount of time the service is not provided and it is the sum of the reaction and recovery times.

Now Approaches are Methods to solving Downtime Problem by using Oracle techniques used to evaluate HA solutions



III.CLOUD DOWNTIME

Cloud Computing has become increasingly essential to the live services offered and maintained by many companies. Its infrastructure should attend to unpredictable demand and should always be available (as long as possible) to end-clients. However, assuring high availability has been a major challenge for Cloud providers. To illustrate this issue, we describe technology or methods to solve .

Introduction When business-critical systems are down for a variety of reasons ranging from hardware or software failures to planned downtime for maintenance or upgrades, the costs can be astronomical. This is true for SAP infrastructures on which hundreds of users depend for their work and especially when systems such as materials management, supply chain management (SCM), and customer relationship management (CRM) run concurrently and interact with each other. In this environment, almost every component requires high availability. For example, a call center without access to the CRM application is as unproductive as a shipping company without a logistics tracking system. To help address this difficult dilemma, Oracle offers a best practice procedure using Oracle Solaris Containers and Oracle Solaris ZFS innovative virtualization technologies that are included in the Oracle Solaris 10 operating system (OS). This procedure is designed to dramatically reduce planned downtime of the overall SAP environment and to enhance the efficiency of IT operations.

1 Minimizing Planned Downtime of SAP Systems with the Virtualization Technologies in Oracle Solaris 10

The Imperative to Minimize Downtime

The majority of annual service downtime in datacenters is not a result of unplanned downtime caused by failures and outages of server or software components. Rather, most downtime (80 percent or more) is due to maintenance, upgrades, and similar planned tasks. Predefined maintenance windows are created for these tasks, usually at night or on weekends so that the work affects a minimum number of users. Planned downtime is calculable for all users and thus tolerable to a certain extent. However, the global economy is generating a demand for more 24/7 services, and as a consequence, maintenance windows are becoming smaller and smaller. Thus it is imperative for IT managers to limit or eliminate maintenance windows.

Shortening Maintenance Windows for SAP Portal Infrastructures

The availability of the SAP Enterprise Portal component in SAP NetWeaver is especially critical in SAP environments for the following reasons:

- x It is the central access point for all employees of a company.
- x It serves as the central communication point among several applications and their users.
- x Many SAP features are currently available only through the SAP Enterprise Portal.
- x In addition to being a central internal access point, it is an access point for suppliers, customers, and partners. Downtime can tarnish the reputation of the company.
- x SAP Enterprise Portal creates an open environment that must be as secure as possible, requiring immediate action when security updates are released.

Cloning the SAP Portal Environment Using Oracle Solaris Containers

With Oracle Solaris ZFS, it is possible to create a snapshot, or clone, of the entire runtime environment, including all running SAP applications such as the SAP Enterprise Portal and Oracle Database within that environment. The clone is identical in behavior to the original system and is available for use within seconds; only the Internet Protocol (IP) address is different. Thus, the clone can be redefined as the main system by simply changing the Domain Name System (DNS) entry after the upgrade has been performed.

Procedure to



Minimize Planned Downtime for the SAP System Upgrading SAP Enterprise Portal to a newer version or importing SAP Service Packages typically requires multiple hours of downtime. The following steps allow you to minimize this downtime to a few seconds, from the users' point of view, using the cloning capability of Oracle Solaris Zones and Oracle Solaris ZFS.

2 Minimizing Planned Downtime of SAP Systems with the Virtualization Technologies in Oracle Solaris 10

The procedure consists of three fundamental steps. The following description assumes that the basic zone configuration tasks such as zone definition, CPU resource pool assignments, and the creation of ZFS file systems have already been made.

Step 1: Create the Clone Zone

In this step, you create the clone of the productive system as shown in Figure 1. This is done in online mode using Oracle Solaris ZFS.

- x Create a snapshot and a clone. Oracle Solaris ZFS allows you to create a snapshot of a file system. Based on that snapshot, a clone can be created. A clone is a kind of writeable snapshot. Technically, it is comparable to a delta copy of the source file system. Only the changes from the original file system are saved. Therefore, a clone does not require any disk space initially, but it grows as data blocks are changed within the cloned file system.
- x Create a new zone based on the cloned file system. After the clone is created, a new zone on top of the clone file system is configured.
- x Adapt zone OS configuration entries. To configure a new host name, host entries, and other such identifiers for the zone, you need to apply some configurations within the zone.

Step 2: Upgrade the Clone System

After the clone is created and booted, that system can be started for applying the SAP software update. Perform the SAP Enterprise Portal upgrade on the clone. During this process, the production system is still available.

- x Run tests. You can perform tests such as a quality assurance (QA) evaluation on the newly upgraded clone.
- x Prepare for the switch. After successfully testing, you are ready to switch all users to the cloned upgraded system.

Step 3: Switch the Systems

A switch of the system is performed by changing the virtual network interfaces. Both systems have a public virtual interface. The IP addresses of these interfaces will be switched. Afterward, the cloned system is publicly available, using the virtual address registered in the DNS.

- x Promote the clone (optional). If System B is to become the new `_mastervolume`, promoting the clone is required. For example, this would be necessary if you want to delete the volume of System A afterward. To learn more about Oracle Solaris ZFS capabilities such as promoting a clone, consult Oracle's online resources. This feature is not included in early versions of Oracle Solaris ZFS.
- x Adapt zone configuration files. To switch the systems in a persistent manner, it is necessary to edit the zone configuration files and add the new entries, such as IP addresses.
- x Switch some DIs in advance (optional). To avoid a heavy workload on the primary application server, you may decide to switch some of the application servers in advance.

3 Minimizing Planned Downtime of SAP Systems with the Virtualization Technologies in Oracle Solaris 10

- x Switch resource pools. Before finally performing the switch of the systems, you should switch the resource pools as well to provide the clone system with the resources reserved for the productive system and vice versa.
- x Switch network interface controllers (NICs) and `Esapmnt share`. When you perform the switch of the IP addresses and the `_sapmnt share`, System B becomes the productive system and System A becomes the shadow system.
- x Shut down the old system (optional). To free resources, you may decide to shut down and delete the old system afterward. The steps above represent only a short overview. For detailed

information, including an extensive step-by-step guide, go to sdn.sap.com/irj/scn/weblogs?blog=/pub/wlg/4804. To minimize failures and to speed up the process described above, it also makes sense to automate this procedure. In many cases, SAP Enterprise Portal environments do not store dynamic data. Consequently, there is no loss of data when switching from the primary environment to the clone environment. Therefore, the switch from one system to the other is completely transparent to the users. The system procedure for minimizing planned downtime for SAP has been analyzed by SAP, particularly for SAP portals. The results are documented and posted on the SAP Developer Network (SDN) at sdn.sap.com/irj/scn/go/portal/prtroot/docs/library/uuid/a02397f9-0d8d-2910-e5a0-b8241bbc9381?QuickLink=index&overridelayout=true

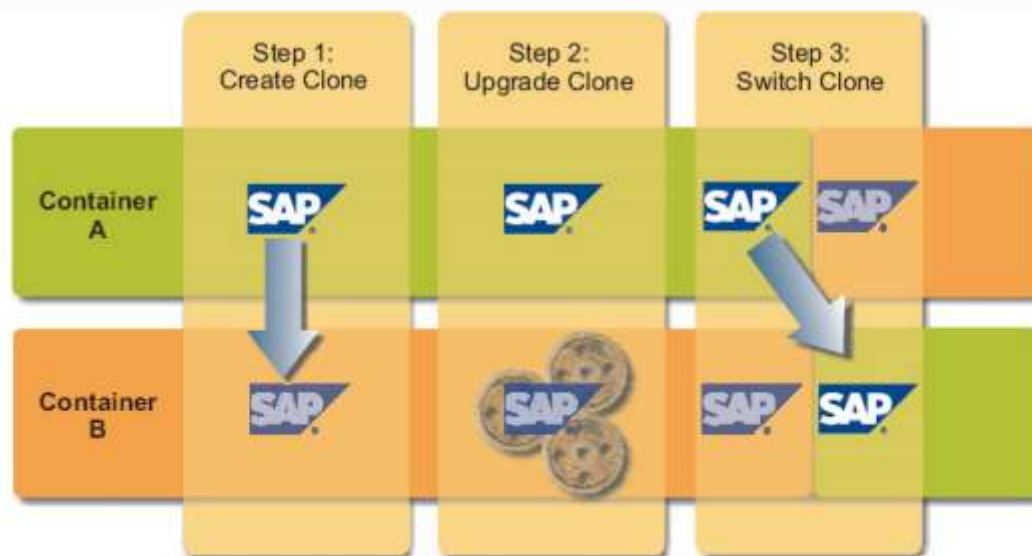


Figure 1. Oracle Solaris Containers and Oracle Solaris ZFS allow IT departments to minimize planned downtime with cloning capability.

Cloning Other SAP Components The clone functionality works with any container in Oracle Solaris Containers, regardless of the application that it contains, so this cloning procedure can be applied to other components in the SAP environment. For example, clones can be used for training, educational, and development purposes.

Minimizing System Maintenance Downtime It is sometimes necessary to take an entire system down for maintenance. The same features of Oracle Solaris Containers that enable cloning can also be employed to move containers on a shared file system from one Oracle Solaris 10 system to another, providing that both systems are at the same patch level. Containers are moved by executing the Oracle Solaris detach and attach commands in the global zone of the system through those scripts. The steps are as follows:

- x Move all containers from one server, which is in maintenance mode, to any other server that has sufficient resources available.
- x Execute all maintenance tasks on the idle server.
- x Move the containers back after maintenance is complete.

The ability to

move containers between systems enables system maintenance without downtime, even on systems that house many consolidated applications. Oracle Solaris Containers An integral part of Oracle Solaris 10, Oracle Solaris Containers isolate software applications and services by using flexible, software-defined boundaries. They allow many private execution environments to be created within a single instance of Oracle Solaris 10. Each environment has its own identity, separate from the underlying hardware, so it behaves as if it is running on its own systemW making application virtualization simple, safe, and secure. In contrast to other virtualization approaches, Oracle Solaris Containers do not require you to install and maintain a dedicated operating system for each virtualized computing environment. Unlike hypervisor based virtualization, the loss of computing power resources is nearly zero when Oracle Solaris Containers are deployed.



Figure 2: Oracle Solaris Containers enable simple consolidation with isolated, software-defined zones for application and services.

Oracle Solaris Containers can do the following:

- x Dynamically move or replicate applications to fit the changes of the business
- x Lower administrative costs by safely combining multiple applications onto a single system
- x Reduce conflicts among applications that are running on the same system by isolating them from one another
- x Offer the predictive self healing feature to minimize fault propagation and unplanned downtime
- x Enhance security by preventing unauthorized access and unintended intrusions

Oracle Solaris ZFS Oracle Solaris ZFS is a file system included in Oracle Solaris 10. It provides very high levels of data integrity and performance, and it improves the ease of file system management by eliminating the need for a volume manager. Furthermore, because it is 128 bit based, it opens the door to virtually unlimited data scalability. Oracle Solaris ZFS has been integrated with Oracle Solaris Containers technology, allowing container administrators to take full advantage of the features of Oracle Solaris ZFS. Designed with the administrator in mind, Oracle Solaris ZFS is the only self-managing, general-purpose file system with predictive self healing




Figure 3. Oracle Solaris ZFS offers simpler administration of high-integrity data on a greater scale and with higher performance than traditional file systems.

Oracle Solaris ZFS offers the following:

- x Simple administration. It automates and consolidates complicated storage administration concepts such as copy-onwrite and snapshots, reducing administrative overhead by 80 percent.
- x Provable data integrity. It protects all data with 64-bit checksums that detect and correct silent data corruption.
- x Unlimited scalability. It offers 16 billion times the capacity of 32 bit or 64 bit systems.
- x Blazing performance. It is based on a transactional object model that removes most of the traditional constraints on the order in which an I/O is issued, resulting in huge performance gains.

Other Uses for Oracle Solaris Containers and Oracle Solaris ZFS Minimizing planned downtime is only one area in which Oracle Solaris Containers and Oracle Solaris ZFS offer a smart solution. Across the complete SAP system landscape, an average of 70 to 80 percent of installed systems are operated to support development, testing, upgrade, and evaluation tasks However, those systems must still be managed and maintained. To effectively operate these nonproductive environments, IT managers need to be able to provision SAP systems for SAP users as quickly and flexibly as possible. Versioning a complete SAP infrastructure after an upgrade is a typical requirement of an ongoing, leading-edge SAP operation as well. All of these requirements can be supported with Oracle Solaris Containers and Oracle Solaris ZFS.

Reference architecture can also serve as a building block toward next-level platforms for production environments. Leveraging Oracle Solaris Cluster helps you implement higher availability, a necessity in production and nonproduction environments alike. Oracle offers a full range of services from capacity planning, solution design, architecture, and deployment to migration, consolidation, and integration of independent software vendor (ISV) enterprise applications. This is all handled by Oracle consulting and SAP solution experts.

IV.CONCLUSION

In today's competitive business world of 24/7 production and operations, downtime whether planned or unplanned is becoming increasingly difficult to tolerate. The availability of all SAP applications, especially

SAP Enterprise Portal, is critical. The ability to optimize SAP processes even during production enables maximum availability in heterogeneous SAP landscapes. Oracle Solaris Containers and Oracle Solaris ZFS minimize downtime during upgrades of SAP portals and optimize SAP processes. In addition, these technologies are included in Oracle Solaris 10. The ability to clone, copy, and move containers and systems on demand and with minimal downtime allows your datacenter to continue operating at optimal performance and at a lower cost. The best practices described in this white paper help IT managers to utilize the technologies in their SAP environments, automating the datacenter for better management.

V.FINAL CONSIDERATIONS

Cloud Downtime , no matter how long, are responsible for large financial losses. Cloud providers look for solutions that provide high availability even in failure cases. In this paper, we proposed a classification for HA Cloud solutions based on 3 layers. We also described and discussed some existing commercial and non-commercial solutions focused on techniques.

High availability is a great challenge for Cloud providers due to its complexity (from the infrastructure to the application level). There are many issues to study in order to minimize Clouds outages, such as portability, feasibility, and security. A next step could be the implementation of HA-as-a-service, highlighting even more the importance of this research area for Cloud providers. Consider these two key points:

→ Which of your business processes can be delayed or halted if the service provider goes down?

→ When your internet connection is down, all your applications drop offline.

Best Practices for minimizing planned downtime in an SAP environment: Demand a service level agreement (SLA) from your provider guaranteeing uptimes in excess of 99.55% (which equals 1.83 days of downtime a year, or 3.60 hours of downtime a month).

Endnotes

- <http://iwgcr.org/>
- <https://www.cloudfoundry.org>
- <https://blogs.dropbox.com/tech/2014/01/outage-post-mortem/>
- <http://www.verizonwireless.com/solutions-and-services/verizon-cloud/>
- <http://ieeexplore.ieee.org/Xplore/home.jsp>
- <http://www.sciencedirect.com/>
- <http://dl.acm.org/>
- <https://www.xenproject.org/>



- http://www.linux-kvm.org/page/Main_Page
- <http://www.vmware.com/>
- <https://www.microsoft.com/en-us/cloud-platform/virtualization>
- <http://docs.openstack.org/developer/heat/>
- <https://wiki.openstack.org/wiki/Heat/HA>
- <https://cloud.google.com/appengine/>
- <https://link.springer.com/article/10.1186/s13677-016-0066-8>
- <https://cloudacademy.com/blog/disadvantages-of-cloud-computing/>
- <http://searchitchannel.techtarget.com/definition/service-level-agreement>

REFERENCES

- [1.] Dillon T, Wu C, Chang E (2010) Cloud computing: issues and challenges In: Advanced Information Networking and Applications (AINA), 2010 24th IEEE International Conference On, 27–33. IEEE. <http://ieeexplore.ieee.org/document/5474674/?arnumber=5474674&tag=1>.
- [2.] Puthal D, Sahoo B, Mishra S, Swain S (2015) Cloud computing features, issues, and challenges: a big picture In: Computational Intelligence and Networks (CINE), 2015 International Conference On, 116–123. IEEE. <http://ieeexplore.ieee.org/document/7053814/?arnumber=7053814>.
- [3.] da Fonseca NL, Boutaba R (2015) Cloud Architectures, Networks, Services, and Management In: Cloud Services, Networking, and Management. Wiley-IEEE Press. <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7090261>.
- [4.] Cérin C, Coti C, Delort P, Diaz F, Gagnaire M, Gaumer Q, Guillaume N, Lous J, Lubiars S, Raffaelli J, et al. (2013) Downtime statistics of current cloud solutions. International Working Group on Cloud Computing Resiliency, Tech. Rep. <http://iwgcr.org/wp-content/uploads/2013/06/IWGCR-Paris.Ranking-003.2-en.pdf>. Accessed Oct 2016.
- [5.] Toeroe M, Tam F (2012) Service Availability: Principles and Practice. John Wiley & Sons. <http://www.wiley.com/WileyCDA/WileyTitle/productCd-1119954088.html>.
- [6.] Coutinho EF, de Carvalho Sousa FR, Rego PAL, Gomes DG, de Souza JN (2015) Elasticity in cloud computing: a survey. Ann. Telecommunications-Annales des télécommunications 70(7–8): 289–309.
- [7.] Kanso A, Lemieux Y (2013) Achieving high availability at the application level in the cloud In: Cloud Computing (CLOUD), 2013 IEEE Sixth International Conference On, 778–785. IEEE. <http://ieeexplore.ieee.org/document/6740222/?arnumber=6740222>.
- [8.] Heidari P, Hormati M, Toeroe M, Al Ahmad Y, Khendek F (2015) Integrating open saf high availability solution with open stack In: Services (SERVICES), 2015 IEEE World Congress On, 229–236. IEEE. <http://ieeexplore.ieee.org/document/7196529/?arnumber=7196529>.

- [9.] Hormati M, Khendek F, Toeroe M (2014) Towards an evaluation framework for availability solutions in the cloud In: Software Reliability Engineering Workshops (ISSREW), 2014 IEEE International Symposium On, 43–46. IEEE. <http://ieeexplore.ieee.org/document/6983798/?arnumber=6983798>.
- [10.] OpenSAF Overview Release 4.4 Programmer's Reference.
<http://sourceforge.net/projects/opensaf/files/docs/opensaf-documentation-4.4.1.tar.gz/download>. Accessed Oct 2016.