

A large, faint magnifying glass graphic is centered on the page, with its handle pointing towards the bottom left. The lens of the magnifying glass is positioned behind the main title text.

# What It's Like to Work With F3

Client Case Studies



# About F3 Technology Partners

## Complete IT Solutions

F3 Technology Partners specializes in products and services that focus on Server and Storage, Virtualization and Database technologies. With our dedication to superb customer service, F3 Technology Partners can design and deliver a cost effective strategy based on the current – and future – needs of your business. The end result will be a more agile and efficient system, which in turn allows your business to innovate faster and meet your customer's needs.

## Let Our Experience Help Yours.



Oracle Database



Virtualization



Storage Solutions



Systems

# Table of Contents

## 1. Workload Consolidation on Oracle Solaris

**The Problem .....5**

**The Solution ..... 6**

## 2. Oracle Database 11g + ASM on Red Hat 6.3

**The Problem ..... 8**

**Background .....8**

**The Solution .....9**

## 3. Linux Kernel Crash Dump Configuration

**The Problem ..... 13**

**Background ..... 13**

**The Solution .....14**

**Kdump configuration .....15**

## 4. Consolidating and Modernizing an Oracle Sun Environment

**The Problem ..... 17**

**Background ..... 17**

**The Solution ..... 18**

1.

# Workload Consolidation on Oracle Solaris

FEB

# The Background

A large Hartford area insurance company that has been a long standing customer of F3 Technology Partners, had aging Sun hardware that was approaching end of service life and needed to be refreshed. In addition, a datacenter consolidation effort for a secondary location was also being planned out. However, due to a power limitation in their primary datacenter, the solution could not be a simple one to one replacement of physical hardware.

The power limitation severely restricted how many new servers could be hosted in the primary location. The aging hardware and datacenter constraints were impacting the organization's ability to support their business needs.

*“ F3 Technology Partners' engineers realized that a complete understanding of the current IT infrastructure was essential in order for the correct solution to be architected. ”*

F3 Technology Partners' engineers realized that a complete understanding of the current IT infrastructure was essential in order for the correct solution to be architected.

This was accomplished by working with architects and engineers inside the insurance company to learn how their systems were utilized. This information was critical to our understanding of how a consolidated environment would need to perform to meet their organization's demands.

## The Solution

Our recommendation was a Solaris Container Farm that would consist of multiple tiered clusters. We had clusters built around single threaded performance and others for multi-threaded performance, along with Production and Non-production variants. These clusters were controlled by Veritas VCS which was chosen to align with customer knowledge and that environment was designed to support both branded Solaris 8 and 9 containers along with Solaris 10 native containers.



The customer's main goal for this was to prolong the Solaris environment with negligible impact. Branded containers allowed a smooth migration from the physical environment to the new consolidated environment by means of OS flash archives. SANS LUNS for application data were able to be masked to the new environment. This left us with the ability to move 100's of GB's of data into the consolidated environment without having to do a data copy. It also created a solid fall back plan to the physical server if application issues were encountered.

CPU pools were implemented for workload segregation and QoS (Quality of Service). We tried multiple configurations of CPU controls and found this method the most beneficial. Applications such as Oracle, that were sensitive to CPU count changes along with performance reporting, benefited.

Memory caps were then implemented and sized according to workload demand. Many applications were on servers with larger memory footprints than what we could accommodate in the new environment. This was mediated by working with the organization's Performance and Capacity team to analyze what the true demand of the workload was. In many cases, the demand was much smaller than what the physical server was sized for. As there is no native mechanism to prevent memory oversubscription with memory capping, it was required that accurate reports and audits be performed as the environment grew.

Custom modifications to VCS trigger scripts allowed for automation of pool creation and destruction. By working with this Veritas framework which maximized automation, less administrative overhead was necessary, in addition to the high availability of production systems at the container level via VCS. The container concept also enhanced the customers' ability to deploy net new Solaris workloads more quickly and efficiently than before, while saving money on hardware costs. Future hardware upgrades will also be less impactful as the containers can be migrated to new hosts very quickly. The lifespan of legacy applications was prolonged as well.

Each cluster of 2 nodes shared the same LUNS. This allowed a container to be portable between the 2 nodes in the event one of the nodes failed. Another benefit was portability when a container needed to be moved from T series hardware cluster to M series for single threaded performance demands. The Disaster Recovery strategy to their remote facility was also improved since it is now more easily achieved at the operating system level with SRDF replication of the container root LUN and corresponding application LUNS.

F3 Technology engineers then worked with the in-house project managers and their teams to successfully migrate their applications from the secondary data center into the consolidated environment and migrate from existing physical servers to virtual in the primary datacenter. Finally, guidelines for growth were developed to limit the number of containers in a single cluster and for when

*“ F3 Technology engineers then worked with the in-house project managers and their teams to successfully migrate their applications from the secondary data center... ”*

a new cluster would need to be deployed to accommodate incoming migrations. Documentation of all of these processes was created to help the customer support the restructured environment going forward.

F3 Technology Partners provided both expertise and resources to complete this project. Our ability to design, implement, and support the solution enabled the project to be completed efficiently and with minimal disruption to the customer's environment. While improving the company's data backup and disaster recovery systems, both current and future hardware costs were significantly reduced. The new infrastructure has the added benefit of new features and capabilities which improve the customer's ability to run their business and meet their customers' needs.

2.


# Oracle Database 11gR2 + ASM on Red Hat 6.3

FEB



# The Problem

A west coast credit union that specializes in banking services, insurance, loans, and investment services, had made a strategic decision to move all of their legacy hardware and software from Windows to newer versions of Linux. They chose Oracle Enterprise Linux (OEL) 6.3 on DELL servers as their investment to support their growing business needs.



They needed to stick with a Red Hat Compatible Kernel (RHCK) as there was a downstream dependency of NetApp snapshotting tools such as Snap Manager for Oracle, Snapdrive, etc., that needed RHCK and were not certified on the Oracle-proprietary Unbreakable Enterprise Kernel (UEK). The customer also made a decision to support their business on database stored on the industry standard storage management of the database: Oracle Automatic Storage Management, aka Oracle ASM. The ASM Library is a set of tools and kernel modules designed to ease management of ASM device discovery and persistent naming of the storage devices from the host.

It used to be the case that Oracle regularly released with ASMLib for all flavors of Linux as a separate installable rpm in Oracle Linux / RHEL 5.x. But, starting with RHEL / Oracle Linux 6.x, Oracle provides ASMLib as part of the standard kernel build ONLY FOR UEK. It no longer releases or plans to release the ASMLib for non-UEK kernels, as detailed on the Oracle website in the Metalink note titled "Oracle ASMLib Software Update Policy for Red Hat Enterprise Linux Supported by Red Hat (Doc ID 1089399.1)."

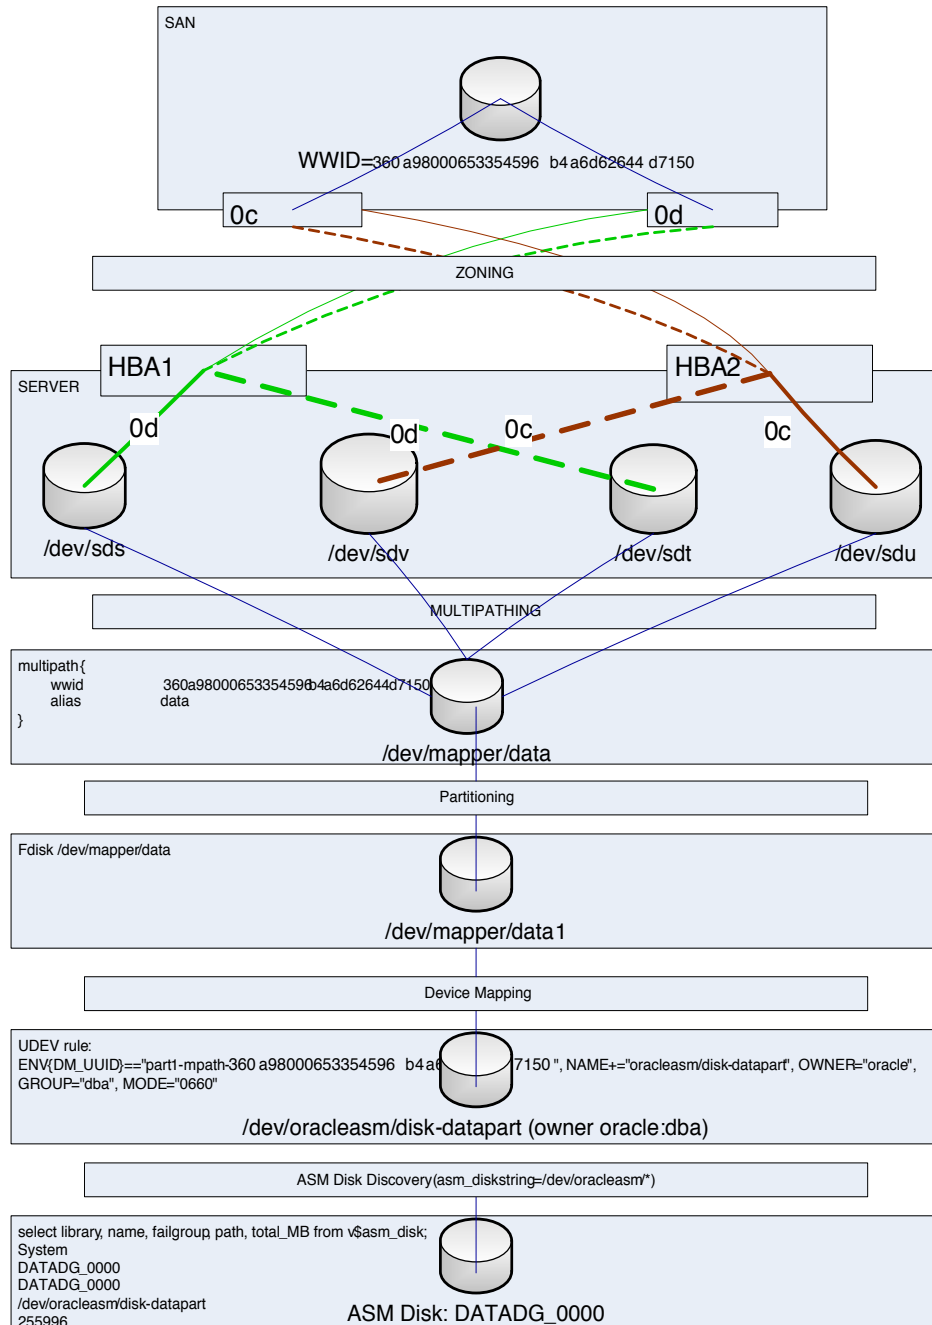
The customer then found themselves in the unenviable position of either giving up their NetApp snapshot tools to accommodate Oracle ASM or give up Oracle ASM to stick with their NetApp tools. The end-state solution implementation is captured in the attached schematic diagram. The customer has

*“ The F3 Technology Partners’ engineering team, after extensive study and research, came up with a design that allowed these two systems to work together. ”*

benefited immensely from the solution that was discovered and implemented by the F3 team. They now have functionality of all the previous NetApp backup technology in addition to Oracle ASM on Oracle Linux. They also now enjoy the performance gains that ASM brings to the table; the elimination of OS from I/O and allowing Oracle RDBMS kernel handle the I/O as if it was a raw device.

# The Solution

This was a totally unique scenario with no information, let alone solution, on the internet. It became a case of F3 Technology Partners discovering our own solution. While it is worthwhile to note that there was a ton of official information available for Oracle 5.x, there was no official documentation on how to go about configuring ASM disks on the partitioned multipath-ed disk from Oracle for OEL6.3. This document attempts to provide that information: how to configure ASM on a multipath-ed disk partition using udev rules.



# The Solution Details

There were two challenges that needed to be addressed: the first was partition alignment recommendation by NetApp and the other was the fact that in OEL6, you cannot define permissions for block devices in the `/etc/multipath.conf` file.

## Partition Alignment:

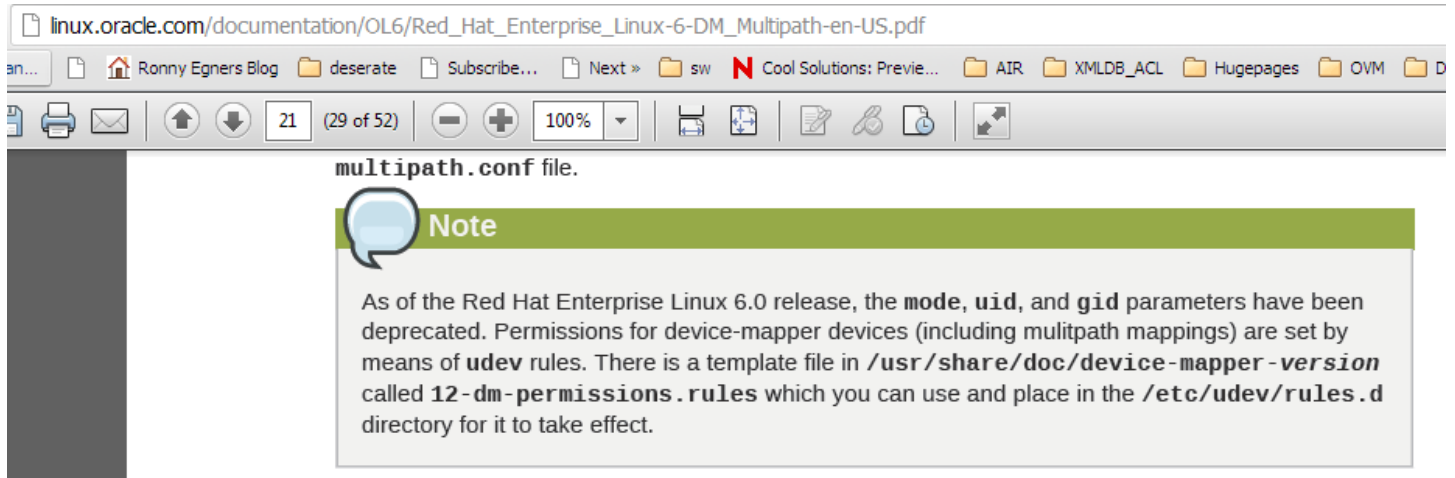
One challenge to the entire solution architecture was identifying the correct udev rule that will consistently create the device node under `/dev` based out of the partition, not the whole device. This is because of complying with NetApp-recommended partition alignment architecture for LUNs sourced from NetApp filers as detailed in Knowledge Base article KB ID:1010717 and Bug 156121. This is necessary to achieve optimal performance of read and write I/Os. The reason necessitating the alignment is that if the ASM Disk partition aligns with (begins at the same) boundary of the LUN's volume block, the I/O performance will be optimized. In the absence of this partition alignment, there's a potential that 2 LUN volume blocks will need to be read (or written for writes) for a read operation that crosses the unaligned boundary which could've been served by a single volume block read in a aligned case.

## Defining persistent disks with correct permissions on the pseudo device:

It used to be the case in OEL5.x that the permissions can be defined in the `/etc/multipath.conf` file for a multipath-ed device as below:

```
multipath {  
  wwid 360a91234567890596b4a6d62644d7250  
  alias ASMDISK_Voting_1  
  mode 660  
  uid 301  
  gid 300  
}
```

Starting with OEL6.x, that feature is deprecated.



The screenshot shows a PDF viewer window with the title bar "linux.oracle.com/documentation/OL6/Red\_Hat\_Enterprise\_Linux-6-DM\_Multipath-en-US.pdf". The document content is titled "multipath.conf file." and features a green "Note" box with a speech bubble icon. The note text reads: "As of the Red Hat Enterprise Linux 6.0 release, the **mode**, **uid**, and **gid** parameters have been deprecated. Permissions for device-mapper devices (including multipath mappings) are set by means of **udev** rules. There is a template file in `/usr/share/doc/device-mapper-version` called `12-dm-permissions.rules` which you can use and place in the `/etc/udev/rules.d` directory for it to take effect."

So, the only recourse is to use correctly configured udev rules.

Both the factors of partitioning and correct permissions on the block device are taken care by the correctly configured udev rules. The following rule-block is the central piece of the whole architecture:

```
[root@oralinb rules.d]# cat /etc/udev/rules.d/99-oracle-asmdevices.rules
ENV{DM_UUID}=="part1-mpath-360a91234567890596b4a6d62644d7250",
NAME+="oracleasm/disk-datapart", OWNER="oracle", GROUP="dba", MODE="0660"
```

The above rule will create the `/dev/oracleasm/disk-datapart` device node and can be used as a ASM disk with the appropriate disk discovery string.

## References:

[http://linux.oracle.com/documentation/OL6/Red\\_Hat\\_Enterprise\\_Linux-6-DM\\_Multipath-en-US.pdf](http://linux.oracle.com/documentation/OL6/Red_Hat_Enterprise_Linux-6-DM_Multipath-en-US.pdf)

<https://kb.netapp.com/support/index?page=content&id=1010717>

<http://murtazahabib.wordpress.com/2012/08/20/4/>

Metalink Article: Configuring Oracle ASMLib on Multipath Disks [309815.1]

Metalink Article: White Paper: Oracle ASM and Multi-Pathing Technologies [294869.1]

Metalink Article: How To: How to Configure LUNs for ASM Disks using WWID, DM-Multipathing, and ASMLIB on RHEL 5/OEL 5 [1365511.1]

Metalink Article: Configuring non-raw multipath devices for Oracle Clusterware 11g (11.1.0, 11.2.0) on RHEL5/OEL5 [ID 605828.1]

Metalink Article: Configuring raw devices (multipath) for Oracle Clusterware 10g Release 2 (10.2.0) on RHEL5/OEL5 [ID 564580.1]

3.

# Linux Kernel Crash Dump Configuration

FE3

# The Problem

A long standing enterprise client of ours was searching for a way to implement a standard kdump configuration for storing system crash information for its Redhat Linux infrastructure. This client, a leading provider of insurance and wealth management services for millions of consumers and businesses worldwide, has a rapidly growing Linux infrastructure. The IT departments needed to have the ability to store system crash information and have it available when support needed it. A standardized kdump configuration was important to be able to roll it out to all of their varied systems.

## Background


The Linux operating environment is generally stable. It does not usually suffer irrecoverable errors that crash the system and require a reboot to restore to normal production. Nevertheless, these kinds of problems do occur from time to time. They are known as kernel crashes and are of utmost interest and importance to administrators in charge of these systems. Being able to collect them and analyze them, provides the system expert with a powerful tool in finding the root cause to crashes and possibly solving critical bugs.

Without the kdump information for their systems, root cause analysis was, in most cases, delayed. To collect the information, they would enable kdump for that system individually and wait for it to happen again. This procedure caused delays in problem resolution and more outages as the system might not run into the problem for a period of time.

Kdump is a flexible tool, with extended network-aware capabilities which gives it the ability to store crash data over the network. Kdump requires additional storage, either locally or remote to the system. In this company's current technological environment, a mixture of blade and individual systems, some systems did not have the required disk space or the ability to add storage to the system. Adding more storage would be costly considering the number of systems that would needed it and the configuration would need to be modified for systems that could not add the additional storage.

The customer decision to use Network Attached Storage would eliminate the need for costly local and SAN storage. It would allow a single, standard configuration, and would provide a single repository for storing the kdump data.

However, the organization did have concerns about using NAS as a storage location. The size of the



kdump, depending on the system, can be large. The time that it would take to write the data over the network could affect the recovery time of the systems since the kdump must complete before the system finishes rebooting and becomes available to its customers.

## The Solution

Once the NAS storage space was allocated through the company's own resources, F3 Technology Partners configured a couple of systems to test writing the kdump information. For testing reasons, the systems were purposely crashed multiple times to force a dump to the NAS. The test results were presented to the architecture team and F3 engineers then developed a script to help push out the kdump configuration via Bladelogic to the customer's Redhat environment. During the testing, F3 noted that there was a difference in kdump configuration between Redhat 5.6 and Redhat 5.8. Our engineers did the necessary updates to the configuration and the script in order to accommodate the two different Redhat releases and facilitate the roll out.

This enterprise customer now has a central location to collect kdump information when a system does crash, allowing for a speedier fix and a decrease in downtime. There is now one configuration that can be rolled out to all Redhat systems whether they are a blade or individually racked systems. The customer also has a script that will roll out the configuration for new servers and can handle multiple versions of Redhat. With the new configurations, the time it took to do the write to the NAS was acceptable and would not affect their SLAs with their customer base.

*“ F3 Technology Partners was able to provide resources that could configure, test, and work through anomalies found during the testing to provide a working solution based on the customers requirement for a standard configuration for multiple platforms. ”*

Throughout our long standing relationship with this customer, F3 has proven that our professional services continue to identify and fulfill the customer's requirements.

For the final solution, will list which files were modified, most of this is standard configuration from Redhat, but what will be noted is the differences found between Redhat 5.6 and 5.8. Also noted in the steps is that in Redhat 5.8, there were some performance issues when using the bond0 network interface. Parameters were changed to increase the performance.

## Kdump Configuration:

1. Check that the kdump package is installed on the servers.

```
# rpm -qa | grep -i dump
system-config-kdump-1.0.14-4.el5
```

2. Modify the `/boot/grub/grub.conf`

Need to add the `crashkernel=512@256M` parameter to the kernel.

```
From: kernel /vmlinuz-2.6.18-274.7.1.el5 ro root=LABEL=/ elavator=deadline
To: kernel /vmlinuz-2.6.18-274.7.1.el5 ro root=LABEL=/ elavator=deadline
crashkernel=512M@256M
```

3. Reboot to load the kdump module in kernel. `reboot`

4. Make sure the kdump module is loaded

```
# cat /proc/iomem | grep Crash\ kernel
```

5. Modify the `/etc/kdump.conf` file.

```
net hfdnas011:/vol/flx_496/VAR_CRASH
path /var/crash
core_collector makedumpfile -d 31 -c
options bonding mode=1 - Redhat 5.8
```

The options `bonding mode=1` was added as there was a performance problem when using the bond interface versus using a single network interface. This removed the performance problem.

6. In the bond0 file: `/etc/sysconfig/network-scripts/ifcfg-bond0`

```
BOOTPROTO=static - Redhat 5.8
```

Without `BOOTPROTO` the network information such as `/etc/resolv.conf` will not be saved in the kdump boot image as it was in version 5.6. Also, attempting to put the IP address information in the `kdump.conf` file did not resolve the problem either.

7. Add the `/var/crash` mount point for the NAS filer in the `/etc/fstab` file.

8. Restart Kdump service

```
service kdump restart
```

This will create a new kdump boot image.

9. Make the service starts up upon every reboot

```
chkconfig kdump on
```



4.

# Consolidating and Modernizing an Oracle Sun Environment

FEB

# The Problem

A well-known publishing company was interested in consolidating their forty existing, aging Sun servers into a small number of new Sun servers. They were not interested in making any application-level changes and wanted only the minimum amount of operating system-level changes necessary in order to achieve consolidation. The cost of ownership of their aging server infrastructure had become unacceptably high, and they were researching options to cut costs while avoiding major disruptions to their business-critical systems.

This customer was unfamiliar with the technology and processes required to consolidate their infrastructure. It was important to them to know up-front that F3 Technology Partners had prior experience in delivering this solution, could implement it effectively, and could provide knowledge transfer and post-implementation support. Prior to committing to the purchase of new hardware, the customer wanted to be confident that their investment would maintain or exceed current service levels which reducing their total cost of ownership.



## Background

*“ F3 Technology Partners took the time to learn about the customer’s current server infrastructure, document their requirements, constraints, and concerns, and gather data about their inventory.”*

We then analyzed the data and proposed a future state configuration which the customer approved. As a result of our analysis, our first step was to provide project management services to determine when the consolidations would take place, what the migration strategy would be for each existing host to the new environment, and what technical resources would be required for each migration. The F3 solution consisted of five new Oracle Solaris T4-2 servers connected to the customer’s existing storage infrastructure. Approximately 35 aging Sun hosts were virtualized onto five servers, a consolidation ratio of 7-to-1.

# The Solution

Using Oracle Solaris 10 containers, ZFS, and resource management technologies, our engineers were able to have the existing hosts be easily consolidated onto a small physical footprint. The identities, configuration, and service levels of the IT services, which had been running in the environment prior to consolidation, were all maintained during this process. We then performed the actual migration of each in-scope host onto the new Sun platform and helped ensure that all IT services running on the new infrastructure ran as well as - or better - than on the old.

By working in partnership with F3 Technology Partners, the customer gained valuable knowledge and insight into industry best practices for consolidating workloads in an Oracle/Sun environment. This solution, designed and delivered by F3, will enable the client to greatly reduce ongoing hardware support costs by reducing their physical server count from 35 to 5. The client will also benefit from a server platform in which new workloads can be deployed more quickly than before and service resiliency and availability are improved.

*“ Throughout this project, F3 Technology Partners demonstrated deep technical expertise, industry background, and customer focus. ”*

In addition, server resources, such as memory and processing, are more efficiently utilized and scalable than previously.

F3 was able to accommodate unforeseen circumstances and address these issues with speed and flexibility. In addition

to the implementation, our engineer's trained the customer's technical resources and managed the schedule which helped accelerate the completion of the project.

The customer was impressed with our ability to design the most optimal solution that fit both their business needs and budget, with a minimum of business interruption. The end result for this publishing company was a technology environment that reduced their costs and increased their efficiency resulting in increased productivity and business agility.

5.

**Storage**

**Modernization Project**

**FEB**

# The Background

F3 Technology Partners has had a long-standing business relationship with a large automotive auction company. This company's data was stored on unsupported storage which, while performing adequately and reliably, exposed the customer to major business risk in the event of an unplanned outage or failure.

## The Problem

The customer is heavily dependent upon their server infrastructure to run their business. They did have a number of SAN and NAS storage arrays but they were approaching end-of-support. In addition, they did not have the staff resources to migrate their data from their current arrays onto more modern hardware. Previous attempts at migrating to new storage had been unsuccessful due to both a lack of adequate documentation and knowledge about how and where the storage was being used. This storage issue exposed them to significant business risk in the event of a service outage, since the storage arrays held most of their business-critical data.

In collaboration with the customer, F3 Technology Partners performed a survey of all of their servers and produced documentation detailing all of the business's storage quantities, usage, and affected hosts. We identified how much new storage would be required to continue to meet their business needs. We then used this documentation to prioritize which storage would be migrated and when.

## The Solution

Our recommendation was to use Oracle ZFS storage appliances as the target storage platform. Our engineers then consolidated all of the customer's server data onto one Oracle 6140 SAN array, one Oracle 7320 clustered ZFS storage appliance, and one Oracle 7210 ZFS storage appliance. This solution provided three tiers of storage: SAN for data which requires the highest availability and performance; clustered NFS/CIFS for data which required high availability and access from Linux, Solaris and Windows clients; and non-clustered NFS/CIFS for less critical data which requires access from Linux, Solaris and Windows clients. The design also features a fourth storage unit, an Oracle 7210 ZFS storage appliance, which is located in a remote site. The two primary appliances replicate their data to this appliance for backup and disaster recovery purposes.

Following review with the customer, we then implemented the storage migration onto the new appliances. F3's system engineers performed all migrations over the course of several months, resolving several issues in the process. Data migration was achieved either through direct host-based copying of data from the old SAN arrays to the new storage, or by using the ZFS appliances' Shadow Migration feature which allows incremental copy of data via NFS.

Our recommendation was to use Oracle ZFS storage appliances as the target storage platform. Our engineers then consolidated all of the customer's server data onto one Oracle 6140 SAN array, one Oracle 7320 clustered ZFS storage appliance, and one Oracle 7210 ZFS storage appliance. This solution provided three tiers of storage: SAN for data which requires the highest availability and performance; clustered NFS/CIFS for data which required high availability and access from Linux, Solaris and Windows clients; and non-clustered NFS/CIFS for less critical data which requires access from Linux, Solaris and Windows clients. The design also features a fourth storage unit, an Oracle 7210 ZFS storage appliance, which is located in a remote site. The two primary appliances replicate their data to this appliance for backup and disaster recovery purposes.

Following review with the customer, we then implemented the storage migration onto the new appliances. F3's system engineers performed all migrations over the course of several months, resolving several issues in the process. Data migration was achieved either through direct host-based copying of data from the old SAN arrays to the new storage, or by using the ZFS appliances' Shadow Migration feature which allows incremental copy of data via NFS.

The new storage is well-documented and is well-understood by both the client and by F3's support resources. F3 and the customer maintain a close relationship and we continue to provide ongoing second-level support services for their entire server infrastructure. This results in faster diagnosis and resolution of issues.

## Conclusion

F3 provided both expertise and resources to complete the project. Our ability to design, manage, implement and support the solution enabled the project to be completed efficiently and with minimal disruption to the customer's IT infrastructure. The risk of an outage in the storage infrastructure, which previously could have caused a major disruption to the customer's business, has been eliminated. While improving the company's data backup and disaster recovery systems, the new infrastructure has the added benefit of new features and capabilities which improve the customer's ability to run their business and meet their customers' needs.

# Are you looking for similar results?

## **SPEAK WITH A F3 PROFESSIONAL**

F3 brings IT experts, business flexibility and exceptional service to help our clients strengthen their business.

**[Request an Assessment](#)**