

# ORACLE MAGAZINE

**First Line of Defense** Oracle Database Firewall monitors database conversations /19 **Upgrade to Oracle Business Intelligence Enterprise Edition 11g** Take advantage of the latest OLAP analysis, mapping, KPI, and scorecard features /43 **Add Advanced List Interaction** Use model-driven LOVs in Oracle ADF to implement smart lists /47 **Control the Flow of Execution** Part 2 in a series of articles on understanding and using PL/SQL /50 **On Better Loading and Fact-Checking** Our technologist enjoys the rule of external tables and debunks index scan myths /57

## STORAGE FROM A TO ZFS

How organizations optimize storage with tiered Sun **flash**, **disk**, and **tape** solutions from Oracle and enable unified storage with the Sun ZFS Storage Appliance




### Rise of the Machines

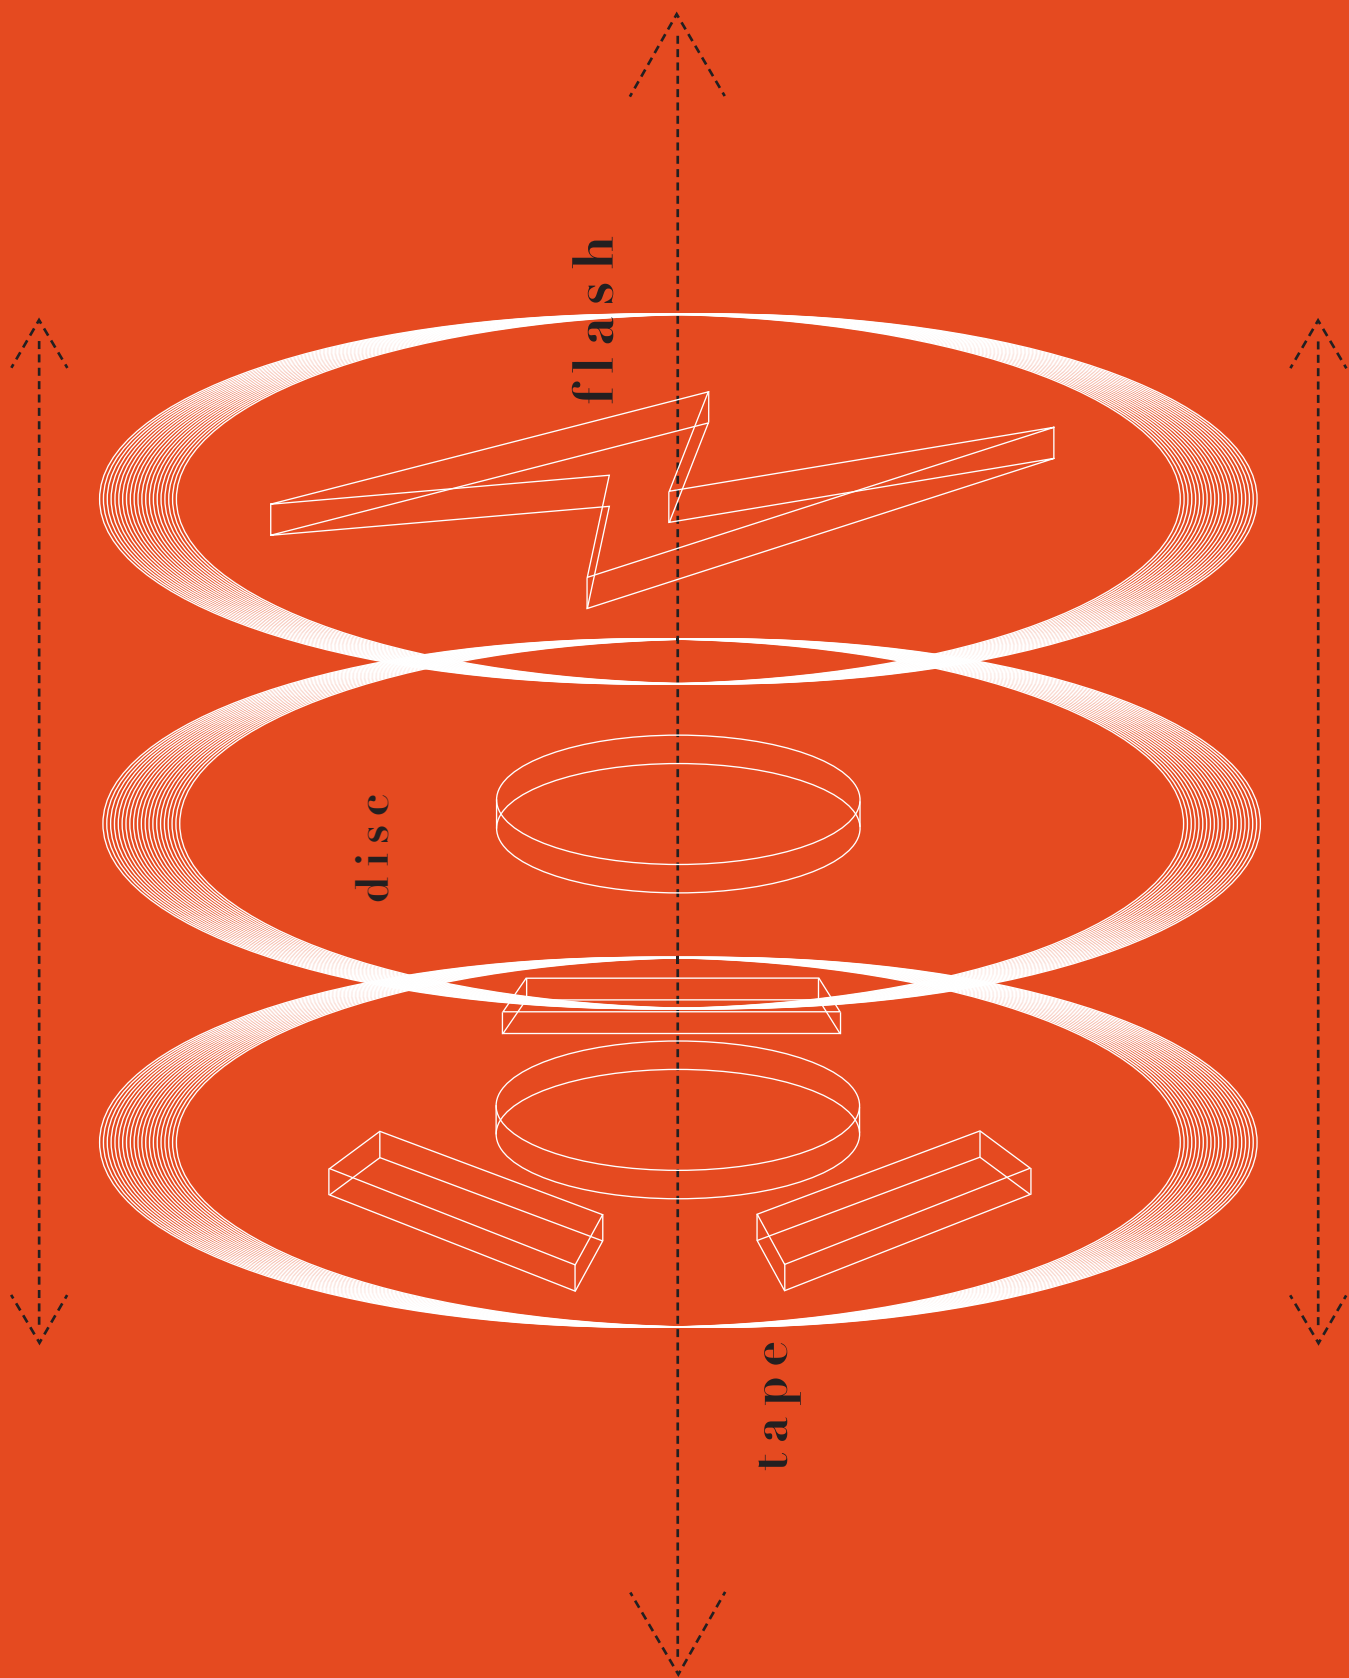
One organization lays a foundation for shared services with Oracle Exalogic and Oracle Exadata

### First Line of Defense

Oracle Database Firewall monitors database traffic and blocks unauthorized activity

SUN ZFS STORAGE 7420 APPLIANCE









# Storage from A to ZFS

How organizations optimize storage with tiered Sun **flash**, **disk**, and **tape** solutions from Oracle and enable unified storage with the Sun ZFS Storage Appliance

Most business users don't spend time thinking about where their files are stored. They log in to the applications they need, create documents, retrieve information, enter transactions, and periodically hit the Save button. They have been told that if they follow a few simple guidelines, all of their critical information will be saved, backed up, and duplicated behind the scenes. The vast storage infrastructure that automates these activities is largely hidden from view.

BY DAVID BAUM

Storage infrastructures grow as the demand for information grows, but strategic storage solutions, from tape libraries to flash and disk arrays to unified storage solutions, can simplify the essential tasks of storing and managing enterprise data. Two organizations, Novant Health and the Australian Bureau of Meteorology, have created business-ready storage systems that exemplify best practices for information management as they efficiently deal with many terabytes of new data every day.

### KEEPING CRITICAL DATA ONLINE

Novant Health is a nonprofit integrated healthcare system that serves communities and patients from Virginia to South Carolina. Like most of today's healthcare providers, Novant needs to provide physicians and administrators with rapid access to clinical and business information, comply with regulations regarding the security of patient information, and reduce the cost of capturing, storing, and delivering an immense volume of structured and unstructured data.

Novant's storage odyssey began nearly 10 years ago, when the IT department created a digital archiving system that enabled more than a dozen hospitals to capture, store, and retrieve images electronically. Led by Jim Grossman, manager of system services, and Robert Dick, storage and UNIX administrator, Novant created a radiographic imaging system that stores files for X-ray, ultrasound, nuclear medicine, magnetic resonance imaging (MRI), and computed tomography (CT) scans. It wasn't long before Grossman and Dick realized that they could broaden this solution for other applications that needed to store and archive patient data. They designed an extensible storage infrastructure around Sun Storage Archive Manager, which now manages 1.4 petabytes of information and supports 38 clinical applications with differing data sets.

"Our storage system was architected around the principle that an application's access to data is top priority," explains Dick. "That means not only being able to safeguard the data in more than one location, but also being able to ensure that it is accessible and survivable through any disaster, so authorized users can always access clinical data."

#### Archive, Access, and Protect with Sun Storage Archive Manager

Oracle's Sun Storage Archive Manager software provides data classification, centralized metadata management, policy-based data placement, protection, migration, long-term retention, and recovery to help organizations effectively manage and utilize data according to business requirements. The software enables users to reduce the cost of storing vast data repositories by providing a powerful, easily managed, cost-effective way to access, retain, and protect business data over its entire lifecycle. This self-protecting file system offers continuous backup and fast recovery features to help enhance productivity and improve resource utilization.

#### SNAPSHOTS

##### Novant Health

[novanthealth.org](http://novanthealth.org)

**Industry:** Healthcare

**Employees:** 25,000

**Revenue:** US\$3.4 billion

**Oracle products:** Sun Storage 6000 series arrays, Oracle's Sun x86 Rackmount servers, Oracle Solaris Cluster, StorageTek T10000 tape drive, StorageTek SL8500 modular library system, Sun Storage Archive Manager, Oracle Database

##### Australian Bureau of Meteorology

[bom.gov.au](http://bom.gov.au)

**Industry:** Government

**Employees:** 1,700

**Oracle products:** Sun technical computing, Sun Storage 6000 series arrays, Sun Storage 6780 array, StorageTek SL8500 modular library system, StorageTek T10000 tape drives, Sun Blade server, Sun x86 Rackmount servers, Sun Storage Archive Manager, Oracle Database

As files are created by Novant's radiographic equipment and enterprise information systems, one copy is immediately written to a high-speed FC Sun Storage 6000 series array from Oracle, and two copies are written to mirrored tape archives in two different datacenters.

"We're backing up data as we create it and archive it," explains Dick. "Data is instantly available to the applications, and within 15 to 30 minutes, we have two backup copies in identical file systems in secure, remote locations. Thanks to this mirrored data storage environment, requests can be virtually redirected from one file system to the other without the user or the application being aware of the difference."

Primary data resides in the FC storage environment for 15 to 30 days, sometimes longer, depending on the service-level agree-

ment determined for each application. After that it is relegated to secondary storage on a Sun Storage 6000 series array. Sun Storage Archive Manager moves data automatically from tier to tier. "Once you set up your storage policies, Sun Storage Archive Manager automates the movement of data from drive to drive and from primary to archival storage," Dick says.

Each of Novant's two datacenters runs identical archiving systems, powered by a pair of clustered Sun servers running Sun Storage Archive Manager. The servers are set as an active/passive pair so that if one of them goes down, it will fail over to the other one. Behind the clustered servers is StorageTek Automated Cartridge System Library Software that manages data storage and retrieval in conjunction with StorageTek T10000 tape drives and a StorageTek SL8500 modular library system divided between the two datacenters.

Thanks to a high degree of automation, Dick is able to maintain the entire archival storage environment by himself (with part-time help from two associate storage managers when he is away). "We bring in, on average, 2.5 terabytes of new data each day," he says. "This system is very stable. In general, it's relatively hands-off."

"The system is designed to maximize security and mitigate risk," adds Grossman. "We're not putting tapes on a truck and physically shipping them off to another site."

Sun Storage Archive Manager ingests 2 million files per week. "We've been running it for more than 10 years, and to date I have no record of ever having lost a single file," says Dick.

### GROWING DATA VOLUMES

Novant's storage solution improves patient care through faster, more-accurate diagnoses while eliminating the risk of lost data. It provides physicians with extremely fast access to patient imaging studies: a 10 MB file can be retrieved from disk and delivered to a physician or technician in less than three seconds.



## “Oracle continues to enhance its storage solutions and renew our faith in the solution set we started with.”

—Robert Dick, Storage and UNIX Administrator, Novant Health



Novant Health's Robert Dick, storage and UNIX administrator (left), and Jim Grossman, manager of system services, rely on an extensible storage infrastructure. "Our storage system was architected around the principle that an application's access to data is top priority," says Dick.

According to Dick, the stress on an archiving system is not necessarily due to access performance or a growing volume of total data stored, but rather to the total number of files that the system has to manage. The higher the number of files, the more difficult it is to restore the data in the event of a disaster. That's a pressing issue for Novant Health as it strives to eliminate hard-copy patient records and securely move more information online.

"We have a new document imaging application for scanning patient records that is expected to add 450 million files into the

system over the next four years," Dick notes, citing just one example. "Because of these stresses, we're moving to solid-state disks, FC, and SATA [serial ATA] drives in a three-tier configuration."

Novant's Sun servers will be replaced by newer gear—Novant is strongly considering a pair of Oracle's SPARC T3-4 servers with 16 processors per core, tied to the existing FC and SATA arrays. "We're also considering adding Oracle's Sun Storage F5100 Flash Array," Dick adds. "Our Oracle storage solution is so flexible that every time we have an archiving issue, we can solve it with the system that's in

place. Oracle continues to enhance its storage solutions and renew our faith in the solution set that we started with.”

Dick believes that while solid-state disk technology is ideal for high-performance data retrieval, it’s not a suitable solution for every part of a storage environment. “If you’re not going to access data for a period of time, it makes much more sense to let it sit at rest in a tape library, where you can get it at a minute’s notice, as opposed to keeping all that data spinning and burning resources—generating heat and using power,” he points out. “Tape technology has not seen its swan song. It is still a very viable technology.”

Benjamin Woo, program vice president for worldwide storage systems at IDC, agrees with this assessment. He says an effective storage strategy only places truly active data on expensive FC or solid-state disk drives and moves less-active data to more-cost-effective storage media—often tape. “It doesn’t make sense to keep data spinning if you’re not going to be looking at it for a few months,” he says. “It should reside in less-expensive media. Our research shows that less than 20 percent of the data in today’s data-centers is actively used. Thus a very small percentage of information needs to reside on very expensive, very fast storage.”

These multilevel storage systems are often called *tiered storage* environments because they assign different stages of data to different tiers of storage media, with the goal of preserving accessibility and reducing cost. “Once you have a tiered infrastructure, the key is to create policies to automate the movement of data so it requires minimal human intervention,” advises Scott Tracy, senior director of flash and disk products at Oracle.

According to Tracy, most tiered storage strategies have three basic components. *Primary storage* stores new data from mission-critical applications and databases, with access time measured in seconds. Storage media typically includes ultrahigh performance flash or solid-state disk (SSD) drives and FC storage arrays. *Secondary storage* stores 30- to 90-day-old data that needs to be kept on hand for business continuity as well as fixed content and backup/recovery data. Access time is measured in minutes, and storage media includes unified storage solutions such as the Sun ZFS Storage Appliance and less-expensive SATA disks. *Long-term or archival storage* stores data older than 90 days, often for historic

## “Oracle’s Sun offerings presented a unified, efficient, large-scale storage system that was flexible enough to grow.”

—Robert Lavery, CIO, Australian Bureau of Meteorology

reasons or for legal compliance. Access time is measured in hours, and storage media includes capacity disks and tape. “Use of tiered storage strategies typically results in around 75 percent reduction in costs compared to single tiers of disk storage,” notes Tracy.

### A MULTITIERED SOLUTION FOR SCIENTIFIC DATA

Tiered storage is what the Australian Bureau of Meteorology (BOM) uses to manage multiple petabytes of scientific data. As Australia’s national weather, climate, and water agency, BOM’s expertise and services assist Australians in dealing with the harsh realities of their natural environment, including droughts, floods, fires, storms, tsunamis, and tropical cyclones. Through regular forecasts, warnings, monitoring, and advice spanning the Australian region and Antarctic territory, BOM provides one of the most fundamental and widely used services of government. The agency also supplies specialized forecasting to the aviation industry, oil rigs, fire departments, police forces, and other emergency services, in addition to bulk data uploads for universities and scientific organizations.

BOM’s data storage environment was put to the test earlier this year during Cyclone Yasi, the largest tropical storm to strike Australia since Europeans first settled there in the 18th century. With winds of 186 mph and waves on the coast reaching as high as 30 feet, Yasi wrecked more than 10,000 homes and businesses but did not cause any fatalities, thanks in part to BOM’s precise forecasts and reliable information-management infrastructure.

BOM’s Oracle-based data storage environment gathers about 40 TB of data each week from 6,000 meteorological devices including warning buoys, oceanographic sensing gear, flood hydrology sta-

### Unified Storage with the Sun ZFS Storage Appliance

Oracle’s Sun ZFS Storage Appliance product line delivers unified storage for

- Consolidation of multiple storage devices
  - Secondary storage for data protection (when multiple copies of data are required)
  - Test/development environments in which organizations are optimizing copies of primary data
- The product line enables the rapid deployment of new revenue-producing applications and lowers expenses by reducing storage complexity and its associated administrative costs.

Sun ZFS Storage Appliances consolidate files and block I/O on a single high-capacity, high-performance storage appliance. This consolida-

tion supports 10 different protocols across three interconnects—Ethernet, FC, and InfiniBand. Sun ZFS Storage Appliances are built on the Oracle Solaris operating system, and they offer data services including deduplication, compression, replication, snapshots, and clones.

Key features in the Sun ZFS Storage Appliance include storage analytics and hybrid storage pools. Storage analytics provide real-time comprehensive analysis of storage workloads that assist in capacity planning and provide detailed insight into storage networks. Hybrid storage pools combine DRAM, read-optimized flash devices, write-optimized flash devices, and native

SAS high-density, low-power spindles in a single transparently managed caching hierarchy. Data is automatically migrated from flash to disk based on access patterns and requires no human intervention. This “autotiering” keeps the most-used data in fast-access storage and the least-used data in inexpensive high-density, low-power spindles, reducing overall TCO.

The Sun ZFS Storage Appliance product line combines cloud-ready software and hardware, designed to enable customers to start small, deploy applications faster and at a lower cost, and grow into a next-generation cloud-computing infrastructure.





**Robert Lavery, CIO at the Australian Bureau of Meteorology, says Oracle's large-scale storage offerings provide efficiency and the flexibility to grow with the organization's rapidly evolving data processing workload.**

tions, weather balloons, satellites, and aircraft. After being digitized, cleansed, transformed, and stored in an Oracle database, the data is processed by nearly 300 applications—some of which can predict the path and trajectory of cyclones days in advance.

In 2008, Sun systems were selected to provide the computing and storage facilities for this massive operation, including service, support, and regular equipment upgrades over a five-year period.

"We wanted a one-stop shop—one vendor that could provide a complete solution," states Robert Lavery, chief information officer at the Bureau of Meteorology. "We didn't want to have disparate issues with vendor finger-pointing. Oracle's Sun offerings presented a unified, efficient, large-scale storage system that was flexible enough to grow with our rapidly evolving data processing workload."

Today BOM's core meteorological atmospheric modeling prediction services are housed on a Sun technical computing solution and the output is archived on a large-scale data storage system. "About 10 percent of the data is corporate structured data, and 90 percent is unstructured scientific data such as satellite images," says John de la Lande, the high-performance computing, storage, and facilities manager for BOM's IT environment.

The technical computing system shares resources among many nodes as it processes thousands of inputs from meteorology stations. It stores the output in four tiers. Tier 1, also called *main storage*, is mission-critical storage, with 99.999 percent availability. For this tier BOM uses a primary storage system from Sun. Tier 2, also called *deep storage*, utilizes Sun Storage 6000 series arrays with FC disk drives. Tier 3, also called *bulk storage*, utilizes Sun Storage 6000

series arrays with SATA drives. Tier 4 maintains archival data on a StorageTek SL8500 modular library system in conjunction with a StorageTek T10000 tape drive.

A Sun Blade server and Sun x86 Rackmount servers from Oracle process the data under the direction of Supervisor Monitor Scheduler. The output is then passed to Sun Storage Archive Manager, which automatically moves data from tier to tier according to previously defined storage policies. According to Richard Oxbrow, storage manager at BOM, Sun Storage Archive Manager automates the entire information lifecycle, from fast disk to tape. It includes continuous backup and fast recovery features to enhance productivity and improve resource utilization. Thanks to a high degree of automation resident in the Oracle-based systems, two or three people maintain the entire storage infrastructure.

### EVOLVING TO UNIFIED AND FLASH STORAGE

BOM's storage infrastructure is constantly evolving as Lavery and his team add new hardware and software to improve speed and capacity. For example, the Sun Blade server features solid-state drives (SSDs) that deliver I/O performance of up to 400 hard disk drives. BOM recently purchased two of Oracle's Sun ZFS Storage 7420 appliances, which use hybrid storage pools with flash-based caches to dramatically improve application response times.

Oracle's Tracy says these unified storage solutions are ideal for organizations that wish to consolidate and virtualize their storage infrastructure, as well as for companies adopting cloud computing. Oracle's unified storage product line, anchored by the Sun ZFS Storage series, can transparently manage where data is placed in a multitiered storage environment with hybrid storage pool technology, holding copies of frequently used data in fast SSDs while storing less frequently used data in less-expensive, high-capacity SAS disks. (For more information on Oracle's Sun ZFS Storage product line, see "Unified Storage with the Sun ZFS Storage Appliance" on page 34.)

As BOM evolves its information systems, Lavery believes the move to engineered systems, such as storage appliances, is right for the organization. "We are extending our key vendor relationships to partnerships," he says, "because we believe that our services are critical enough and important enough to merit this level of commitment." ◀

---

**David Baum** ([david@dbaumcomm.com](mailto:david@dbaumcomm.com)) is a freelance business writer based in Santa Barbara, California.

### NEXT STEPS

**LEARN more about**  
**Sun storage and tape**  
[oracle.com/storage](http://oracle.com/storage)

**Sun ZFS Storage Appliances**  
[bit.ly/9WV4rH](http://bit.ly/9WV4rH)

**Sun Storage Archive Manager**  
[bit.ly/kNWeae](http://bit.ly/kNWeae)