

ORACLE®

Zero Data Loss Recovery Appliance Deep Dive: Direct from Development

Timothy Chien
Principal Product Manager
Oracle

Raymond Guzman
Consultant Member of Technical Staff
Oracle

Fernando Simon
DBA Manager
Brazilian Justice Tribunal of Santa Catarina

ORACLE

**ORACLE
OPEN
WORLD**

October 25–29, 2015
San Francisco

Agenda

- 1 ➤ Today's Backup & Recovery Challenges
- 2 ➤ Recovery Appliance Architecture
- 3 ➤ Virtual Full Backup & Real-Time Redo Transport
- 4 ➤ Tape Backup & Replication
- 5 ➤ Managing Time and Space
- 6 ➤ Customer Case Studies
- 7 ➤ Summary / Q&A

Agenda

- 1 Today's Backup & Recovery Challenges
- 2 Recovery Appliance Architecture
- 3 Virtual Full Backup & Real-Time Redo Transport
- 4 Tape Backup & Replication
- 5 Managing Time and Space
- 6 Customer Case Studies
- 7 Summary / Q&A

All-Purpose Storage HW & SW Not Designed for Database

Treat Databases as Just Files to Periodically Copy




Data Loss Exposure

Lose all data since last backup



Daily Backup Window

Large performance impact on production



Poor Database Recoverability

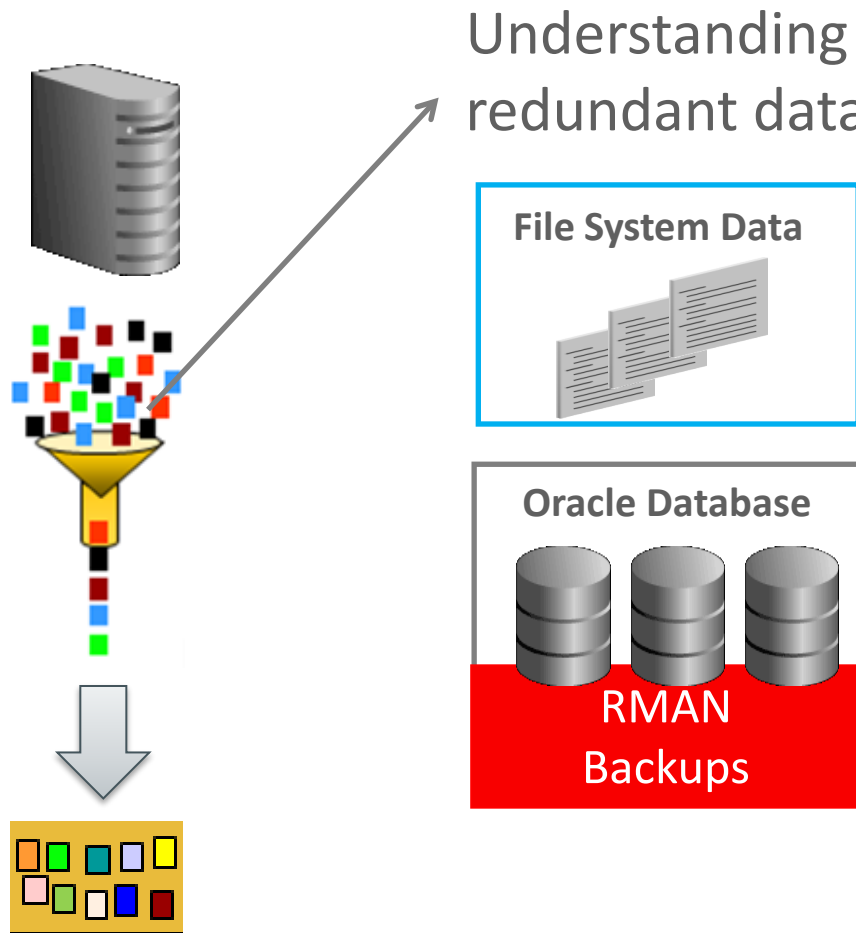
Many files are copied but protection state of database is unknown



Many Systems to Manage

Scale by deploying more backup appliances

Oracle Deduplication & Validation Challenges




Understanding the backup stream is critical to identifying redundant data

- 1. File system data** format does not change in the backup process
 - ✓ Ideal for storage deduplication technologies
 - ✓ Deduplication claims of **10 – 50x**
- 2. RMAN** passes a pre-packaged backup set
 - Backup stream / block format is **Oracle proprietary**
 - RMAN backup stream is largely opaque to external backup applications and storage
 - Full backups (production overhead)** needed for dedupe
 - Minimal deduplication (< 6x)** vs. file system data
- 3. Storage technologies** cannot validate Oracle blocks nor interpret transactional change data

Fundamentally Different Approach to Protect Business Critical Database Data

Zero Data Loss Recovery Appliance



Recovery Appliance Unique Benefits for Business and I.T.




Eliminate Data Loss

Real-time redo transport provides instant protection of ongoing transactions




Minimal Impact Backups

Production databases only send changes. All backup and tape processing offloaded



Database Level Recoverability

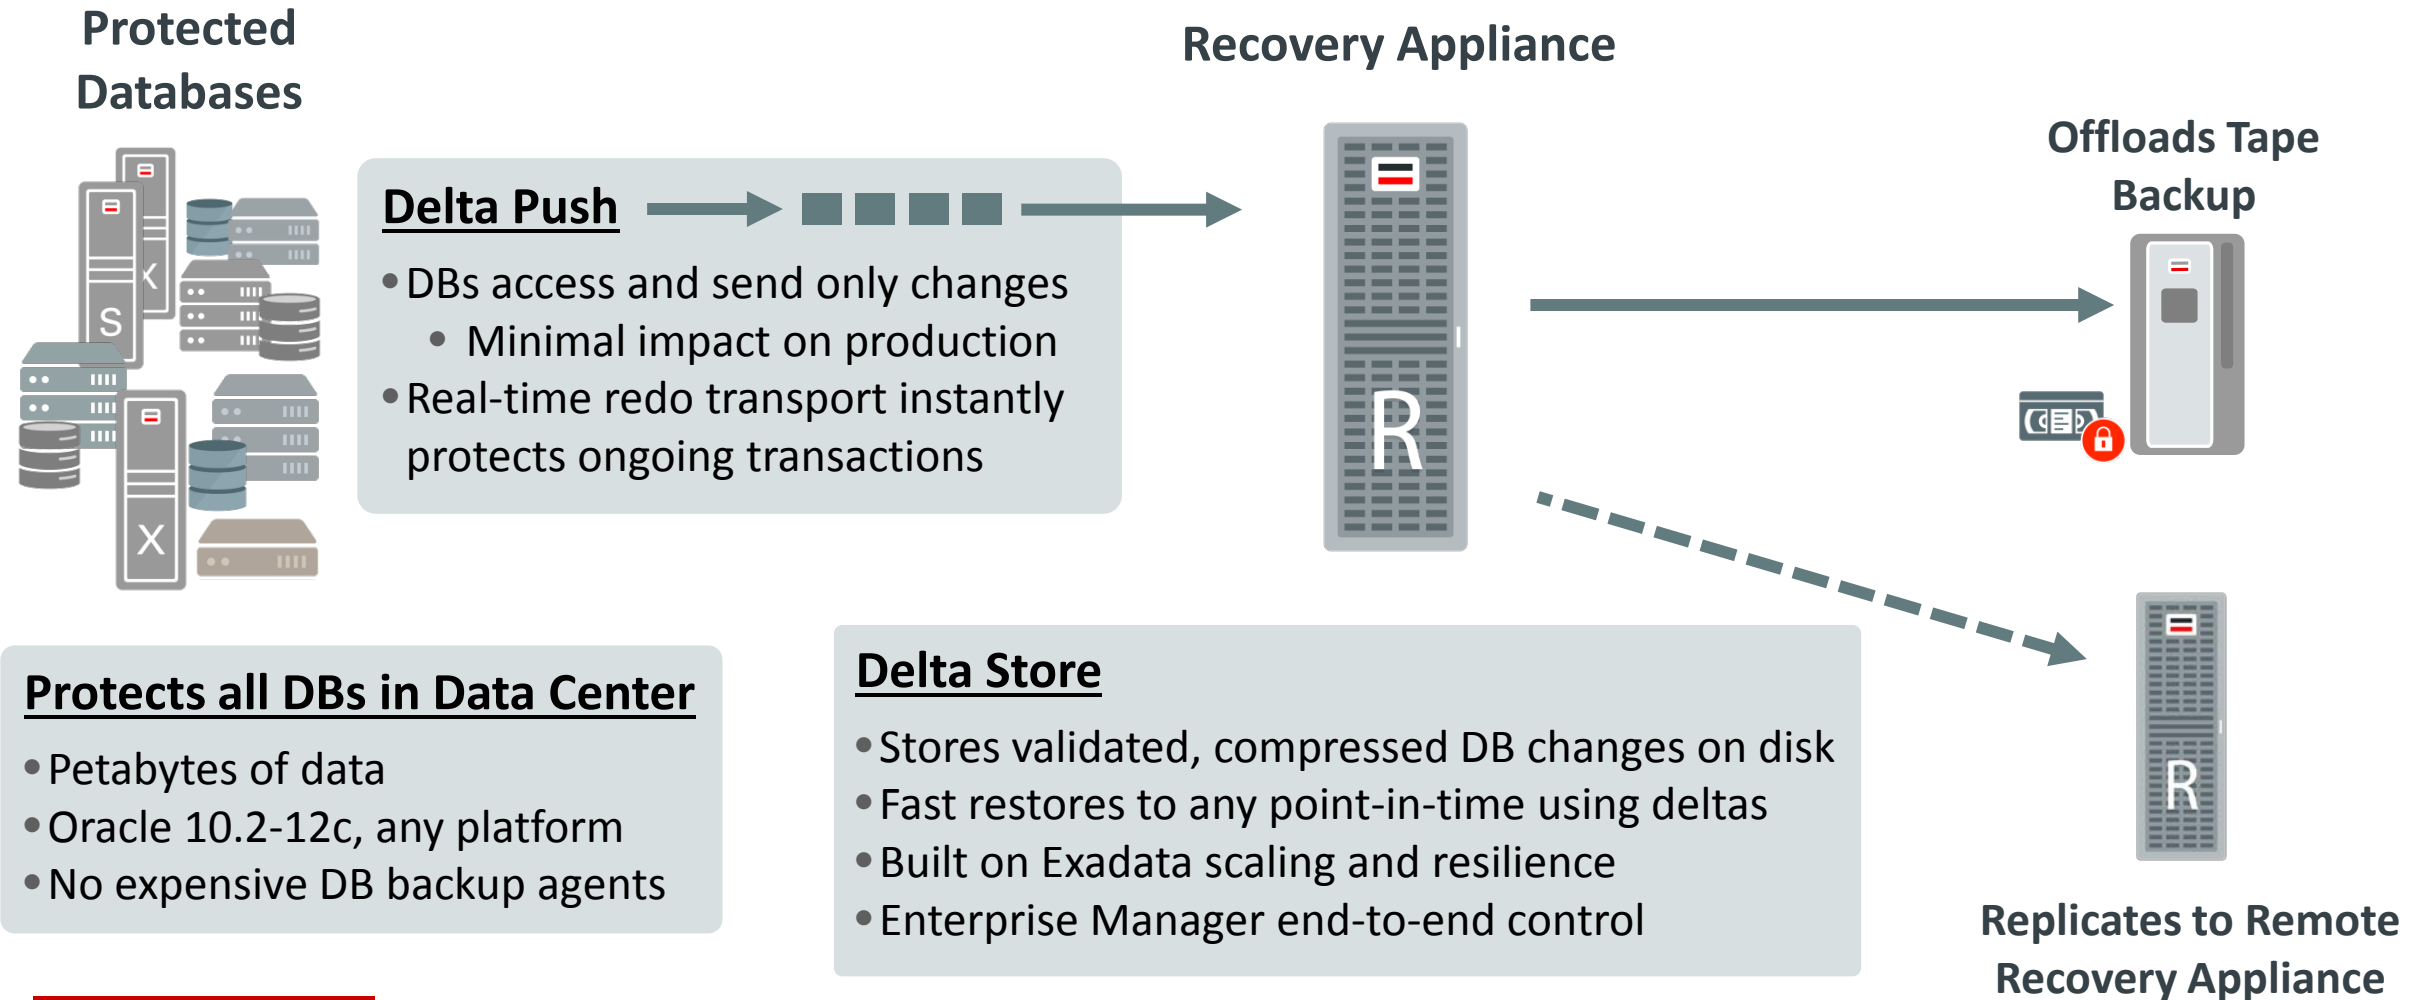
End-to-end reliability, visibility, and control of databases - not disjoint files



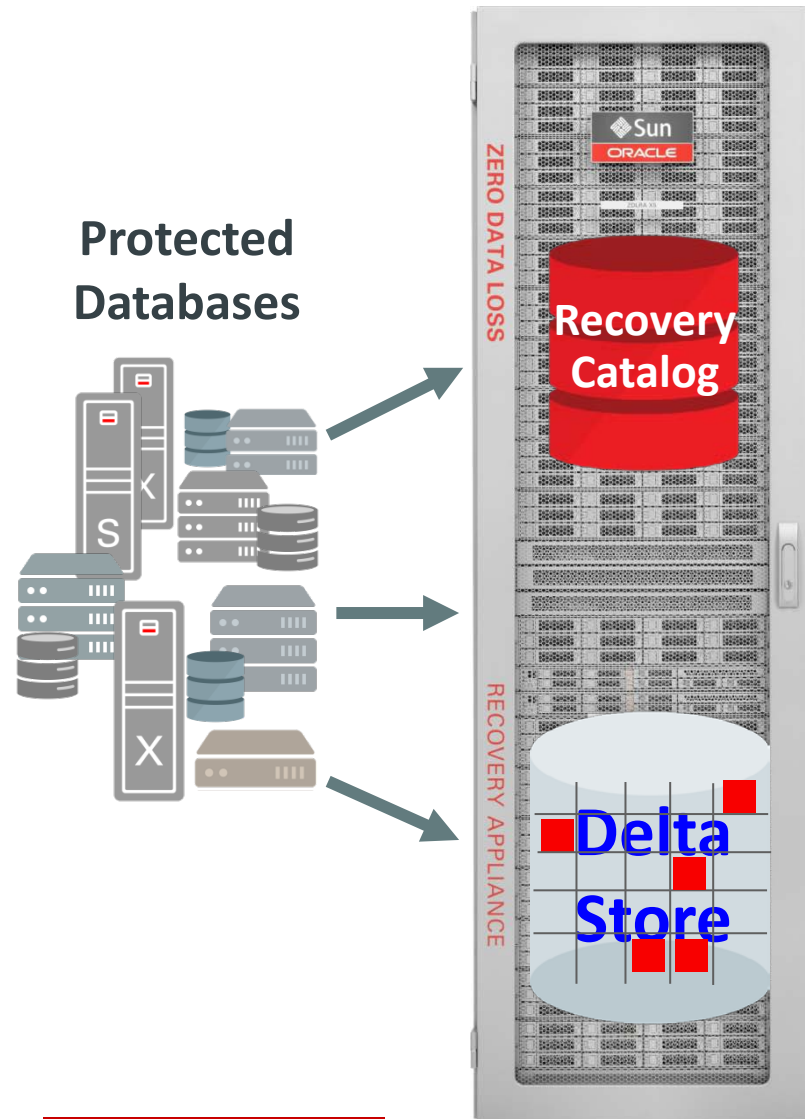
Cloud-Scale Protection

Easily protect all databases in the data center using massively scalable service

Zero Data Loss Recovery Appliance Overview



Key Architecture Components



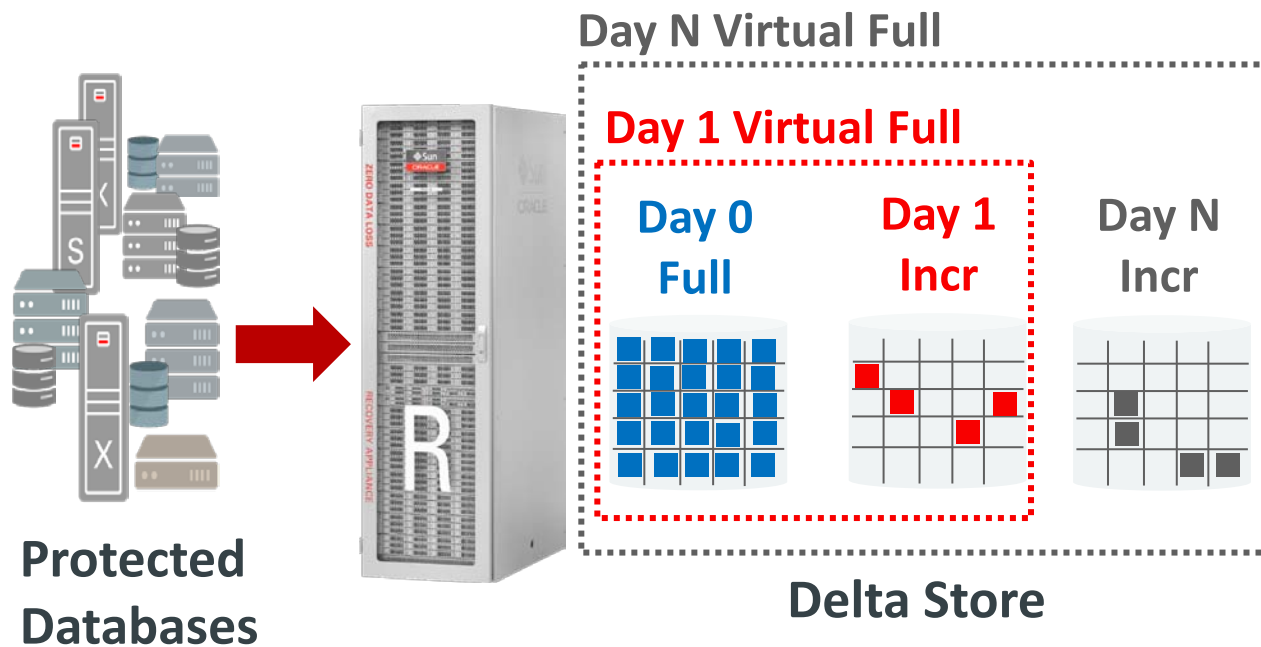
- Based on Exadata X5-2 HC with embedded 2-node RAC+ASM database
 - Serves as central RMAN Recovery Catalog, storing all backup metadata
- Delta Store 'backup data' configured in separate ASM disk group, stored using compressed format
- Pre-bundled Oracle Secure Backup tape software
 - 16Gb QLogic Fiber Cards supported for tape connectivity
- RMAN + Recovery Appliance Backup Module on protected databases
- EM Cloud Control 12c Release 4 + Recovery Appliance plug-in

Agenda

- 1 Today's Backup & Recovery Challenges
- 2 Recovery Appliance Architecture
- 3 Virtual Full Backup & Real-Time Redo Transport**
- 4 Tape Backup & Replication
- 5 Managing Time and Space
- 6 Customer Case Studies
- 7 Summary / Q&A

Space-Efficient “Virtual” Full Backups

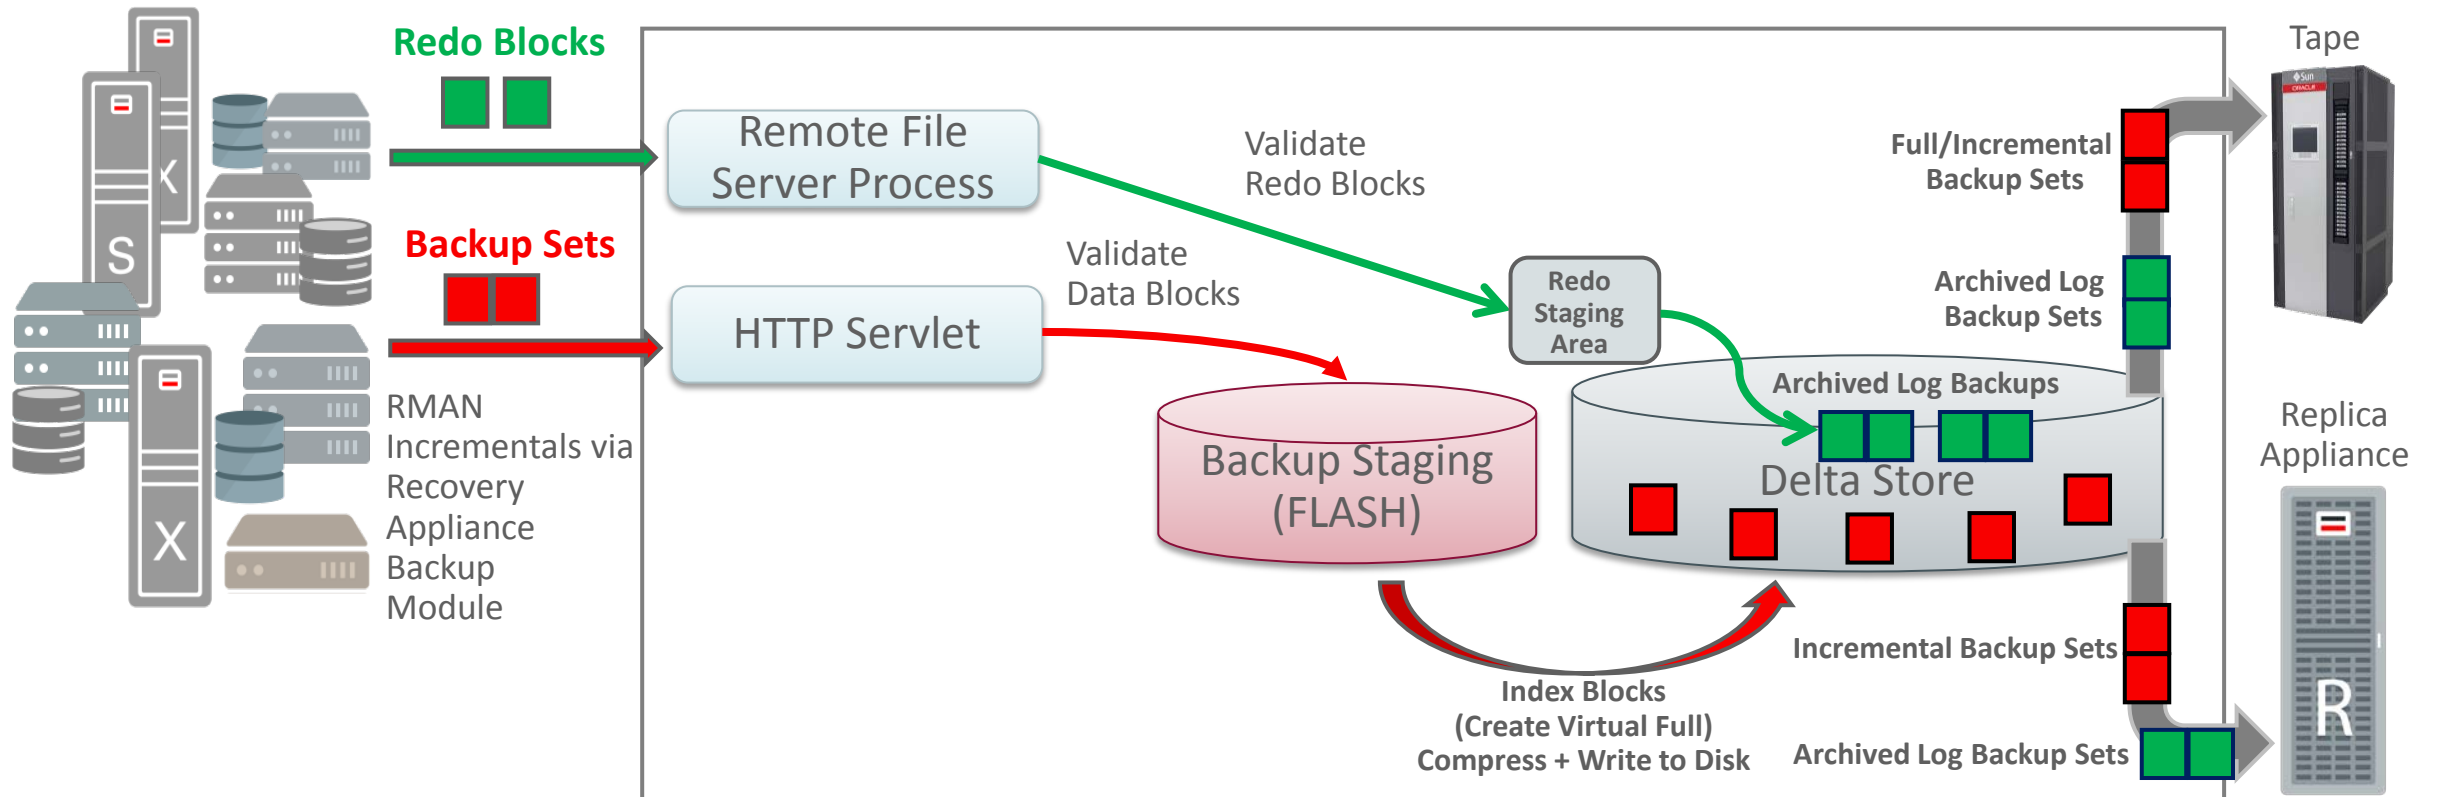
No More Full Backups: Incremental Forever Architecture



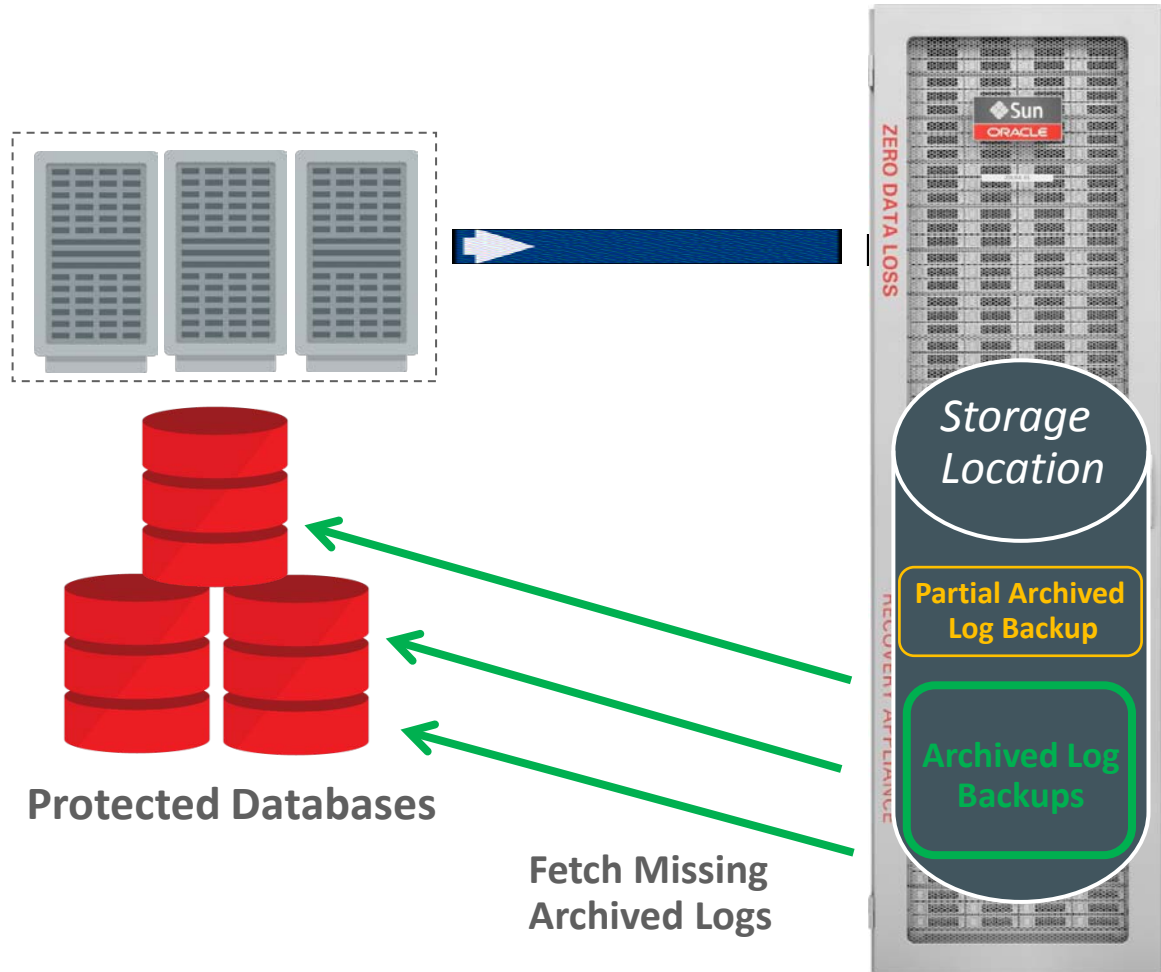
- After one-time full backup, incrementals used to create ***virtual*** full database backups on a daily basis
 - Pointer-based representation of physical full backup as of incremental backup time
 - Virtual backups typically 10x space efficient
 - Enables long backup history to be kept with the smallest possible space consumption
 - “Time Machine” for database

Backup & Redo (Delta Push) Workflow

Disk/Tape/Replica Backup Lifecycle



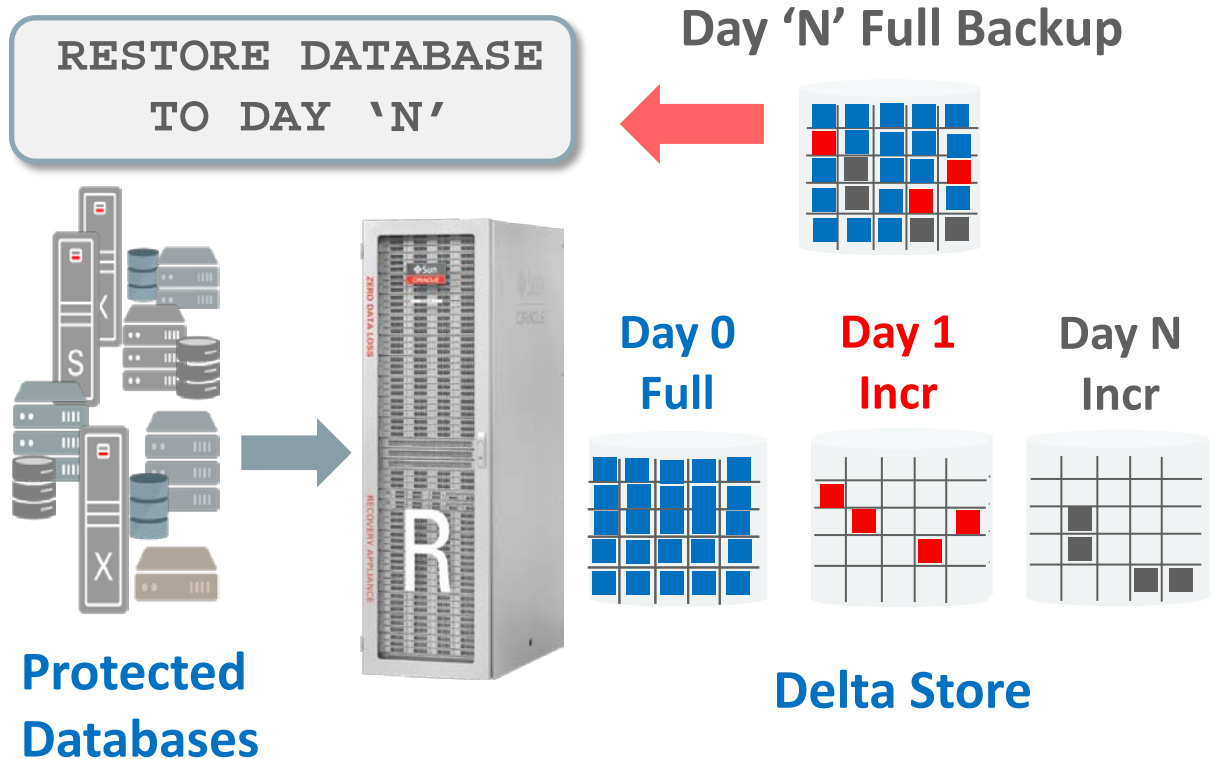
Real Time Redo Transport



- Redo transferred from memory buffers on database server – very low impact
- If redo connection is lost, the appliance ‘buttons up’ the received redo and creates a ‘partial archived log backup’
 - Preserves recovery until the last change received by the appliance (i.e. 0-1 second RPO)
- When connection is reinstated, gap detection process on the appliance automatically fetches missing archived logs from protected databases

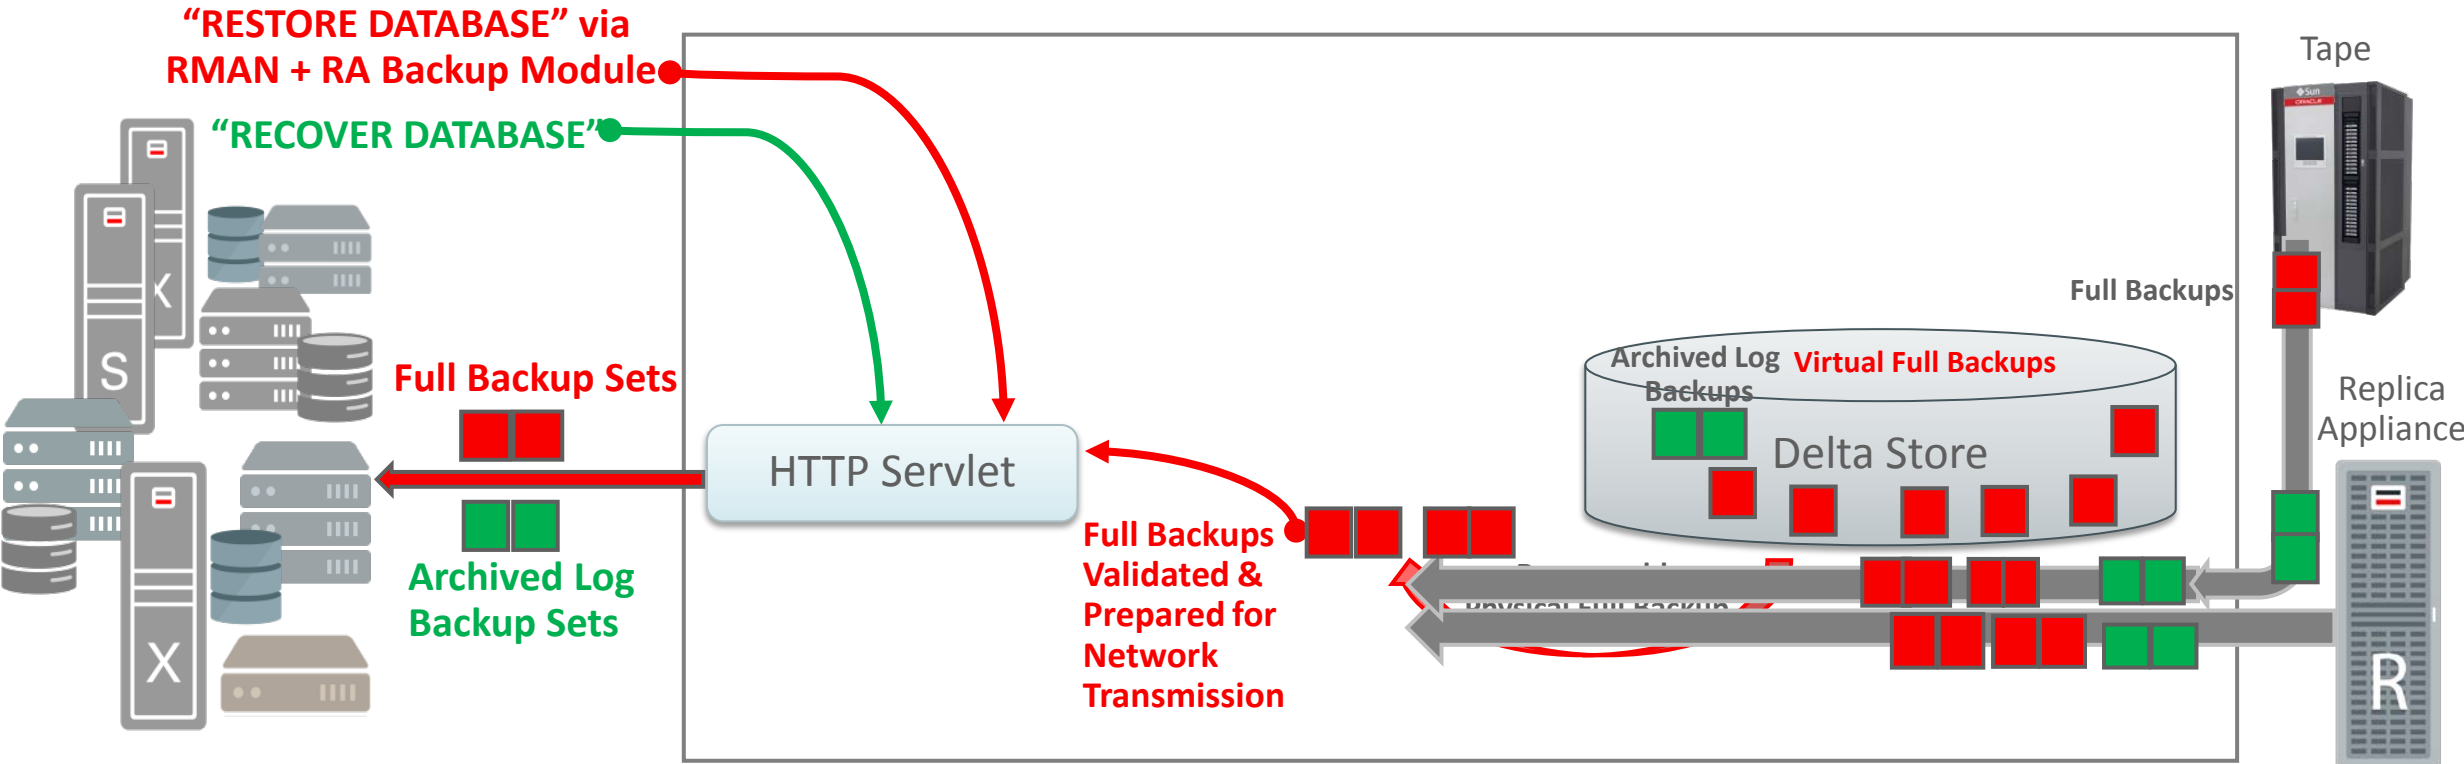
Fast Restore to Any Point-in-Time

No Load On Production Servers to Merge Old Backups



- Directly restore any virtual full backup
 - All blocks referenced from virtual full are efficiently retrieved
 - Eliminates production server overhead of traditional restore and merge of incrementals
- Supported by the scalability and performance of the underlying Exadata hardware architecture


Restore Workflow



Agenda

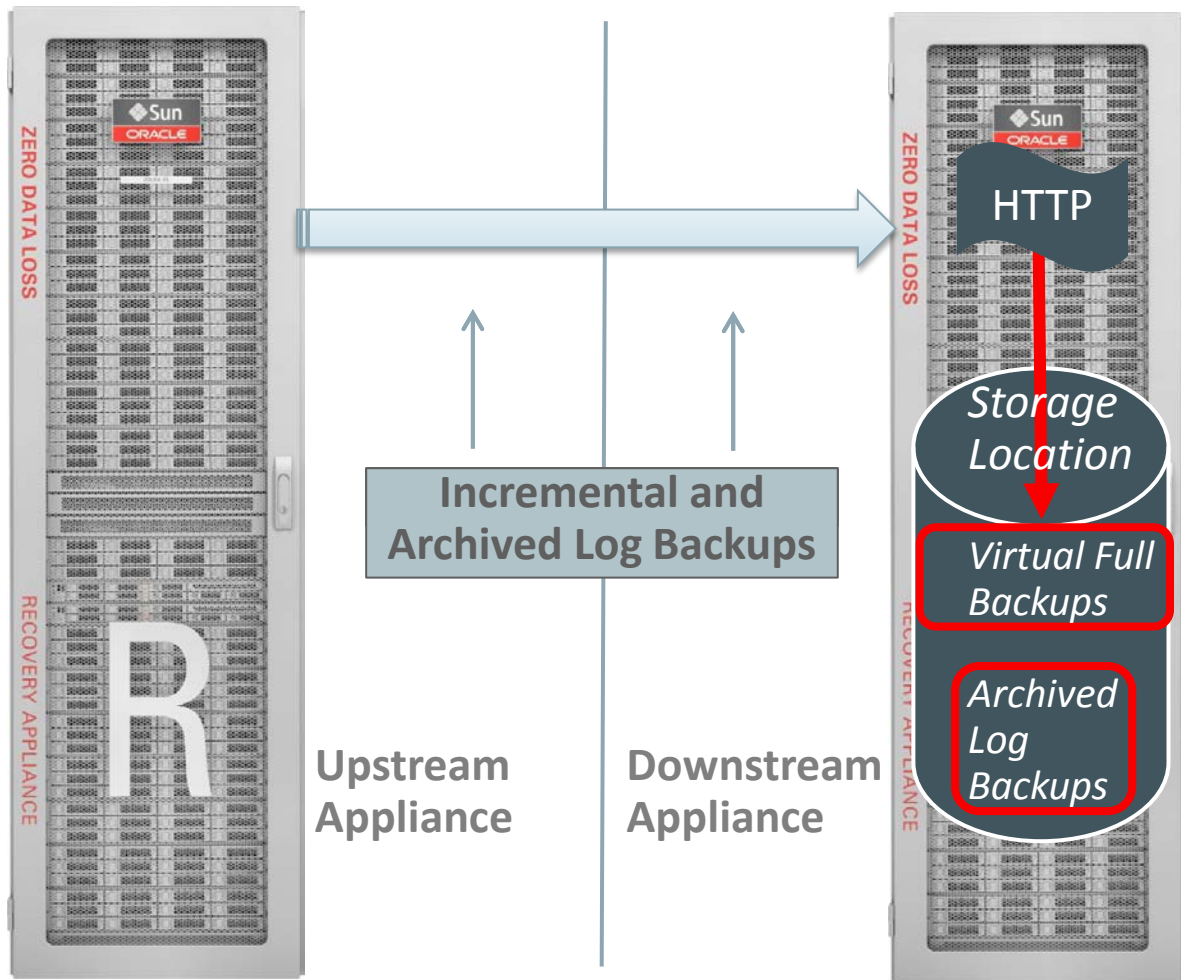
- 1 Today's Backup & Recovery Challenges
- 2 Recovery Appliance Architecture
- 3 Virtual Full Backup & Real-Time Redo Transport
- 4 Tape Backup & Replication**
- 5 Managing Time and Space
- 6 Customer Case Studies
- 7 Summary / Q&A

Tape Backup Data Flow



- Create Media Manager
 - SBT library pathname
 - Number of total tape drives
 - Number of dedicated tape drives for restore
- Create Attribute Set (media family, copies, etc.)
- Create Copy to Tape Job Template
 - Media Manager, Attribute Set
 - Protection Policy or Protected Database
 - **FULL / INCREMENTAL / ARCHIVED LOG / ALL**
- Queue / schedule tape job for CTDB at TIME 'T1'
 - **FULL** – Most recent full backup of CTDB prior to TIME 'T1' (assembled from corresponding virtual full)
 - **INCREMENTAL** – Incremental backups of CTDB relative to FULL backup
 - **ARCHIVED LOG** – Archived log backups of CTDB relative to FULL backup


Replication Data Flow



- When a Protection Policy is added to Replication Server on Upstream Appliance, all future incremental and archived log backups for its protected databases will be queued for replication
- Level 0 backup is first replicated
- Each level 1 incremental received by downstream generates a new virtual full
- New virtual full records on downstream are periodically synchronized to upstream catalog
- Number of concurrent replication streams can be adjusted, based on network consumption / environment

Agenda

- 1 Today's Backup & Recovery Challenges
- 2 Recovery Appliance Architecture
- 3 Virtual Full Backup & Real-Time Redo Transport
- 4 Tape Backup & Replication
- 5 Managing Time and Space**
- 6 Customer Case Studies
- 7 Summary / Q&A



Policy-Based Cloud-Scale Database Protection

Gold Policy – Customer Critical




Disk: 35 days
Tape: 90 days



Silver Policy – Internal Critical



Disk: 10 days
Tape: 45 days



Bronze Policy - Test/Dev



Disk: 3 days
Tape: 30 days



Tape



Replica



Replica Recovery
Appliance also
Policy-Based

Recovery Appliance Protection Policies

- Standardized recovery window, tape retention, replication policies

Real-Time Database Recoverability & Space Monitoring

Enterprise Manager Provides Key Metrics At Your Fingertips

ZDLRA Nashua > Protected Databases

Protected Databases (29)

Search

View Add Edit Remove Remove Multiple Details

Database	Version	Protection Policy	Database Size (GB)	Recovery Window			Unprotected Data Window		Errors and Warnings	Redo Transport	Copy-to-Tape	Replication	Last Complete Backup
				Goal	Current	Needed Space (GB)	Goal	Current					
EM13REP	12.1.0.2.0	SILVER_DS	13.46	10 days	N/A	100.00	N/A	N/A				Oct 15, 2:32 PM ...	
ASTORYA	12.1.0.2.0	SILVER	5271.89	10 days	6 days 01:09	2688.51	N/A	< 1 sec	✓	✓	✓	Oct 15, 11:08 AM ...	

Columns Hidden 1

Protected Database Detail: ASTORYA

Status Incidents and Events

Unprotected Data and Recovery Window

Backup Space

De-Duplication Ratio 10.2 :1

Keep Backup Space

Next Scheduled Backup View

Last Copy To Tape Oct 13, 11:28 PM EDT

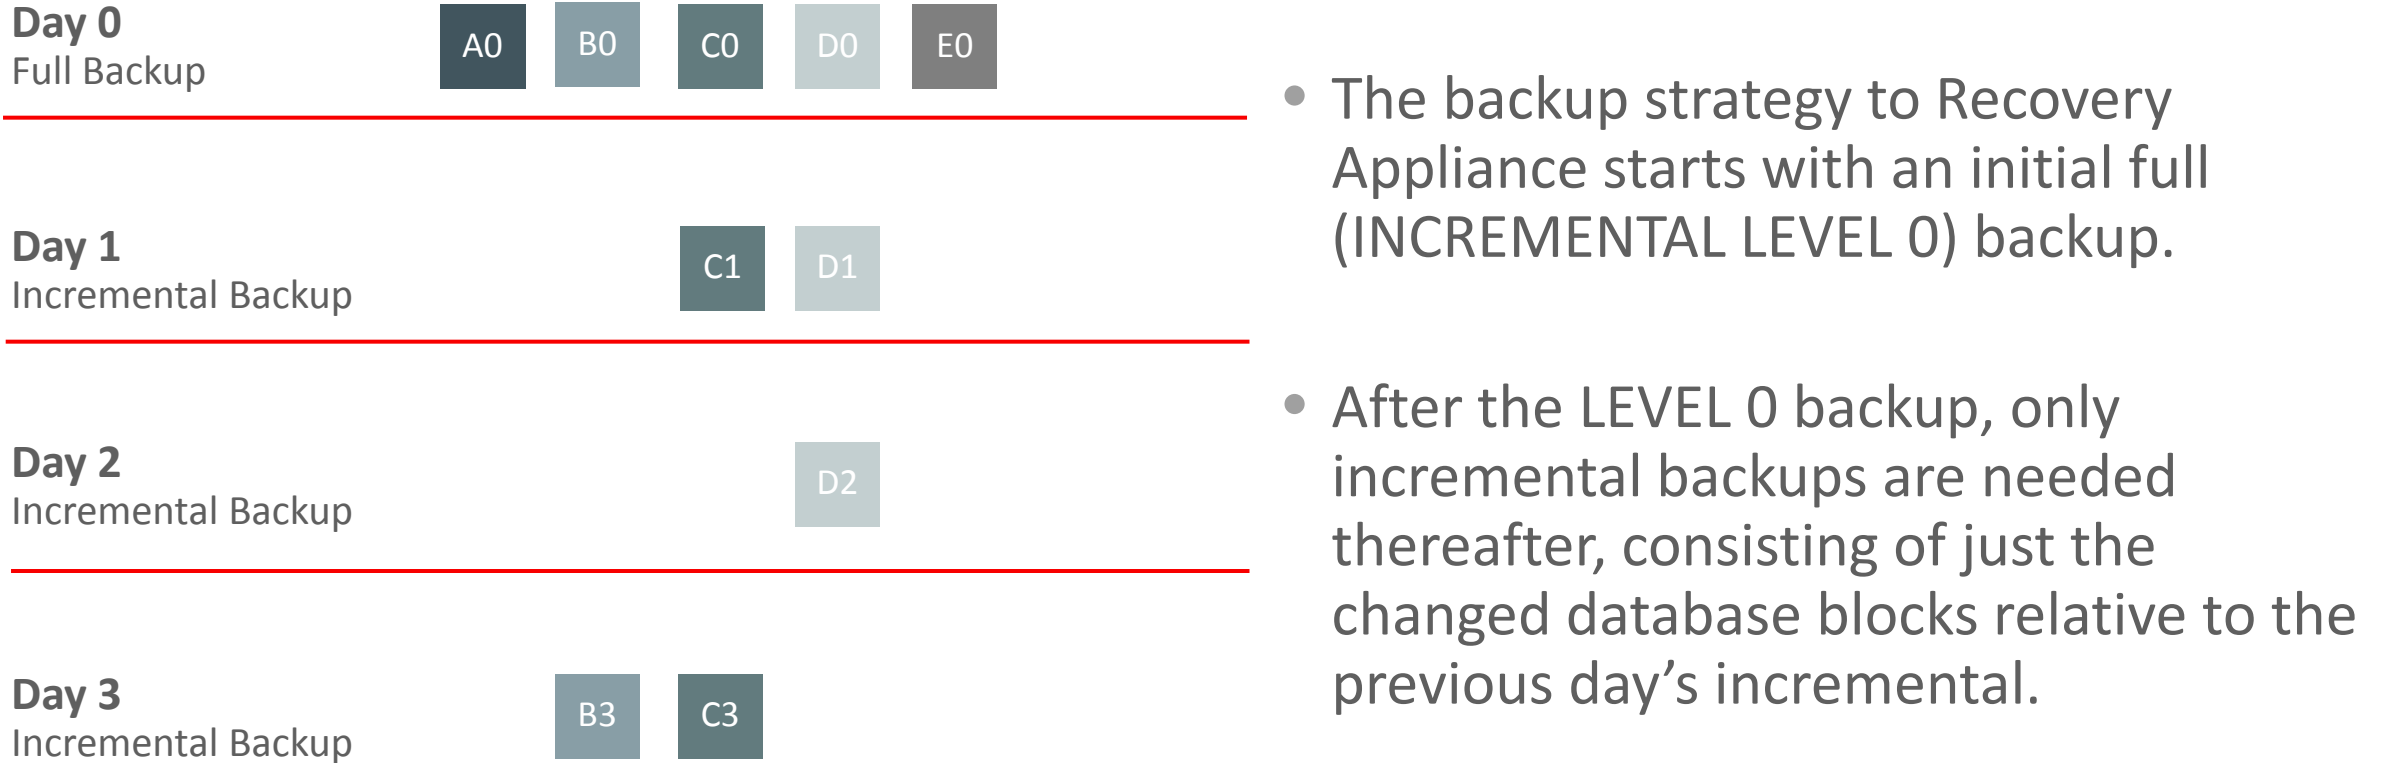
Last Replication Oct 15, 3:11 PM EDT

✓ TIP For more information, view the following reports for this database.

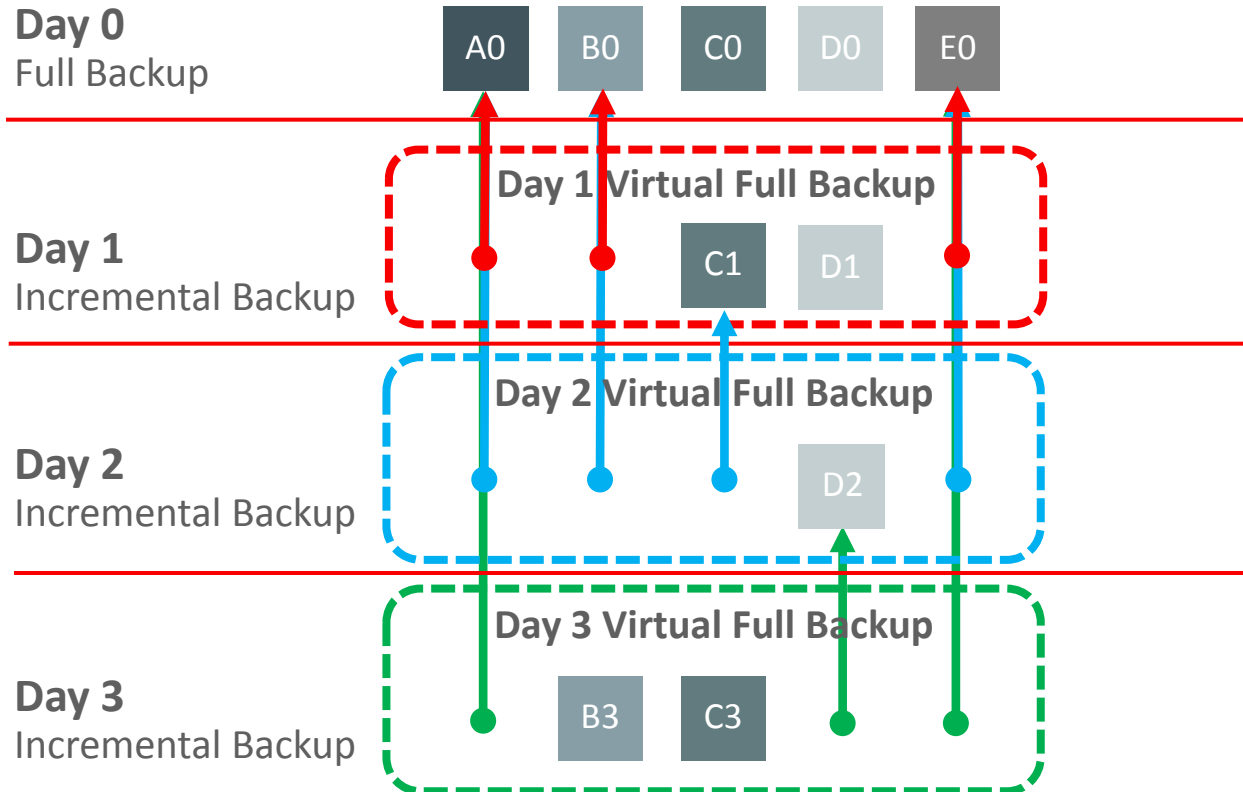
[Protected Database Report](#)

[Backup Report](#)

Delta Store – Incremental Backups

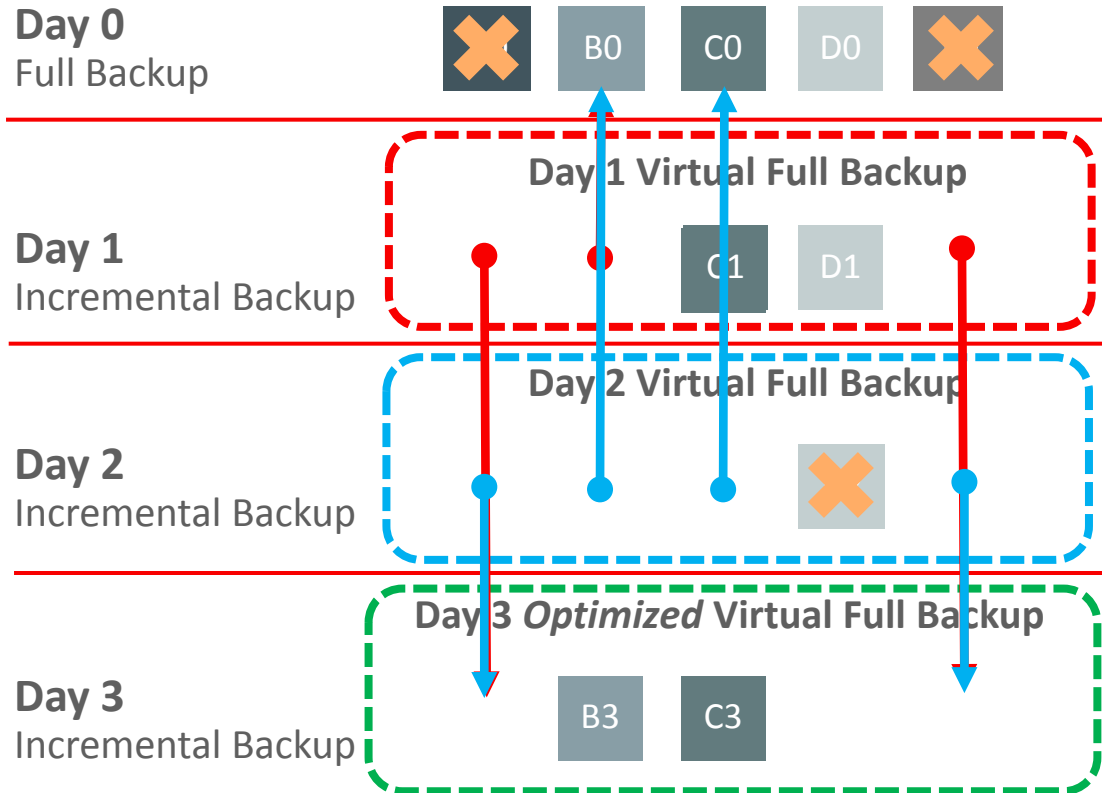


Delta Store – Virtual Full Backup Creation



- A virtual full backup is created in the Recovery Appliance catalog for each received RMAN incremental backup.
- It operates at a data file level and appears as a normal LEVEL 0 backup in RMAN LIST/REPORT & catalog queries.

Delta Store – Virtual Full Backup Optimization



- OPTIMIZE task is a background process that runs on the most recent virtual full backup to 'optimize' its read I/O access for future restore operations.
- Virtual full blocks are reordered into contiguous sets for fewer & larger read I/O requests, which is more efficient than many smaller read I/O requests.

Agenda

- 1 Today's Backup & Recovery Challenges
- 2 Recovery Appliance Architecture
- 3 Virtual Full Backup & Real-Time Redo Transport
- 4 Tape Backup & Replication
- 5 Managing Time and Space
- 6 Customer Case Studies**
- 7 Summary / Q&A

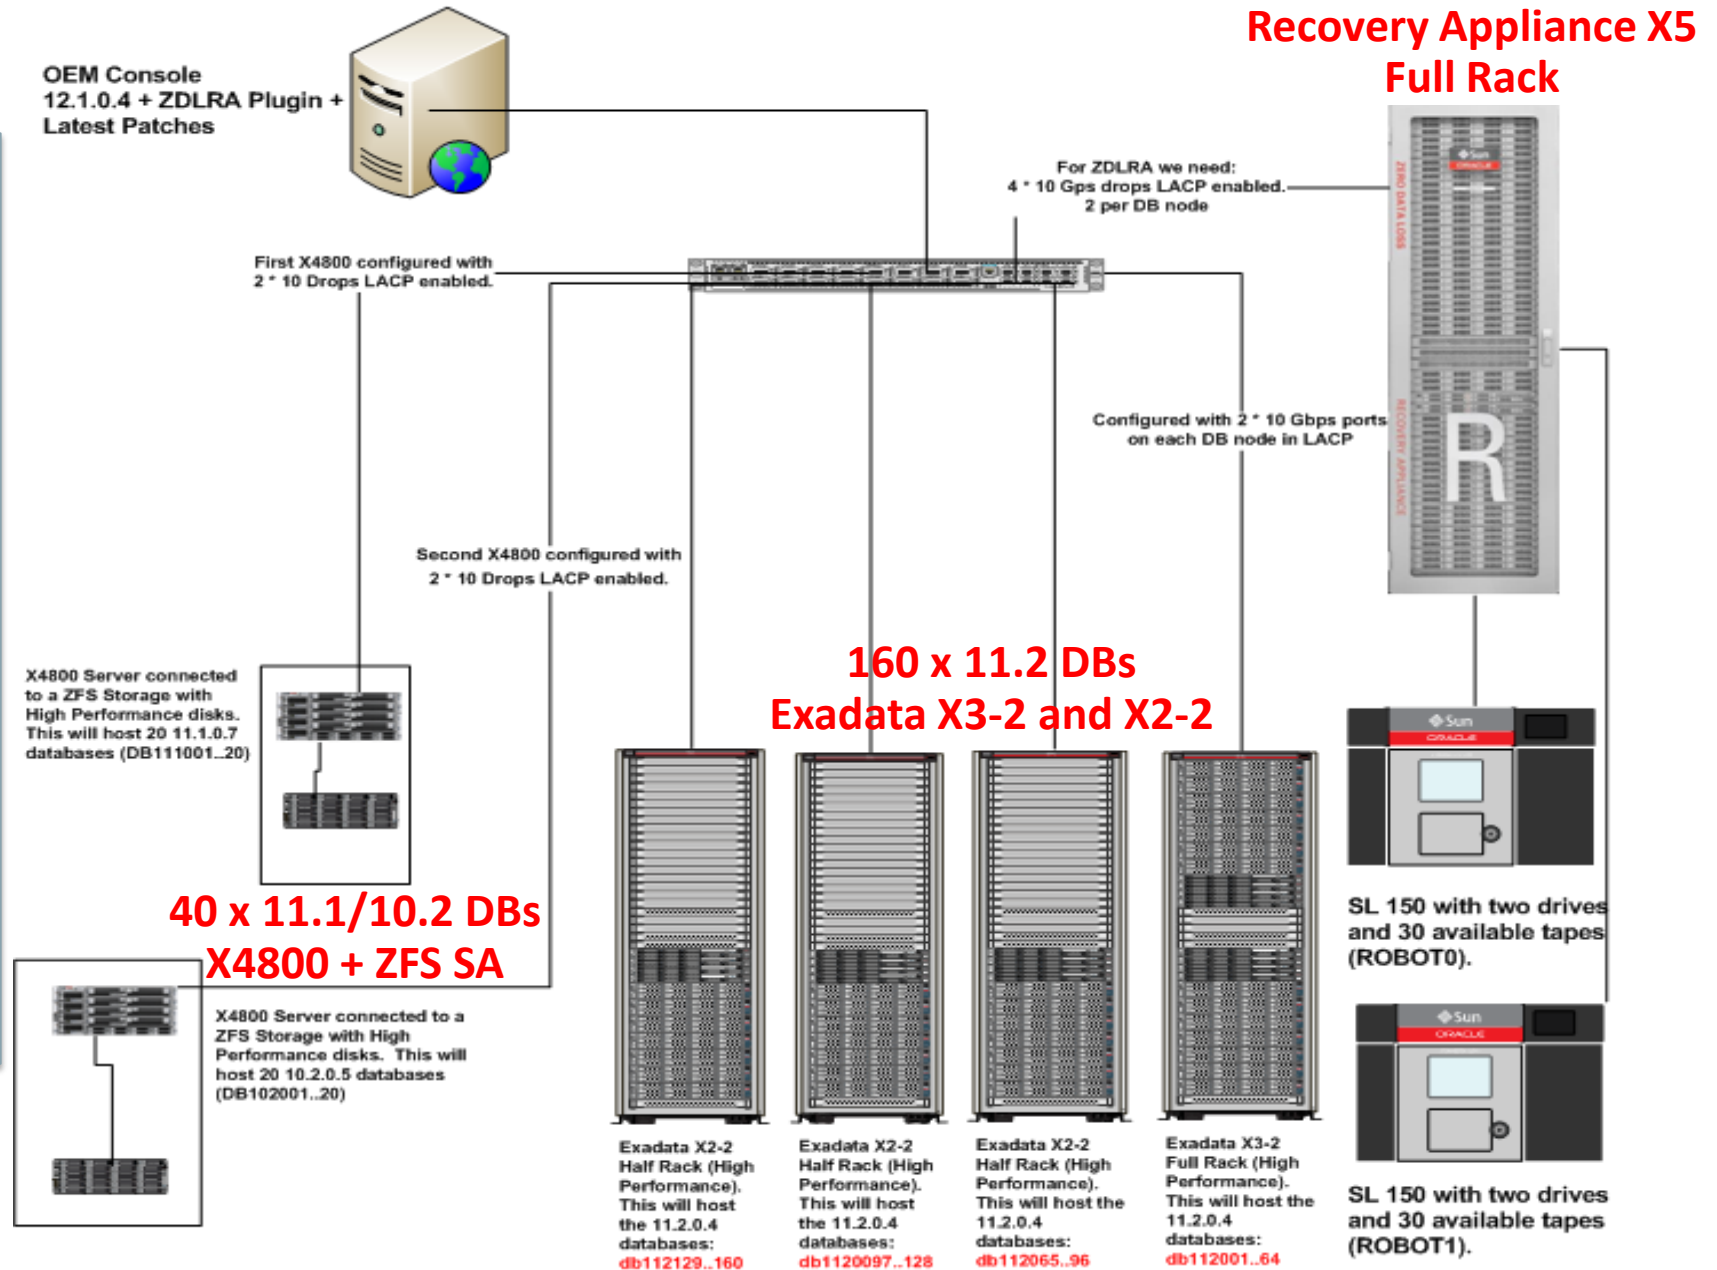
Very Large Oracle Customer

Backup & Recovery for 16,000 Enterprise Databases

- RMAN backups to local SAN/NAS “dump” storage, 2 weeks retention.
- Disk backups swept weekly by Third Party SW to Backup Appliance for 30 day retention
 - Replicated to a remote Backup Appliance residing at the bunker site
 - Backups are copied to physical tape to meet > 30 day retention needs.
- **Pain Point #1:** Very large & costly “dump” storage deployed across 16,000 databases.
- **Pain Point #2:** Inability to coordinate Third Party SW sweep schedule with RMAN backups being *fully completed*, resulting in incomplete backups..**50% restore failure rate**
- **Pain Point #3:** Non-Oracle integrated tape backup strategy..**> 48 hours RTO due to multiple stages in restore process**
- **Pervasive Risk: Recovery Failure + Data Loss = Non-Compliance, \$\$ Penalties, Bad Press**

POC Details

- Concurrent 200 DB backups complete in 8 hours
Done in 2 ½ hours
- Copy virtual full backup of 200 DBs to tape in 7 days
Done in 2 days
- Report real-time RPO of < 1 second on 160 x 11.2 DBs
ZERO redo transport lag
- Restore 2 DBs while rest (198 DBs) are backing up
Done in 2 hours





Customer Case Study

Fernando Simon

DBA Manager

Brazilian Justice Tribunal of Santa Catarina

Brazilian Justice Tribunal of Santa Catarina

- Justice Tribunal of Santa Catarina (TJSC)
 - Civil and Criminal Court
- Santa Catarina in numbers:
 - Nearly 7 million residents
 - 295 cities
 - 6th largest GDP by region in Brazil
- TJSC in numbers:
 - 13,000 employees
 - Internal and public services
 - Courts, Judicial Processes, Appeals, Taxes
 - Nearly 14 million judicial cases recorded since inception of Oracle-based system
 - 723,000 just in 2015
 - **All judicial processes are recorded in digital format..fully paperless**
 - **24x7** services for **Legal Appeals** requirements

Brazilian Justice Tribunal of Santa Catarina

- TJSC and Oracle:
 - Oracle user since 1998
 - Nearly 40 production databases
 - 15 critical databases:
 - 30TB total volume
 - Judicial Process, Appeals, Taxes, HR
 - 24x7 access
 - Major Database:
 - More than 8,000 internal users by day
 - More than 70,000 web access by day
 - 17TB size and 100,000 IOPS
 - DSS/OLTP workloads
- TJSC and Engineered Systems:
 - 4 Exadatas: 1 Half V2 (HP), 1 Half X2 (HP), 1 Full X4 (HP), **1 Full X5 (EF)**
 - Among the first OLTP users of Exadata
 - **2 Recovery Appliances**
 - No outages since 2010 (1st Exadata deployment, 2nd in all of Brazil)
 - All databases run on Exadata with IORM (catPlan and dbPlan), Database Resource Manager, RAC Services and Instance Caging

Brazilian Justice Tribunal of Santa Catarina



- Local Partner Samaia IT
 - System Integrator
 - Platinum Partner
 - Winner of the bidding
 - Concepts & implementation of the project
- Samaia IT projects with TJSC:
 - 4 Exadata racks
 - 2 Recovery Appliances
 - 2 StorageTek SL150 tape libraries



Brazilian Justice Tribunal of Santa Catarina

- TJSC and Other HA Infrastructure:
 - Data Center Bunker (started in 2014)
 - Redundant Network Switches
 - Cisco Nexus 7000
 - 1 and 10Gbps network
 - Redundant Storage
 - EMC VNX 5500, VNX 5400
 - Redundant Media Management
 - 2 StorageTek SL150 (8 LTO-6 Drives)
 - HP Blades for Applications and VMware

– *You don't want to remain in jail just because of an Oracle Database outage, right?*


Brazilian Justice Tribunal of Santa Catarina

- TJSC and Backup/Recovery – without Recovery Appliance
 - EMC Networker
 - 8.0.2
 - Two Node Clustered Media Servers
 - 1Gbps communication network
 - Data Domain
 - DD670 configured as VTL
 - 50TB usable capacity
 - **2 Gbps SAN network**
 - **Shared: Databases and File Servers**

- Problems:
 - **Low deduplication** for databases
 - Even after following EMC best practices
 - **Slow, nearly 3 days** to backup most critical 17 TB database
 - Large DBA effort to manage backups
 - Needed to expire old backups from DD every day, due to space constraints (no additional savings from deduplication).
 - Nearly 3 days to recover 17 TB database
 - Difficult to deliver adequate RPO and RTO due to network and HW constraints

Brazilian Justice Tribunal of Santa Catarina

- TJSC and Backup/Recovery – without Recovery Appliance



Brazilian Justice Tribunal of Santa Catarina

Before Recovery Appliance..50, 65, 90+ hours to complete full backups

ORACLE Enterprise Manager Cloud Control 12c

Enterprise Targets Favorites History Search Target Name

opg5 Cluster Database Performance Availability Security Schema Administration

Results

Total 91 (Completed 84 Failed 7)

Backup Name	Status	Start Time	Time Taken	Input Size	Output Size	Output Rate (Per Sec)
BACKUP_OPG5_INCREM_082715013001	COMPLETED	Aug 26, 2015 7:31:00 PM BRT	01:33:06	306.63G	426.72G	78.23M
2015-08-26T00:35:13	COMPLETED	Aug 26, 2015 12:35:14 AM BRT	00:46:18	207.80G	364.06G	134.20M
2015-08-24T23:47:43	COMPLETED	Aug 24, 2015 11:47:44 PM BRT	00:58:06	110.49G	312.58G	91.82M
2015-08-21T19:15:33	COMPLETED	Aug 21, 2015 7:15:34 PM BRT	51:35:07	17.46T	17.34T	97.93M
2015-08-21T00:21:14	COMPLETED	Aug 21, 2015 12:21:26 AM BRT	02:47:30	1010.49G	1.12T	117.32M
BACKUP_OPG5_INCREM_082015050002	COMPLETED	Aug 19, 2015 11:00:39 PM BRT	01:20:49	110.03G	268.63G	56.73M
BACKUP_OPG5_INCREM_081915050001	COMPLETED	Aug 18, 2015 11:00:23 PM BRT	01:23:00	109.90G	270.28G	55.57M
BACKUP_OPG5_INCREM_081815050004	COMPLETED	Aug 17, 2015 11:00:36 PM BRT	02:58:28	251.26G	582.83G	55.74M
BACKUP_OPG5_INCREM_081415050004	COMPLETED	Aug 13, 2015 11:00:33 PM BRT	01:17:24	108.43G	235.20G	51.86M
BACKUP_OPG5_INCREM_081315050003	COMPLETED	Aug 12, 2015 11:00:34 PM BRT	01:18:38	121.70G	255.12G	55.37M
BACKUP_OPG5_INCREM_081215050003	COMPLETED	Aug 11, 2015 11:00:26 PM BRT	01:28:39	111.39G	246.85G	47.52M
BACKUP_OPG5_INCREM_081115050001	COMPLETED	Aug 10, 2015 11:00:24 PM BRT	01:24:35	92.83G	291.37G	58.79M
BACKUP_OPG5_SEMANA_080815011000	COMPLETED	Aug 7, 2015 7:10:42 PM BRT	65:59:24	16.16T	16.04T	70.82M
2015-08-06T23:50:02	COMPLETED	Aug 6, 2015 11:50:20 PM BRT	18:24:09	289.05G	458.39G	7.09M
2015-08-05T23:50:02	COMPLETED	Aug 5, 2015 11:50:20 PM BRT	02:34:38	139.54G	263.06G	29.03M
2015-08-04T20:13:46	COMPLETED	Aug 4, 2015 8:14:04 PM BRT	04:03:08	44.28G	313.06G	21.98M
2015-07-31T23:30:03	COMPLETED	Jul 31, 2015 11:30:21 PM BRT	90:18:14	16.91T	16.65T	53.69M
2015-07-30T23:50:02	COMPLETED	Jul 30, 2015 11:50:20 PM BRT	03:38:33	337.78G	498.87G	38.96M

50+ HRS!

65+ HRS!

90+ HRS!!

Brazilian Justice Tribunal of Santa Catarina

Data Domain Cleaning Task – 1 day+ to Complete

Management Station: .tjsc.jus.br sysadmin Log Out Refresh
Role: admin
.tjsc.jus.br OS: 5.4.0.8-404909 Model: DD670

File System More Tasks ▾ ?

State **Enabled and running: 15 weeks, 8 hours, 30 minutes** Disable...

Clean Status Cleaning finished at 2015/10/06 05:20:22 Start Cleaning...

Summary Configuration Encryption Space Usage **Consumption** Daily Written

Capacity
Fri Aug 07 2015 12:00 PM
50.0 TiB


Post-Comp
Fri Aug 07 2015 12:00 PM
49.7 TiB

Cleaning
Started: Fri Aug 07 2015 9:48 PM
Duration: 1 day 49 minutes 3.0 seconds



Brazilian Justice Tribunal of Santa Catarina

- TJSC and Backup/Recovery – **NOW WITH RECOVERY APPLIANCE**



Brazilian Justice Tribunal of Santa Catarina

ORACLE Enterprise Manager Cloud Control 12c

Enterprise Targets Favorites History

Search Target Name

opg5 Cluster Database Performance Availability Security Schema Administration

Logged in as sys | spmlx-edgx5-04.tjsc.jus.l

View Backup Report

The following backup jobs are known to the database. The data is retrieved from the recovery catalog. View Backup Sets And Image Copies

Search

Status: All Start Time: Within 1 month Type: All Go

Results

Total 15 (Completed 14 Failed 1)

Backup Name	Status	Start Time	Time Taken	Type	Output Devices	Input Size	Output Size	Output Rate (Per Sec)
-------------	--------	------------	------------	------	----------------	------------	-------------	-----------------------

Protected Database Detail: OPG5

Status Incidents and Events

Unprotected Data and Recovery Window

Unprotected Data: 17.41 hrs
Recovery Window: 11.00 days
Unprotected Data Threshold:
Recovery Window Goal: 10.00 days

Backup Space

Used: 26052.7 GB
Needed: 24022.64 GB
Reserved: 38912.0 GB

De-Duplication Ratio 11.9 : 1
Keep Backup Space None
Next Scheduled Backup View
Last Copy To Tape Oct 6, 7:07 PM BRT
Last Replication N/A
TIP For more information, view the following reports for this database.
Protected Database Report
Backup Report

Brazilian Justice Tribunal of Santa Catarina

- Backup/Recovery Benefits **with Recovery Appliance**

- **Reduce RTO and RPO**

- High Protection and Zero Data Loss

- High Deduplication Factor

- **Incremental Forever**

- New level 0 (virtual full) after only **23 minutes** incremental backup time
- Nearly **200X improvement** in backup performance
- Fast 10 Gbps Network

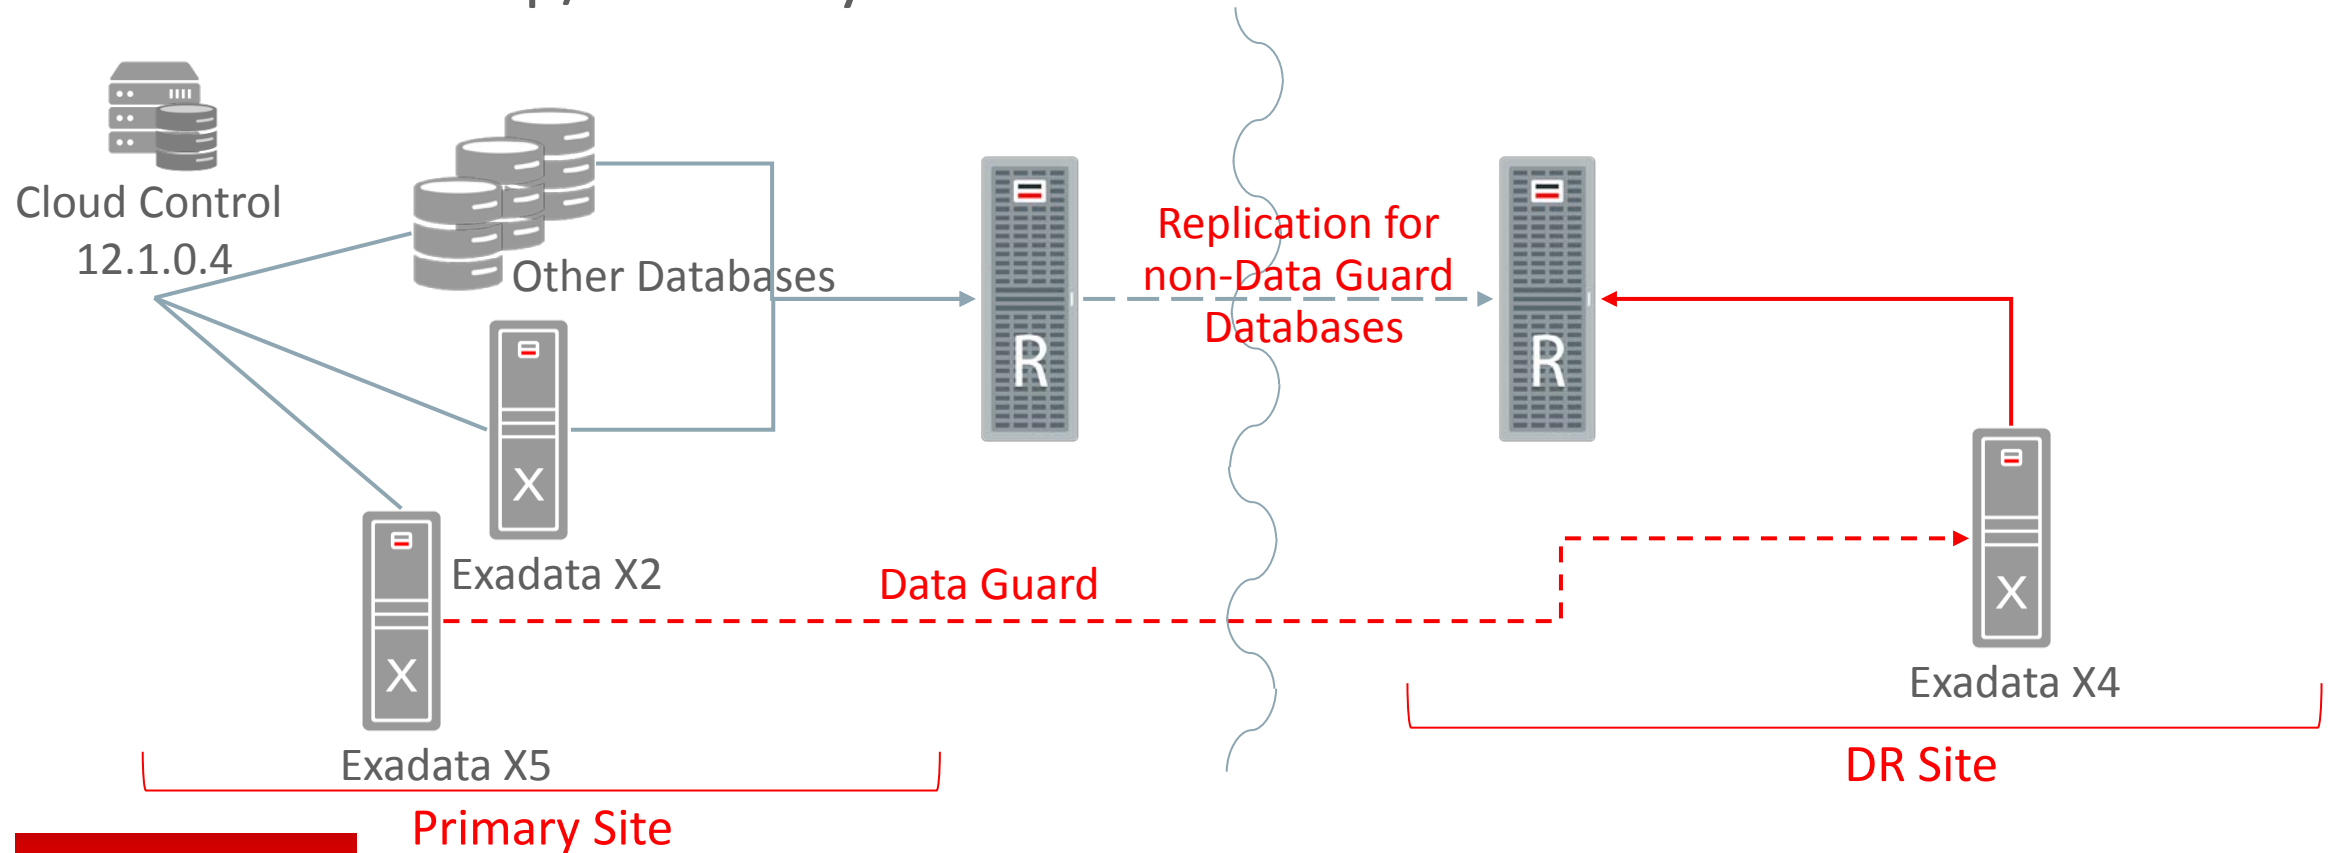
- Facilitated DB Migration from Exadata X4 to X5

- 2 Recovery Appliances less expensive than expanding Data Domain to support same backup requirements

- *The major benefit for us was*
Incremental Forever and Virtual Full

Brazilian Justice Tribunal of Santa Catarina

- TJSC and Backup/Recovery – **UPCOMING FULL MAA + RECOVERY APPLIANCE**



Brazilian Justice Tribunal of Santa Catarina

- TJSC and Recovery Appliance

- WHY:

- Key Features: Virtual Full / Incremental Forever
- Easy / measurable growth as required
- Impossible to sustain Data Domain
 - Lower Deduplication, Lower Performance
- Reduce the complexity of backup infrastructure
 - No more effort to tune and control every VTL of Data Domain
- Oracle on Oracle, ONE support organization for entire environment

- TJSC:

- www.tjsc.jus.br

- <https://www.youtube.com/user/canaltjsc>



Summary / Q&A

Summary

Recovery Appliance **Designed** for Oracle Data Protection

- Leverages proven Oracle Backup & Recovery + Data Guard + Robust, Scalable Exadata Hardware Platform
 - Protect database transactional changes in real-time
 - Reduce backup storage and production overhead via incremental-forever
 - Recover to any point-in-time with standard RMAN commands
 - Offload backups to tape via user-configured policies
 - Monitor space usage based on database recovery window goals
 - Gather real-time data protection status across the entire Oracle enterprise via EM
- For more information: oracle.com/recoveryappliance

Integrated Cloud

Applications & Platform Services

ORACLE®