Technology Insight Paper

Oracle's Engineered Systems Approach to Maximizing Database Protection

By John Webster

January, 2016

Evaluator Group

Enabling you to make the best technology decisions

Data Protection—an Ever-present Requirement

Enterprise IT is transitioning to new application delivery models. The exact nature of those models will be different for every organization and will consist of a blend of commodity infrastructure, converged and hyper-converged application delivery platforms, and cloud services to name only three. However, no matter how much of a transformation IT architects plan to make using what they believe they need as the right blend of technologies and management practices, a common thread will run through all transformation projects—the ever present need to protect and secure data from loss due to such things as administrative error, component and media failures, and infrastructure gaps.

In fact, we believe that the need for data protection and security becomes even more important as transformational projects move forward. While the value of infrastructure ownership is decreasing, as evidenced by the growing use of cloud services by enterprise IT, the value of data ownership is growing. And not only is the value of data ownership increasing, the amount of valuable data is increasing as well as is evidenced by the growth of Big Data applications now targeting predictive marketing and financial risk mitigation.

Protecting against data loss has always been a core enterprise IT function that cannot ultimately be outsourced. Data security and protection remain the sole responsibility of the enterprise. Transformation will not change that. However, new platforms, new applications and application delivery models, and the now constant influx of new technologies will make IT's response to data protection and availability mandates even more challenging.

Here we review those challenges and offer a review of a data protection and recovery solution built specifically for Oracle Databases —Oracle's Zero Data Loss Recovery Appliance (Recovery Appliance).

Not a "One-Size-Fits-All" Challenge

Enterprise database and storage administrators are now responding to a number of divergent technology trends that are complicating their ability to assure data consistency, availability and security:

> The increasing use of all flash and hybrid flash/disk arrays for primary storage while still supporting spinning disk arrays for secondary storage applications;
>
> Hyperscale and scale-out data center architectures that senior IT management now wants to leverage. These include the growing use of software-defined and open source storage software solutions;
>
> The increasing need to store and protect business-critical transactional data and as well as large volumes of unstructured data for Big Data analytics applications.

Normal data growth already challenges IT administrators to assure data access and availability while controlling data protection costs related to hardware, software, and management processes. Sub-second recovery point e demands are increasing as well across a broader range of applications.

This brings into question the continued reliance on general purpose data protection solutions that are used for most if not all applications generating both structured and unstructured data. Without doubt, there are good economic reasons for consolidating backup and recovery processes. However, there are trade-offs to consider. A one-size-fits-all strategy adds complexity when administrators need to give backup and recovery performance priority to the most critical applications. In addition, what works well for large volumes of unstructured data may leave data protection gaps when dealing with transaction-oriented databases. Finally, the ability to automate data protection processes and delegate backups and recoveries—an ability that becomes of increasing importance in more diverse application environments—becomes more complex as well.

Oracle's approach is simply this: If you believe that applications founded on the Oracle Database are critical to ongoing business operations, then address those with a data protection solution that was developed and optimized by Oracle's own engineers rather than using general purpose backup hardware and software. The benefits of doing so include better backup and recovery performance to meet stringent RPO and RTO SLAs, reduced data loss exposure, increased data integrity, and greater application user satisfaction.

Oracle's Zero Data Loss Recovery Appliance

Oracle's Recovery Appliance was designed for database recovery versus the traditional backup centric approach and was co-engineered with Oracle Database. As such, it leverages the Oracle Engineered System's architecture—a group of systems that also includes Exadata and Exalytics. Engineered Systems—with thousands of deployments worldwide—implement specially designed, advanced API calls between hardware and software as part of Oracle's cross-portfolio strategy to converge the two. Indeed, the Recovery Appliance is built on the Oracle Exadata architecture. The design also reflects what Oracle calls its Maximum Availability Architecture that is represented in other Oracle Engineered Systems—a "database-aware" high availability architecture with end-to-end Oracle validation capabilities, active clustering, and no single point of failure.

For Oracle Database, the Recovery Appliance was engineered to go beyond the capabilities of the general purpose backup appliances available today that are often limited in scale and performance by single-controller architectures. First, it is scalable to petabytes of usable capacity to accommodate a growing number of protected databases, backup traffic, and storage usage by adding compute and storage servers to the single appliance configuration. The Base Rack includes two compute servers and three storage servers providing up to 94 TB usable capacity for backups. It can be upgraded on a per storage server basis up to a maximum of 18 storage servers per rack with 580 TB usable capacity. When additional compute servers are required beyond the Base Rack, additional racks are connected via InfiniBand. Up to 18 fully configured racks can be connected together, providing up to 10 PB of usable

capacity in a single system image (see Figure 1 below). Oracle claims an effective (deduped) capacity of up to 5.8 PB of virtual full backups per rack with a 120 TB/hr. virtual backup rate (12 TB/hr. sustained delta ingest) and a maximum of 100 PB of virtual full backups, with a 2 PB/hr. virtual backup rate (216 TB/hr. delta ingest rate). For restorations, Oracle claims that a full rack can restore at the rate of up to 12 TB/hr., and 18 fully configured racks can restore at the rate of up to 216 TB/hr. Ultimately, by scaling both compute and capacity simultaneously, the Recovery Appliance can meet the demands of large-scale enterprises to protect growing data stores with consistent performance expectations.

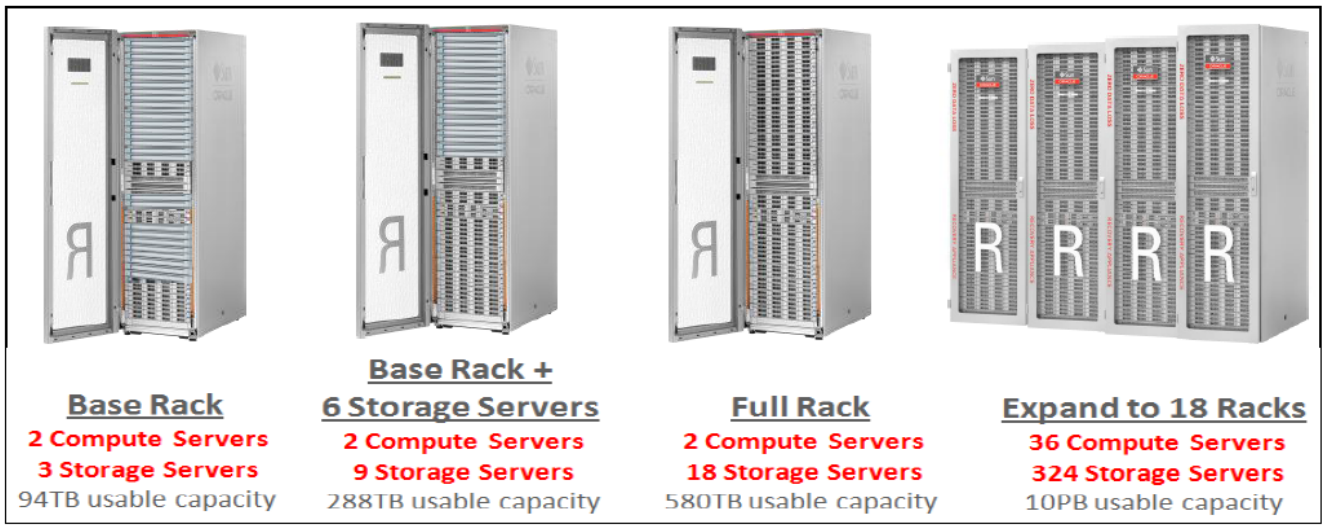

Figure 1. Oracle Recovery Appliance Configurations (Source: Oracle)

Second, it incorporates two key features: Delta Store and Delta Push (described in more detail below), named as such because only changed blocks (deltas) are sent to and stored on the Recovery Appliance. This strategy reduces additional overhead on the database server and enables the most efficient source-side data reduction technique available today for Oracle Database protection.

Eliminating Data Loss—Delta Push and Real-time Redo Log Technology

Delta Push consists of two processes that are run on each protected database: RMAN incremental backups and Real-time Redo Transport. Delta Push **incremental backups** eliminate the need to run weekly or daily full backups of a database. (See Figure 2 below). After an initial RMAN full backup (level 0), only the changed blocks are backed up using standard RMAN processes resulting in an "incremental forever" backup strategy. Backup schedules are defined by the administrator and can be set at daily or shorter time intervals. By sending only the changed blocks and eliminating the need for ongoing full backups, Delta Push virtually eliminates backup window constraints and maximizes resource efficiency.

To protect against more catastrophic data loss exposure such as the loss of a production system and storage, **Real-time Redo Transport** is used to reduce a database's Recovery Point Objective (RPO) to sub-second time intervals. The Recovery Appliance receives redo blocks directly from the database's in-

memory buffers and writes to a pre-configured redo log staging area on the appliance. The Recovery Appliance automatically compresses and creates an archived log backup written to the Delta Store. If a disruption in connectivity occurs, the Recovery Appliance automatically "closes" the redo log and creates a partial archived log backup to assure recoverability up to the point of the last change received by the appliance. Once the connection between the protected database server and the appliance is restored, a gap detection process on the appliance automatically retrieves missing archived logs from the protected databases. This technology protects transactions in real-time, offering zero to sub-second recovery point objectives.

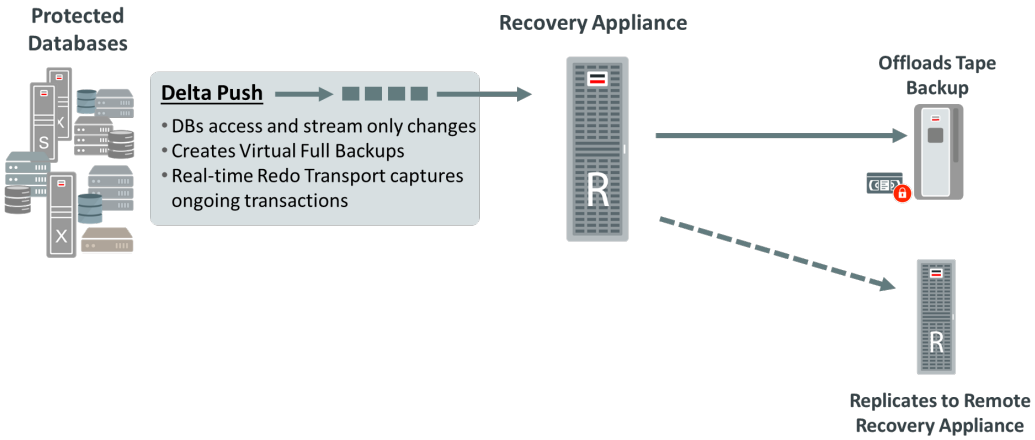

Figure 2. The Oracle Recovery Appliance Data Protection Environment showing Delta Push example. (Source: Oracle)

Point-in-time Recovery using Virtual Full Backups

Virtual full backups offer the same advantages as restoring from a full backup but without the processing overhead or need to apply incremental changes. Backups are immediately validated, compressed, indexed and written to the Recovery Appliance's Delta Store. To accomplish a point-in-time restore, the Recovery Appliance creates a virtual full backup from the current and previous backup blocks that have been received... Administrators can then restore a subset of the database such as a data file or the entire database from the latest virtual full or from a previous virtual full backup (Day N full backup) as is shown in Figure 3 below. Delta Store technology ensures database-aware validation of backup data for data integrity. It also provides the ability to recover full databases to any-point-in-time without the resources required to regularly create full database copies. In addition, backup data can be replicated to an offsite appliance (see Replication discussion below) and/or offloaded to tape or the Oracle Public Cloud for an additional level of protection.

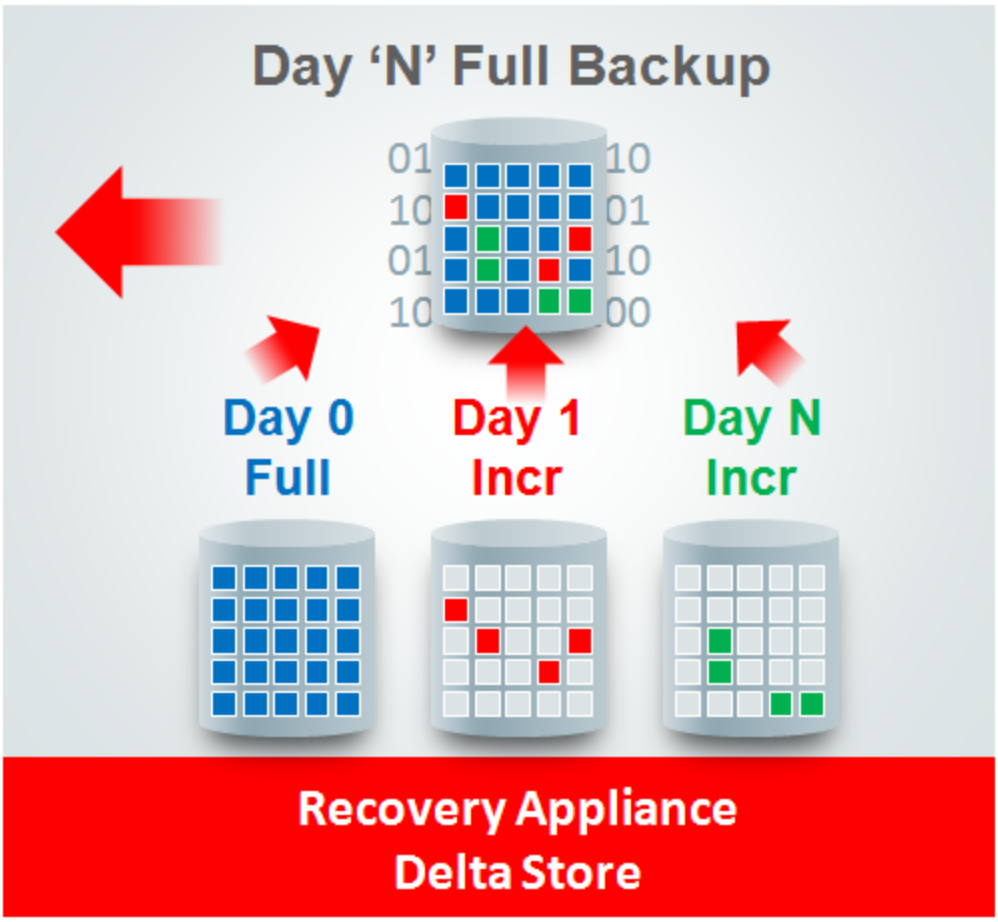

Figure 3. Recovery Appliance Delta Store Virtual Full Backups (Source: Oracle)

Restoration and recovery processes are fully integrated with Oracle RMAN. When an RMAN restore and recovery request is issued, the most appropriate virtual full backup is identified. The Recovery Appliance then constructs the physical backup sets from the Delta Store and sends the backup sets to the RMAN client along with the archived log backups needed to restore data files to the desired point-in-time. With Real-time Redo Transport enabled, the most current archived logs plus any partial archived logs will be available to the RMAN client allowing the database to be recovered with the most current version of data at the point in time required.

Automated Data Protection

Automation of data protection processes improves the efficiency of storage administrators and reduces human error particularly when managing hundreds of databases. To help minimize manual tasks, automation of data protection processes is exploited with the Recovery Appliance in a number of different ways. These include:

- Automatic purging of backups on the Recovery Appliance based on user-defined recovery window goals
- Dynamic allocation of space to meet recovery window goals,
- Backup of archived logs on the appliance when Real-time Redo is enabled,
- Creation of virtual full backups during restore operations
- Tape and/or replication of backups when defined
- Monitoring the backup environment and alerting administrators as needed

Real-time recovery status for all databases under management is automatically displayed via a Recovery Appliance Enterprise Manager plug-in. In addition, out-of-the-box reports by database and/or policy can be automatically created and distributed (e.g. email) based on a user-defined schedule.

The Recovery Appliance is also integrated with Oracle Enterprise Manager Cloud Control which gives an administrator the ability to view and control these automated functions from a centralized management GUI. Integration with Oracle Enterprise Manager Cloud Control via an Enterprise Manager Recovery Appliance plug-in offers administrators an end-to-end view of the backup lifecycle managed by the Recovery Appliance from the initiation of an RMAN backup to final disposition to disk, tape, or to a secondary Recovery Appliance. Administrators can view backup volume/performance and aggregate or per database capacity consumption. Alerts can also be sent when goals determined by policy are not being met such as impending capacity limitations that are threatening the ability to recover to any point-in-time within the stated recovery window goal.

Data Protection-as-a-Service (DPaaS)

The Recovery Appliance's policy-based management environment enables storage administrators to expand data protection process management beyond the storage environment to qualified database administrators and business user groups. An end-to-end view, from RMAN to disk, enables database or storage administrators to deliver policy-based Data Protection-as-a-Service. DPaaS can be instantiated internally within an enterprise to multiple LOBs, or at a cloud service provider offering DPaaS to a global base of external customers. DPaaS policy formulation can be done in collaboration with business units and DBAs while storage administrators implement policy. DBAs have the ability to implement policies as well.

Policy service levels—gold, silver, bronze for example—can be created for groups of databases. Policies are automatically applied to a database when it is added to the group. Defined under policy groups are processes such as recovery window settings for disk and tape, maximum on-disk backup retention

periods, unprotected data window threshold, and guaranteed copy requirement. A real time 'Current Status' dashboard visible to those with access rights, allows users to see and further explore:

- Summary of current and recent activity
- Protected database backup/recovery issues
- Daily data sent / received
- Performance
- Media Manager status
- Replication status
- Storage Location status
- Incidents and Events

Pro-active capacity planning reports can be generated from analysis of the Recovery Appliance's monitoring of disk capacity usage and data growth rates.

Dynamic Space Allocation and Meeting Recovery Window Goals

When a database is added to the Recovery Appliance, it is associated with a protection policy that includes a Recovery Window Goal and given a "reserved space" storage capacity equal to the amount the database would be expected to consume based on its pre-defined recovery window goal. At any given time, a database may be consuming more or less than its reserved space because the Recovery Appliance dynamically allocates space to best meet recovery window goals for all databases under management. However, it is guaranteed to always have its reserved space amount of storage. The Recovery Appliance uses the recovery window goals and reserved space settings to proactively purge blocks of data outside of recovery window goals to free up space for new backups.

Replication for Disaster Recovery and Business Continuance

For disaster recovery and business continuance, different modes of automatic data replication are supported between Recovery Appliances at different geographic locations. Traditional one-way, bi-directional, and multi-site hub and spoke topologies are supported with no limit on distance between sites (see Figure 4 below). As with the Delta Push technology, only the changed data blocks (incremental and archived log backups after the initial full) are replicated. Once they are received at the downstream site, backups are compressed, validated and indexed to mirror the source and the catalogs between the source and target are reconciled. Backups at the downstream site may have the same or different recovery settings as that of the source Recovery Appliance. If a virtual full is requested that cannot be satisfied at the source site, it is automatically retrieved from the downstream —a process that is transparent to DBAs, for example.

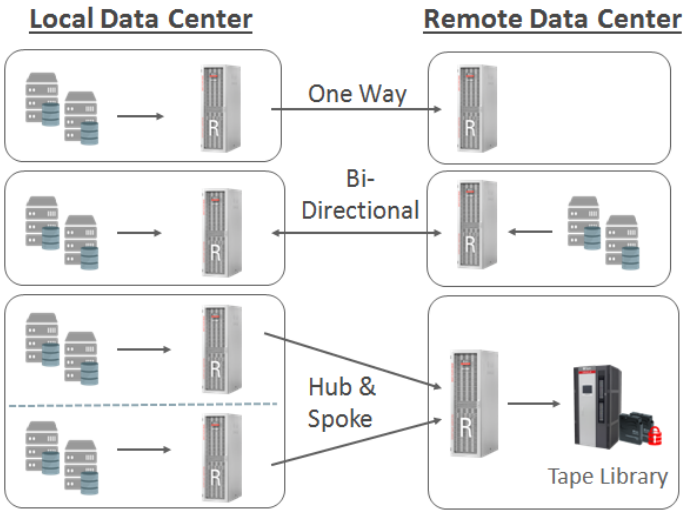

Figure 4. Recovery Appliance Backup Data Replication Topologies (Source: Oracle)

Oracle Recovery Appliance User Experience

At the most recent Oracle OpenWorld event, a Recovery Appliance user reported the results of a side-by-side comparison with a well-known, single controller, general purpose backup appliance. This public sector customer is a civil and criminal court system in Brazil with 13,000 employees serving 7 million residents. All judicial processes are recorded in digital format and 24x7 access to them is a mandatory requirement. Their IT department currently supports 40 Oracle databases used by 8,000 internal users. One of their critical databases also sees an average of 70,000 web accesses per day.

Using a popular general purpose backup appliance, the customer experienced a nearly three-day time requirement to backup and recover their most critical 17 TB Oracle database. They experienced what they considered to be low deduplication rates and needed to expire old backups on a daily basis due to space constraints. They also reported that it was difficult to meet what they considered to be adequate RPO and RTO objectives due to network and hardware constraints. Weekly full backups of one 17TB Oracle Exadata system took 72 hours over a 1Gbps network. Offload of backups to an Oracle StorageTek tape silo was also a weekly requirement.

Using Oracle's Recovery Appliance, they were able to accomplish the initial full backup of the same Exadata system in 12 hours over a 10Gbps network and could complete incremental backups in 23 minutes. Backups across their Oracle Database environment that once averaged 90 hours to complete can now be achieved in approximately 20 minutes with a 12:1 deduplication rate. The customer credits these results to their ability to do "incremental forever" and virtual full backups with the Recovery Appliance. They also reported that their investment in two Recovery Appliances cost less than what they would be required to spend upgrading and expending their current general purpose backup appliances to support the same requirements.

Evaluator Group Assessment:

When data is lost and can't be recovered, it's gone forever. As the growing value of data ownership is now understood by a broad spectrum of business and organizational executives, data protection and recoverabilty have become an increasingly crucial consideration. Losing data equates to losing value in ways that are now expressed as lost revenue, lost business opportunity, lost productivity, and increased risk exposure.

Oracle's Recovery Appliance, co-engineered with Oracle Database and Oracle Recovery Manager, offers IT administrators a single, automated, Oracle-optimized and recovery-focused environment with sub-second RPO for data protection infrastructure and support. Its scale-out architecture puts it in a different class of backup appliance by enabling Oracle DBAs and application administrarors to expand data protection capabilities continuously without adding stand-alone controllers that create islands of backup data and make backup and recovery more complex and time consuming. As one user has reported, an investment in the Recovery Appliance may actually cost less than it would have to add to their existing general purpose backup appliances while achieving 200X improvement in backup performance.

All things considered, Oracle's Recovery Appliance can add significant value for IT organizations tasked with maintaining the continuous avaliabilty of critical Oracle Databases by providing a platform for data protection-as-a-service. They can now offer different service levels according to an application's business criticality and monitor in real-time whether or not they can meet the SLAs. In addition, we note that the Recovery Appliance can use Oracle's StorageTek tape storage or Oracle Public Cloud as additional layers of data protection. This compares favorably with other data protection appliances that have yet to offer this capability.

As such, Oracle's Recovery Appliance is unique among competing general-purpose backup appliances due to its co-engineered capabilities with Oracle Database and the Oracle stack. Indeed, Oracle's Converged Infrastructure Architecture provides customers with a clear view into its future strategy and direction.

About Evaluator Group

*Evaluator Group Inc. is dedicated to helping **IT professionals** and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver **in-depth, unbiased analysis** on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and*

vendor professionals through product and market evaluations, competitive analysis and ***education****.* ***www.evaluatorgroup.com*** *Follow us on Twitter @evaluator_group*