Engineer **More**, Work **Less**

DBS Journey in Autonomous Database

Neighbours first, bankers second

DBS

Whatever your personal goals and family priorities are, they are ours too

Today – DBS, A Leading Financial Services Group in Asia

Headquartered in Singapore, with a growing presence in Greater China, Southeast Asia & South Asia

Best Digital Bank
in the World

33,000 Start-Up
Future-Ready Workforce

Best Bank
in the World

Most Innovative
in Digital Banking

Safest Bank
in Asia, "AA-" Rating

15% ROE
Profitable Performance

S$16.5 billion
In Income

S$743 billion
In Assets

Over 280 branches across 19 markets

DBS

*Financials as of 2021

DBS Journey

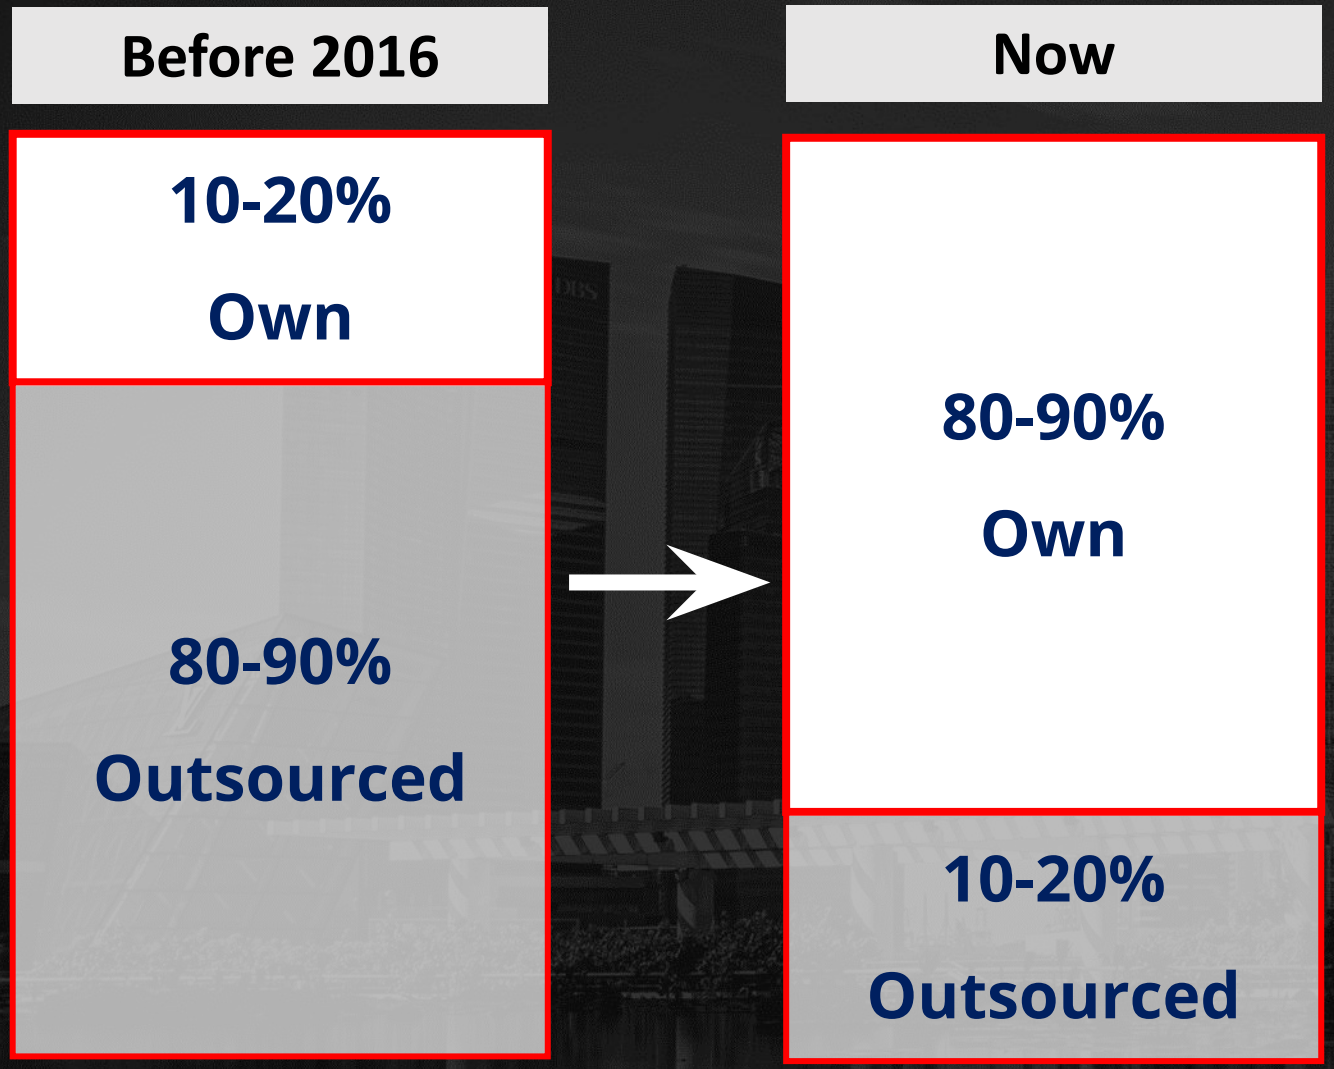

~ 40,000 Staff
~ 10,000 Technologists

- Bank to Technology Centric
- Embedding ourselves in the customer journey
- Thinking and acting like a startup

Being the D in

Tech transformation: History of Big Tech

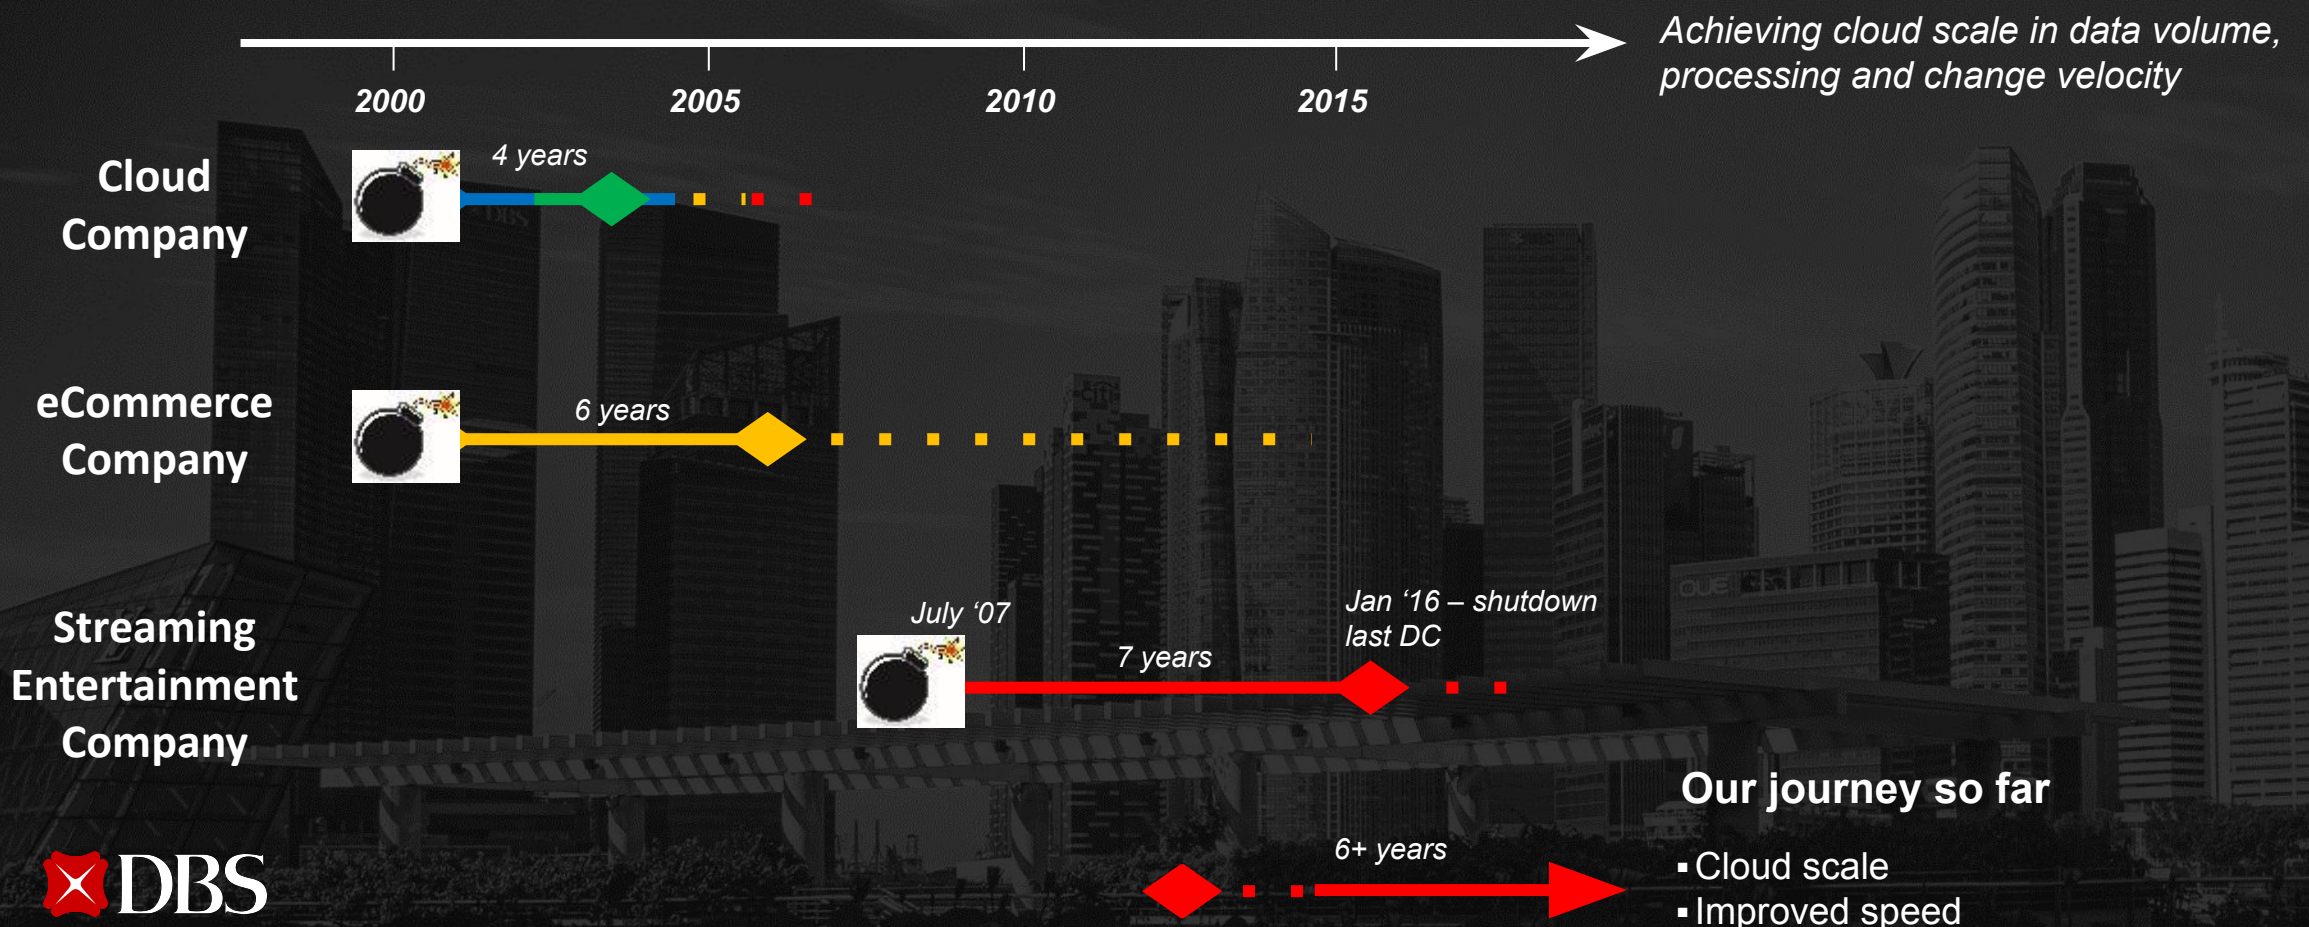

Our journey so far

- Cloud scale
- Improved speed
- Improved resilience
- Open source

We embarked on a strategy focused on 5 key elements to drive change throughout the organization

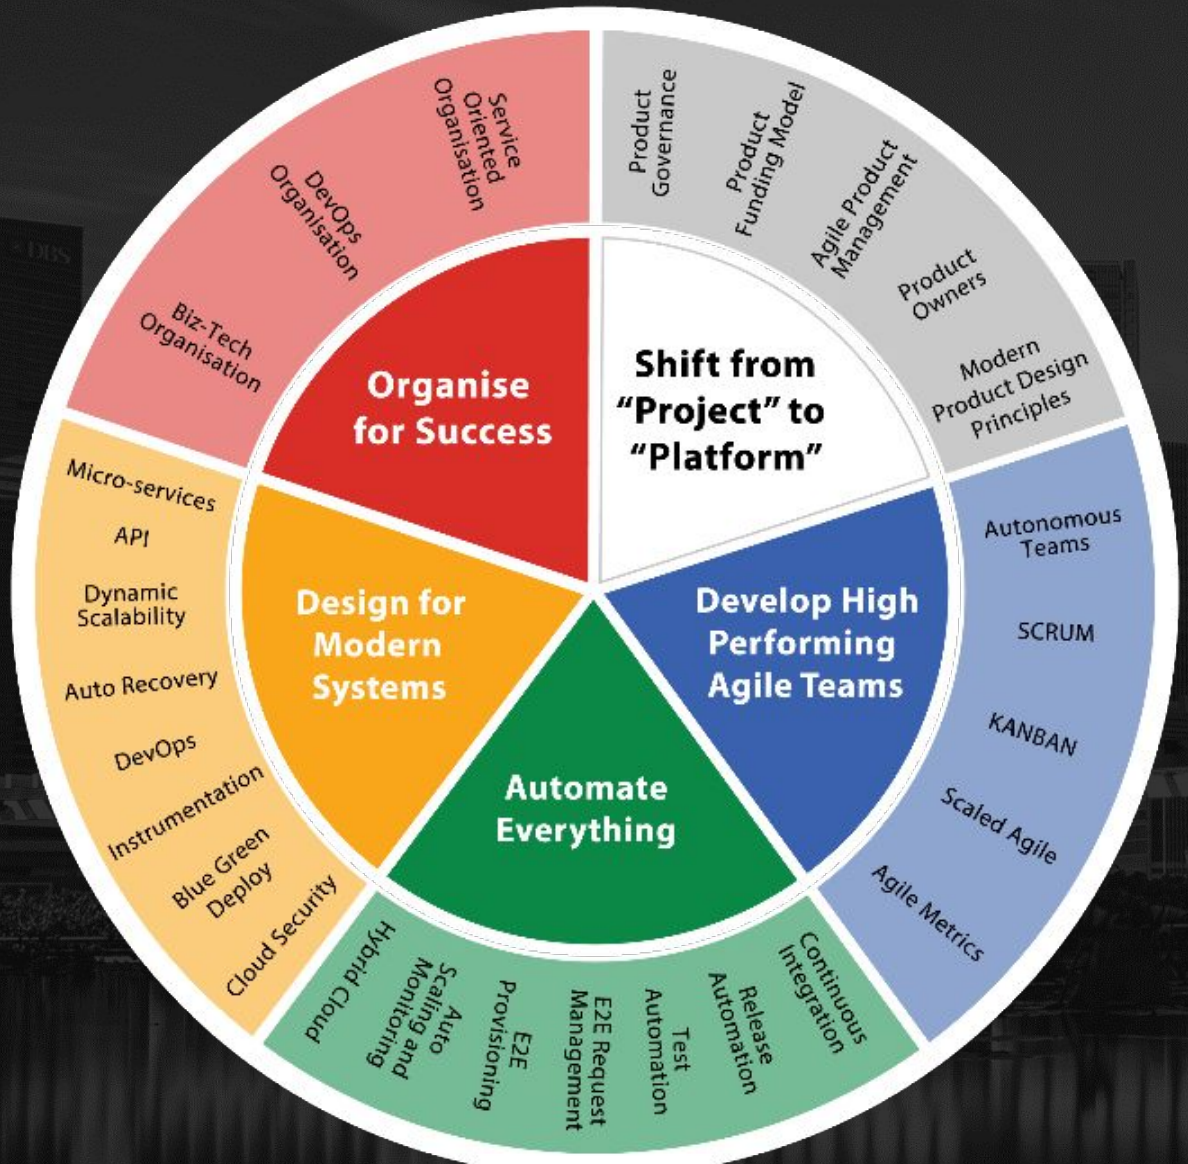

Everyone has a cloud strategy, but not all are the same

CLOUD LIPSTICK?

Typical banks

OR CLOUD TO THE CORE?

DBS

Reliability ≠ Resiliency

Production

Disaster Recovery

99.999%

(1 out of 100,000)

Reliability ≠ Resiliency

> 100,000 flights per day

Commodity Hardware

- *4x to 5x cheaper than enterprise resilient hardware*
- *2.8x more power efficient*
- *Equal or higher processing capacity*

Challenges

- Pros and Cons of Traditional method vs Cloud way
- Mitigating the challenges into easy operation and maintenance
- Automations

	Traditional	Cloud
HW Failure	< 1% failure	5-7% failure
Scale	3,000 OS 0 MariaDB	40,000 OS 6000 MariaDB/Maxscale Build/Change/Operate
Skillset	Oracle DBA MSSQL DBA DB2 DBA	MariaDB DBA (??)
Resiliency	Greenzone = Downtime	Greenzone ≠ Downtime

MariaDB - Autonomous Database in DBS

Spent More Time Innovating/Automating & Less Time Managing

Key Considerations for MariaDB deployment in DBS

1. uto deployment for all regions and all zones
2. ero downtime for security patches and minor bug fix patches
3. Zero downtime upgrade for major version upgrades
4. bility to handle ESX reboots or failures
5. plit read queries and write queries for load balancing
6. eamless deployment in adding more slaves during runtime
7. bility to handle accidental data changes
8. Ability to handle up to 2 level failures

DBS

Day 0 (Deployment Pattern 1) : Standalone DB

For less critical or internal applications

Resiliency = Good

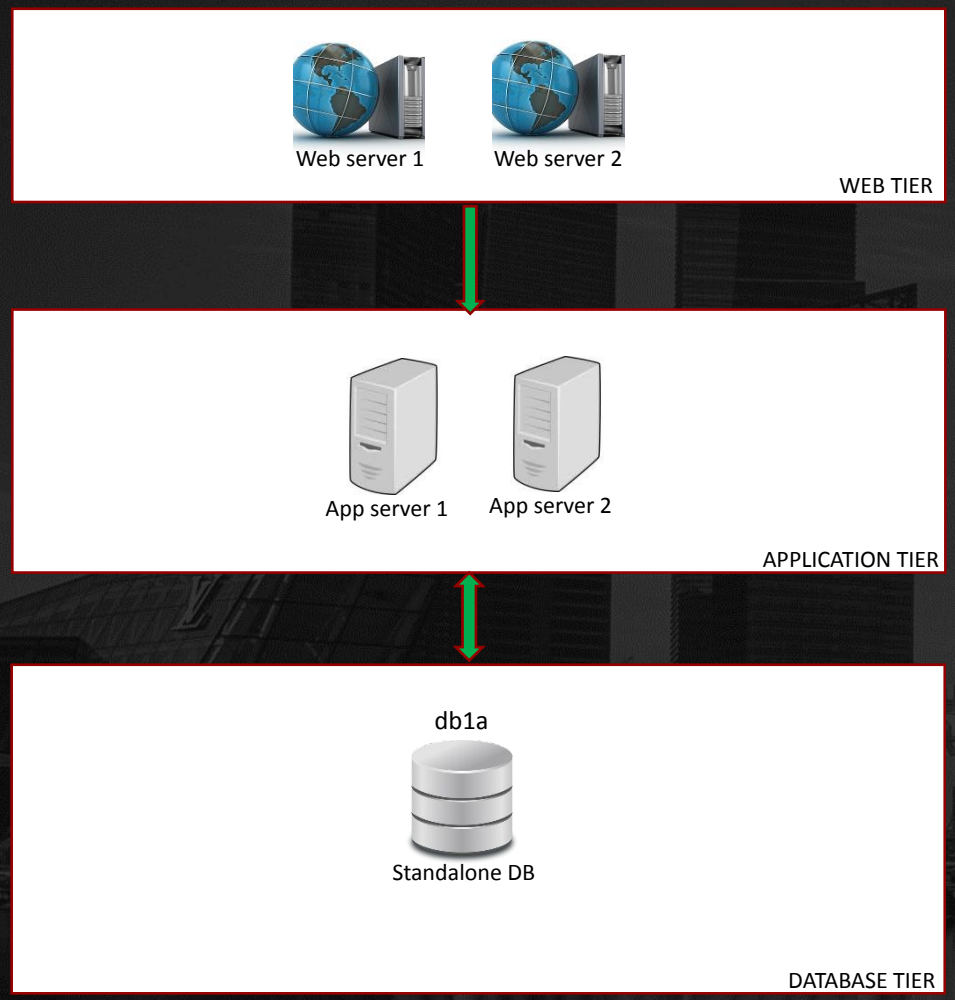

Standalone DB

+ VPC redundancy (5-10mins)
+ Backup redundancy (~4-12hrs)

Constraints: app need to auto rebind to DB connection

1. Leverages on default VPC resiliency. VM reboot within 5-10mins.
2. Application must rebind upon MariaDB VM reboot, else recovery will take longer
3. Application service will be down during 5-10min of VM reboot
4. Standalone deployment is concentrated across less critical systems

Day 0 (Pattern 2) : Master – Slave with Maxscale

For applications requiring higher level of resiliency and redundancy

Resiliency = Better

2/3 Node Master Slave with Maxscale Auto failover

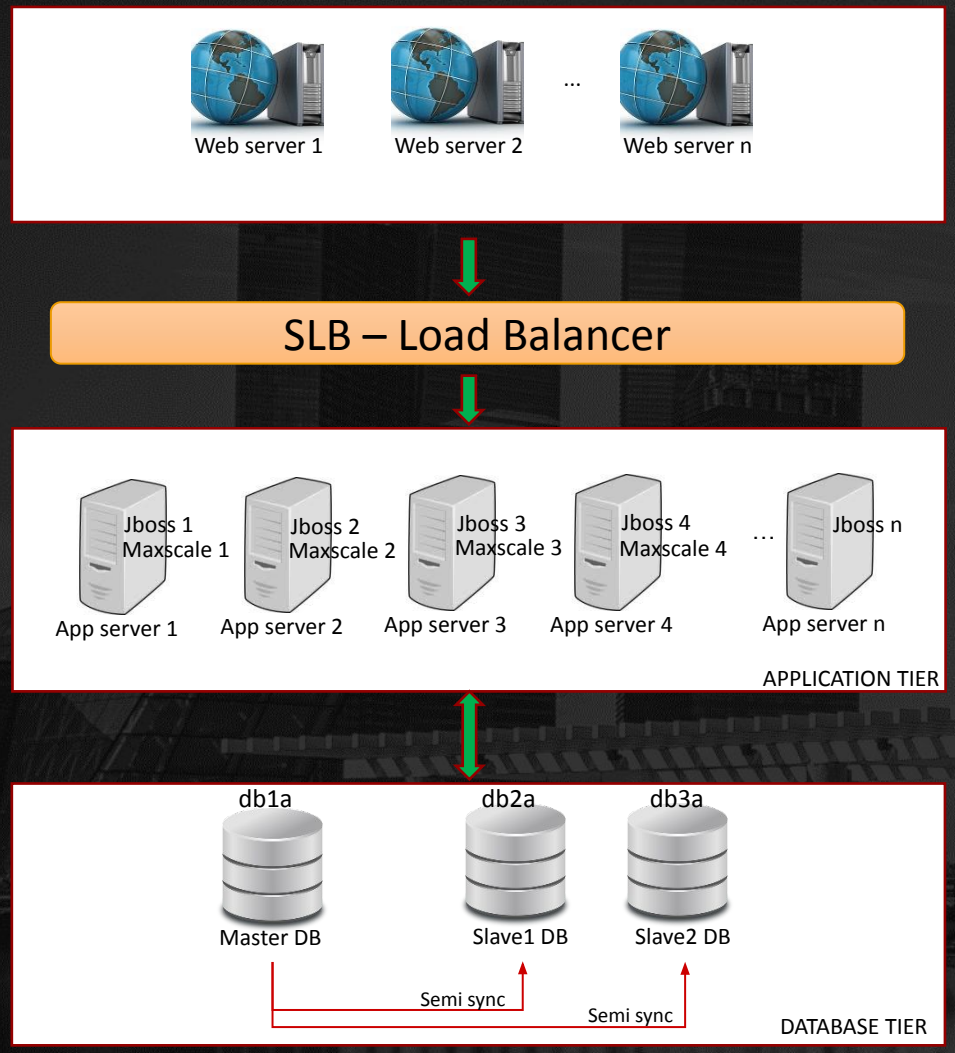

Master – Slave, Maxscale

+ MariaDB redundancy (30-40s unplanned)
+ VPC redundancy (5-10mins)
+ Backup redundancy (~4-12hrs)

Constraints: Will not be able to recover quickly from data corruption

1. Maxscale redundancy kick in to connect to Slave MariaDB (and promote Slave to Master) while VPC redundancy will reboot the failed master
2. Maxscale failover is transparent to app and will take < 1 min
3. Failed master will be bought up as slave and once sync up will provide additional redundancy. Cannot support 2 concurrent failure.
4. Read Write load balancing between Master and Slave nodes
5. App service will be down during Maxscale failover of < 1 min

Day 0 (Pattern 3) : Master – Slave with Maxscale & delayed replicas

For applications requiring highest level of resiliency, redundancy and faster recovery

Resiliency = Best

2/3 Node Master Slave with Maxscale and 2 delayed Slaves (6hrs, 24hrs)

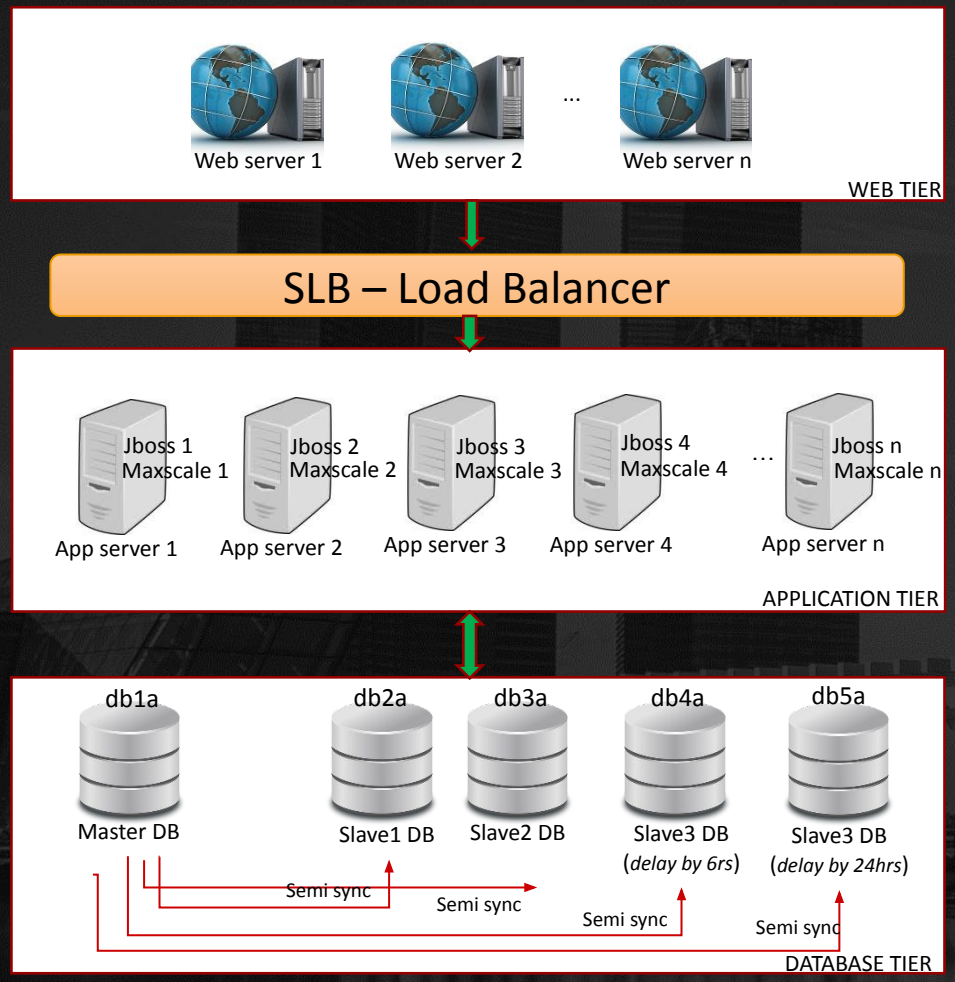

Master – Slave, Maxscale, Delayed Replica

+ MariaDB redundancy (30-40s unplanned)
+ VPC redundancy (5-10mins)
+ Backup redundancy (~4-12hrs)
+ Data Replication redundancy (~1^{++}hr)

Constraints: Decision making should be quicker to activate delayed replica. Resync the other servers in the cluster (~4-12hrs).

1. Default Maxscale redundancy as describe in earlier page.
2. Additional Replica / Slave redundancy lagged by 6 or 24 hr. Protect against intentional or unintentional data corruption.
3. Delayed replica copy can be activated instantly but will depend on app design if need to roll to latest record.
4. Transaction roll forward time will depend on actual data change and likely be between 1-2 hrs.

Day 0 (Various Deployment Patterns)
Provisioning database VMs via DBS Technology Marketplace

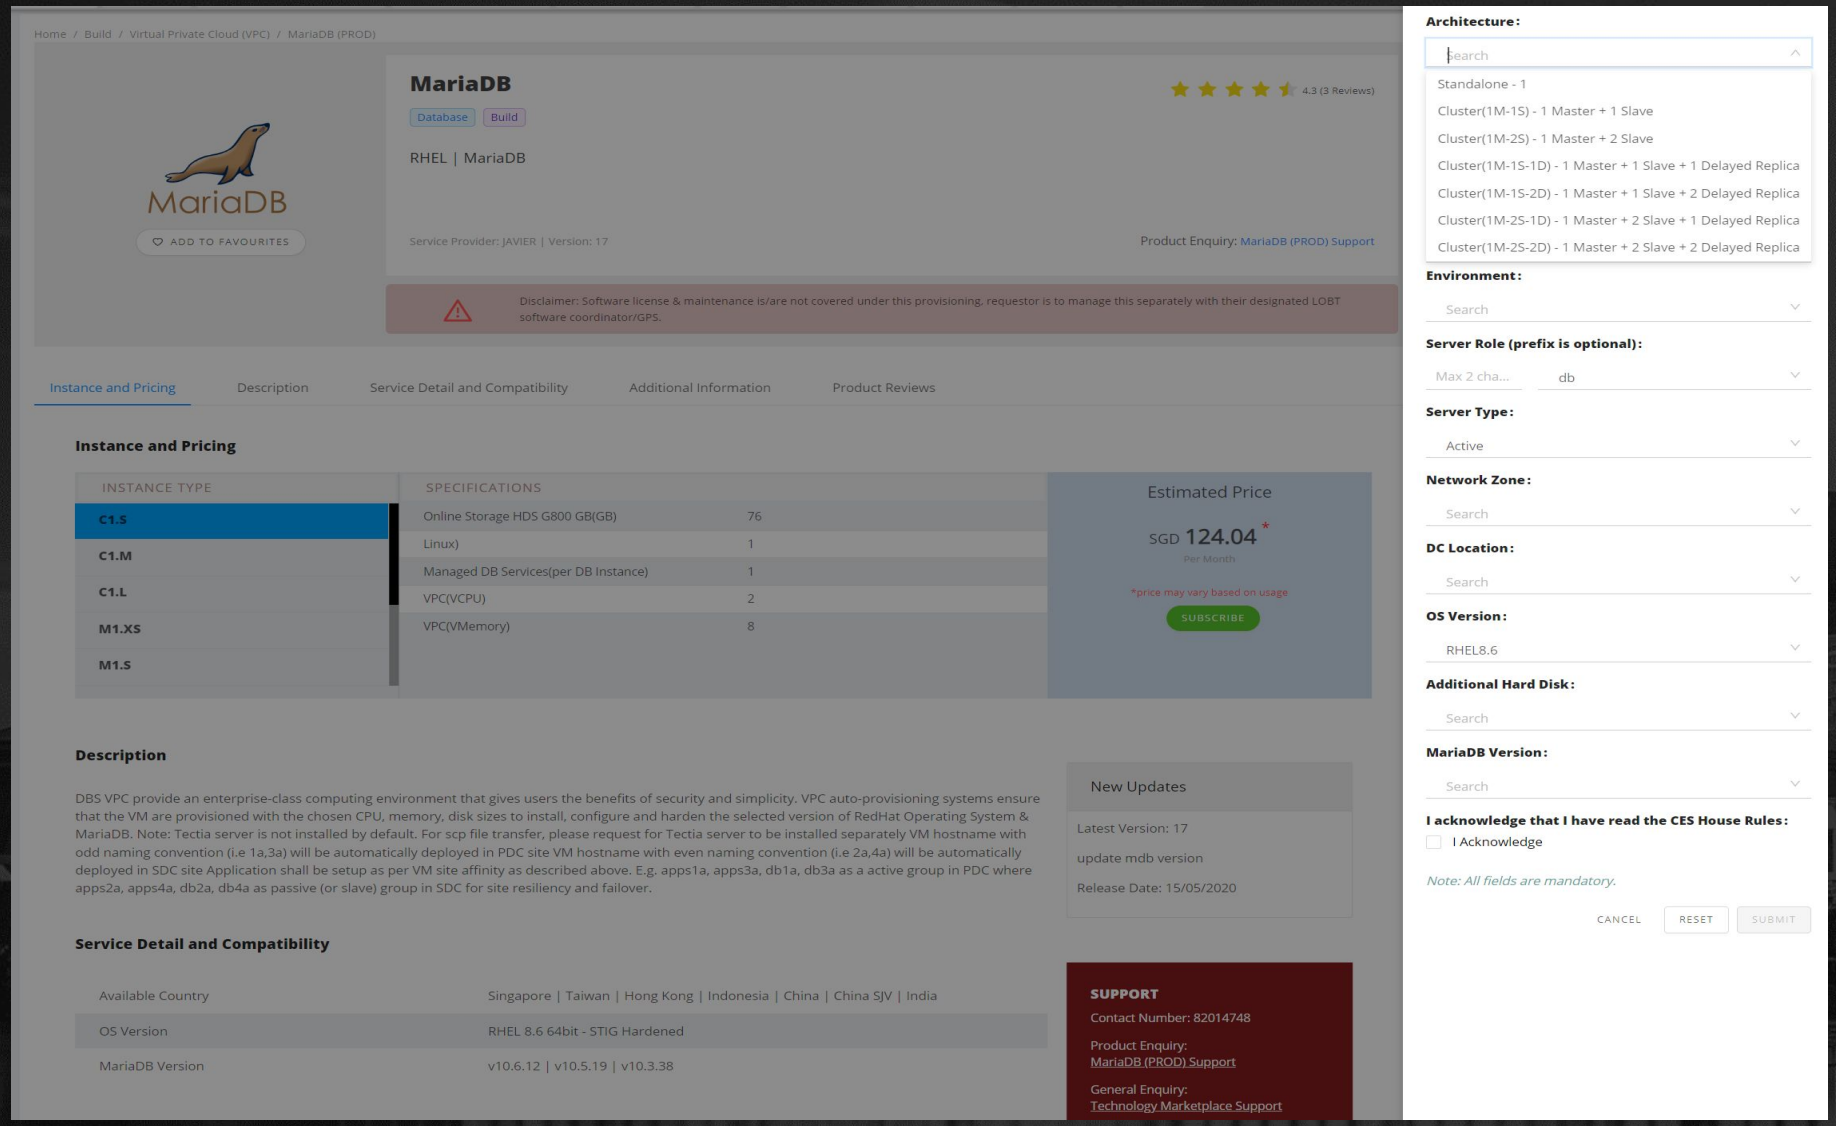

1. Provisioning a database vm is a matter of 5-10 min.
2. Default security hardening is applied.
3. Default maintenance jobs are created and scheduled.
4. Linked with automated portals.
5. Application team can play around anytime they like (provision and decom made easy).

Day 1 - Remaining consideration for application database design – self service portal

To operate MariaDB as a service , self service portal (RunDB) is built.

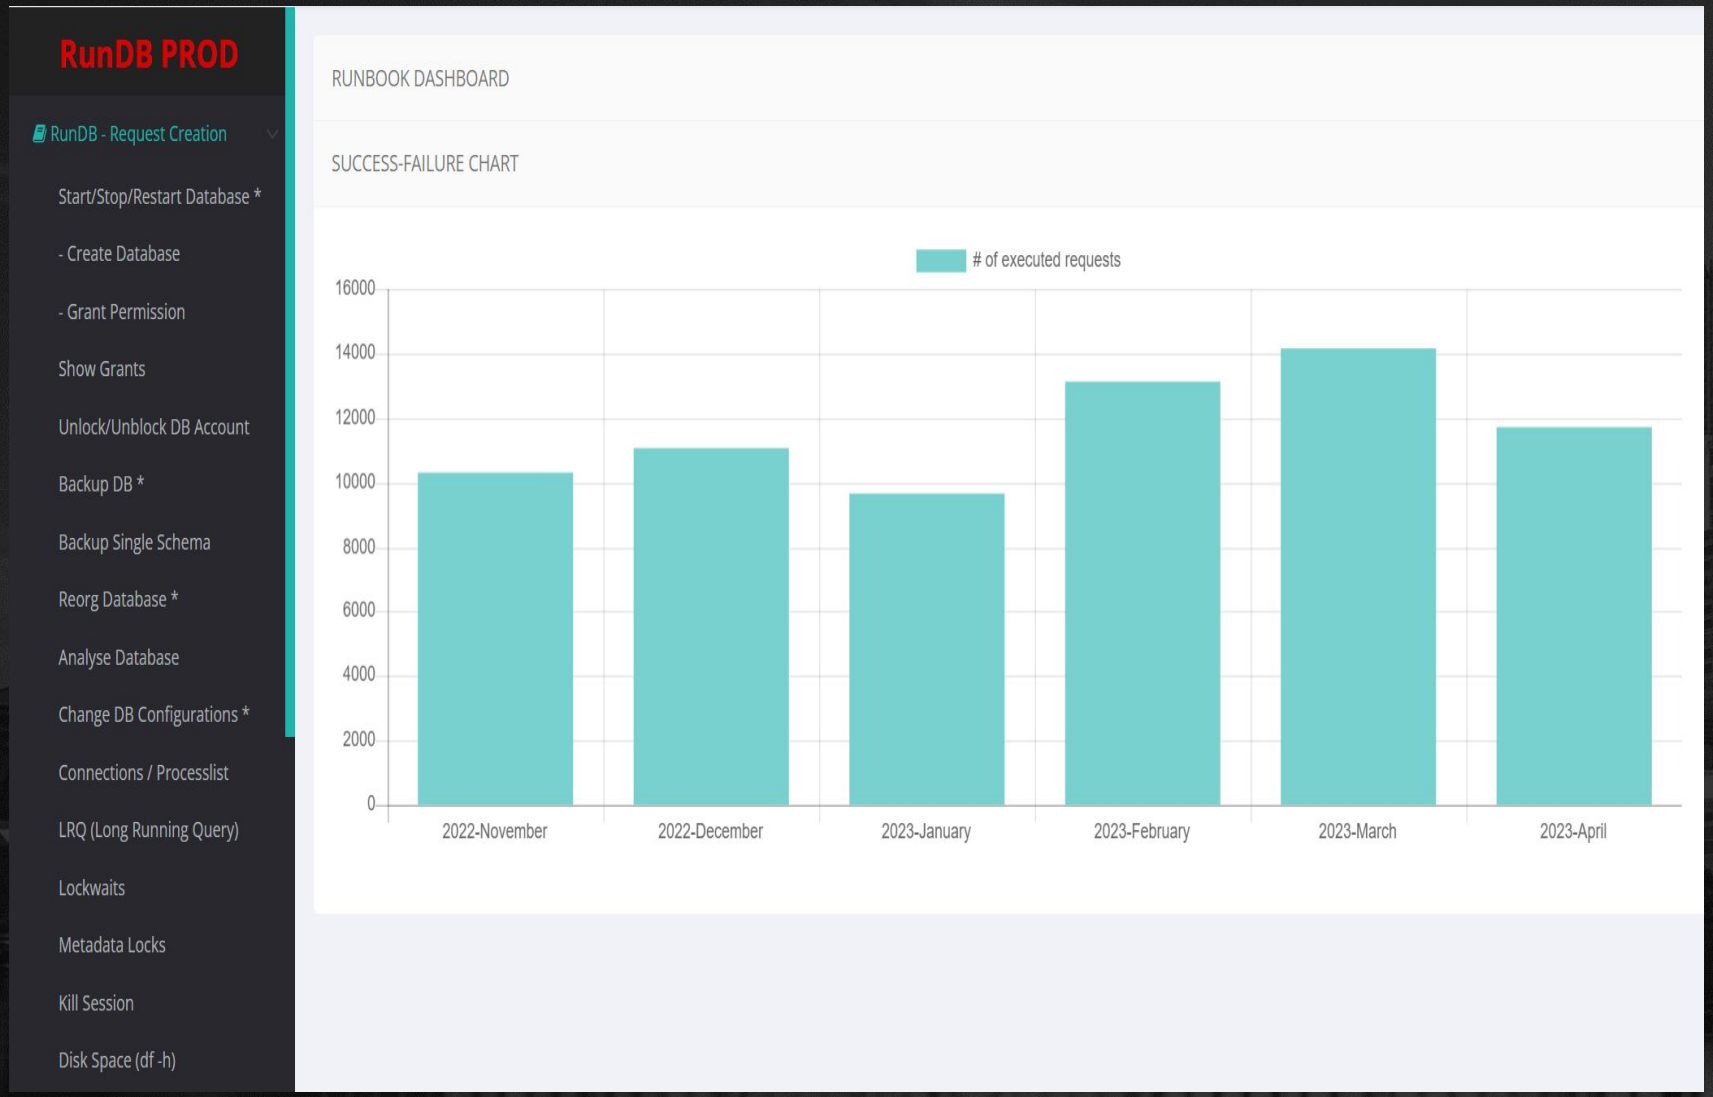

Self service portal will help on,

1. Users to self perform common DBA oriented tasks.
2. DBAs/L1 support teams to easily perform L1 tasks.
3. Avoid request raising for repetitive tasks.
4. Perform tasks without shell access.
5. Perform tasks to avoid human errors.
6. Rollout changes instantly (with maker/checker) to avoid waiting time.

Day 2 - Remaining consideration for application database design – operation dashboard

To operate MariaDB as a critical infra, an operation dashboard is built.

LOB	App Code	Hostname	DB Status	LRQ	Lock W	File sys	Full Bkp	Incr Bkp	Invald Vw	Analyze	ISCD	Pri key	Uptime	Version	EOL	Patch	DB Type
CT	GHKLDR	x01gghkldb1a											40 Days	10.6.9			Master
CBGT	H2OO	x01gh2oodb3a											41 Days	10.5.16			Slave
CBGT	SBIE	x01gsbiedb2a											55 Days	10.6.11			Master
CBGT	SBIE	x01gsbiedb4a											70 Days	10.6.11			Standalone
FR	ISIN	x01gisindb3a											5 Days	10.5.8			Slave
ITT	SPE	x01gspedb1a											47 Days	10.6.11			Slave
CT	GHKLDR	x01gghkldb2a											40 Days	10.6.9			Slave
MOT	ALM	x01galmdb1a											2 Days	10.5.13			Standalone
CBGT	SBIE	x01gsbiedb3a											55 Days	10.6.11			Slave
ITT	SPE	x01gspedb2a											47 Days	10.6.11			Master
IBGT	AOS	x01gaoshkdb1a											145 Days	10.5.8			Slave
MOT	RAMM	x01grammdb1a											13 Days	10.6.12			Standalone
FR	IWF	x01giwfadb1a											60 Days	10.5.16			Slave
MOT	LCRS	x01glcrsdb1a											2 Days	10.6.12			Standalone
IBGT	AOS	x01gaoshkdb2a											145 Days	10.5.8			Master
CT	OCCCIN	x01gocindb1a											6 Days	10.6.11			Master
CT	OCCCIN	x01gocindb2a											6 Days	10.6.11			Slave
CT	BFRS	x06gbfrsdb1a											134 Days	10.5.15			Standalone
FR	IWF	x01giwfadb2a											60 Days	10.5.16			Master
IBGT	RPBS	x01grpbspmdb1a											159 Days	10.6.8			Master
IBGT	SMEP	x01gi3smdb1a											166 Days	10.6.8			Slave
FR	OFCRMTW	x01gcrtwdb3a											5 Days	10.5.8			Slave
FR	PCT-IN	x06gpcindb1a											166 Days	10.6.8			Master
ITT	TRAX-ID	x07gtriddb1a											25 Days	10.6.11			Master
FR	IWF	x01giwfdb1a											60 Days	10.5.8			Master
IBGT	SMEP	x01gi3smdb2a											166 Days	10.6.8			Master
FR	PCT-IN	x06gpcindb2a											166 Days	10.6.8			Slave
CT	RPS	x01grpsdb1a											22 Days	10.6.11			Standalone
ITT	TRAX-ID	x07gtriddb2a											25 Days	10.6.11			Slave
CBGT	TDIN	x01glvbddb2a											25 Days	10.6.11			Slave
CE	BPCP	x11gbpcpdb1a											56 Days	10.3.37			Standalone
FR	IWF	x01giwfdb2a											40 Days	10.5.8			Slave
IBGT	RPBS	x01grpbspmdb2a											32 Days	10.6.8			Slave
FR	PCT-IN	x06gpcindb3a											166 Days	10.6.8			Slave
FR	IWF	x01giwfdb3a											60 Days	10.5.8			Slave
CT	CNGW	x05gcngwdb1a											8 Days	10.6.12			Master

Total : 1470 | Full failed : 1 | Incr failed : 3 | LRQ : 9 | Lockwait : 1 | Uptime : 18 | Pri ky issue : 1067 | File sys issue (Warning) : 39 | Repl lag : 2 | Invalid view : 29 | Analyze fail : 102 |

Dashboard to detect the following

1. Availability and issue detection
2. Long running queries
3. Lock waits
4. Database backup
5. Security hardening
6. File system capacity utilisation
7. Patch Status
8. EOS/EOL Status

Day 2 - Remaining consideration for application database design – operation dashboard

To operate MariaDB as a critical infra, database patch dashboard is built.

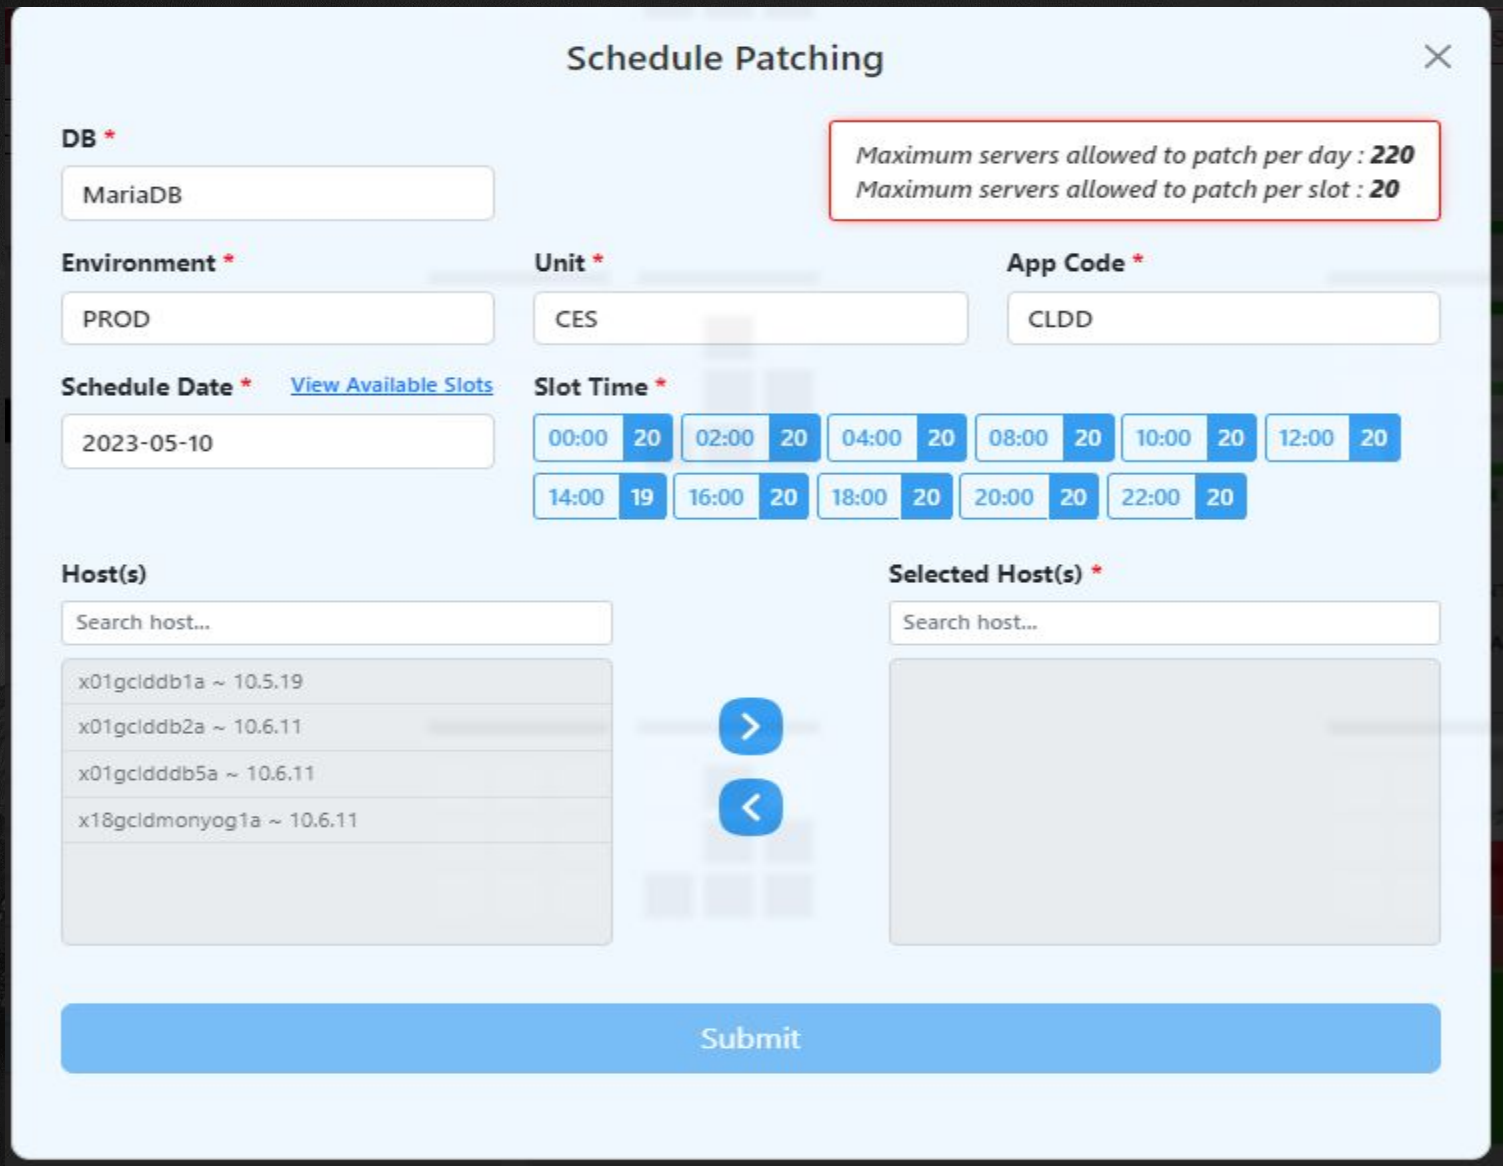

Dashboard to serve the following

1. Database Patch status
2. Patch Reporting
3. Patch Scheduler
4. Patch Tracker

Day 2 - Remaining consideration for application database design – operation dashboard

To operate MariaDB as a critical infra, an operation dashboard is built.

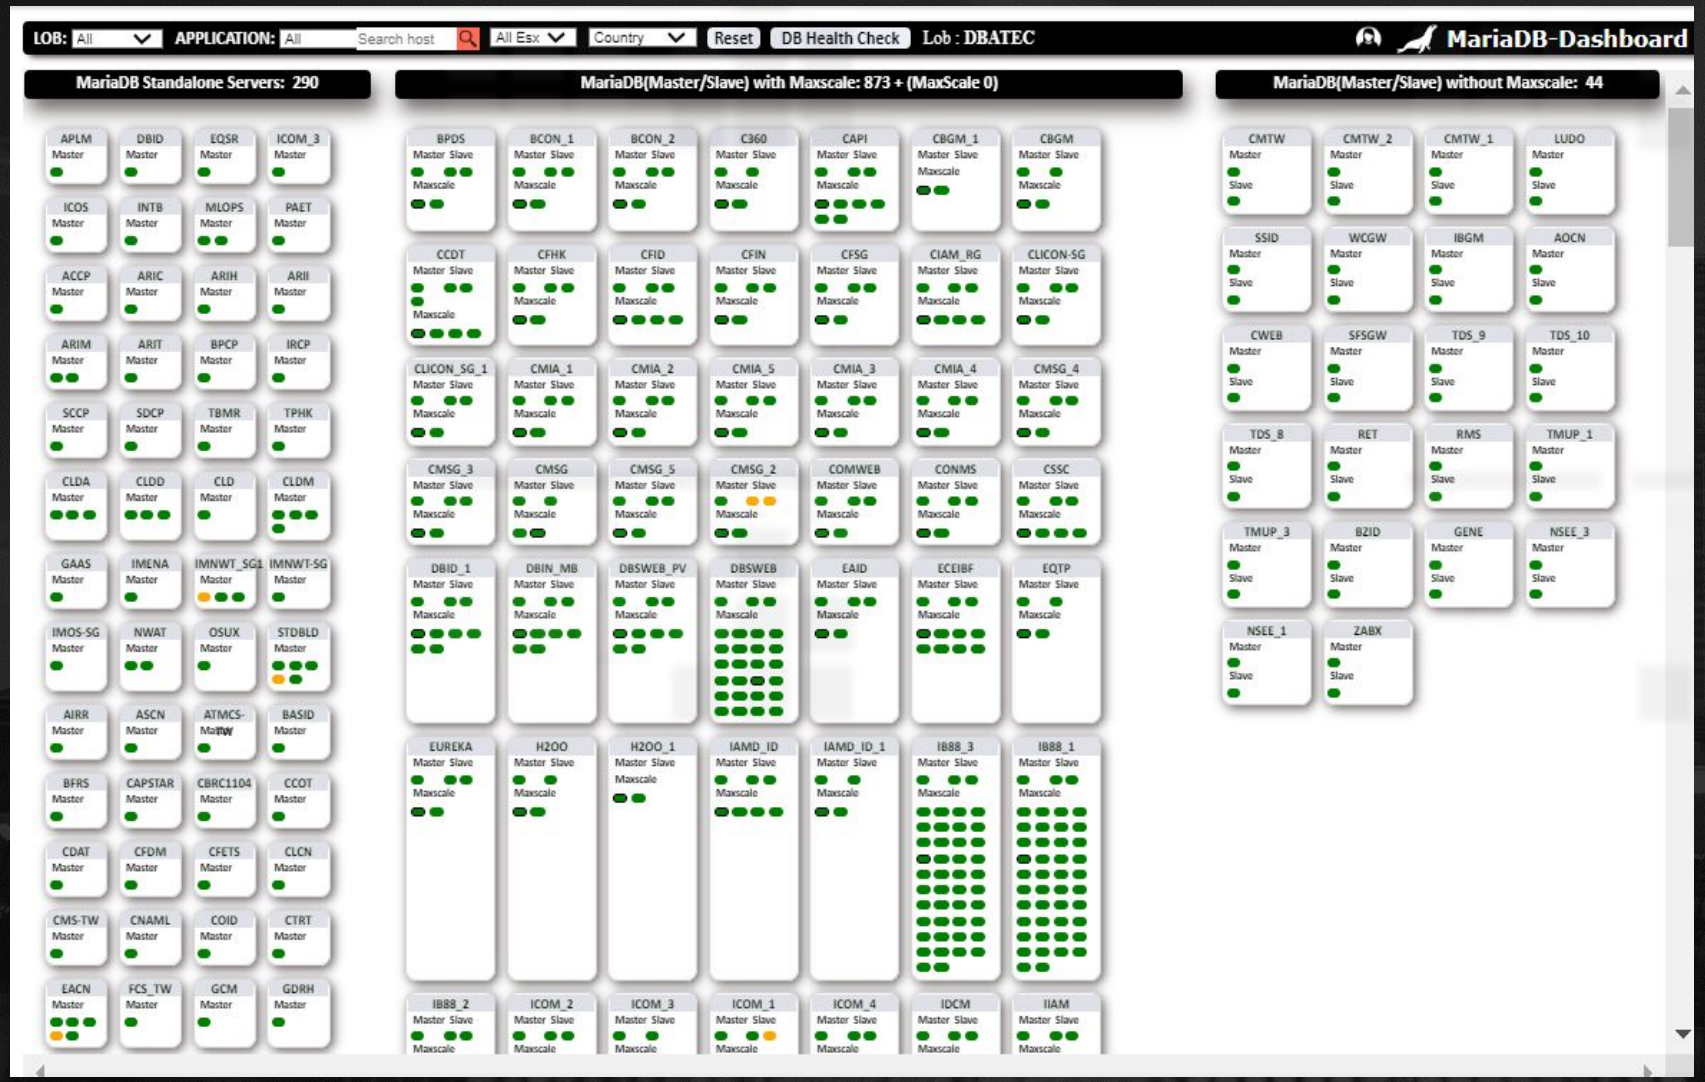

Dashboard to detect the following

1. Availability and issue detection
2. MariaDB Cluster status
3. MariaDB Replication status
4. Database backup
5. File system capacity utilisation

DBS Virtual Private Cloud is engineered for speed, cost efficiency and resiliency

		Before	Now	Industry Average
Cost Efficiency 	**Drive towards cloud** *(% of workload virtualised)*	< 3%	> 99%	50%
	Efficiency of cloud *(# of workload on single hardware)*	1:5 – 1:8	>1:100	1:20^{+}
	Data centre footprint *(in sqft)*	28,000	7,500	67,500
Resiliency & Scalability 	**Resiliency of cloud infra** *(Ability to swing workload to alternate site)*	Passive DR	Active-active	Passive DR
	Capacity to scale *(buffer in data centre site for future growth)*	< 10%	10x	Maybe
Change Cadence 	**System provisioning** *(Provision infra for new applications)*	4 months	Minutes	Minutes
	Aggressive automation *(System admin to workload ratio)*	< 1:50	1:1,200 (OS) 1:650 (DBA)	1:272 (OS) 1:154 (DBA)

What next ?

1. MariaDB, Maxscale on Containers/Kubernetes
2. MariaDB 23.x
3. Maxscale 23.x
4. Ransomware protection
5. Seamless downgrade

etc

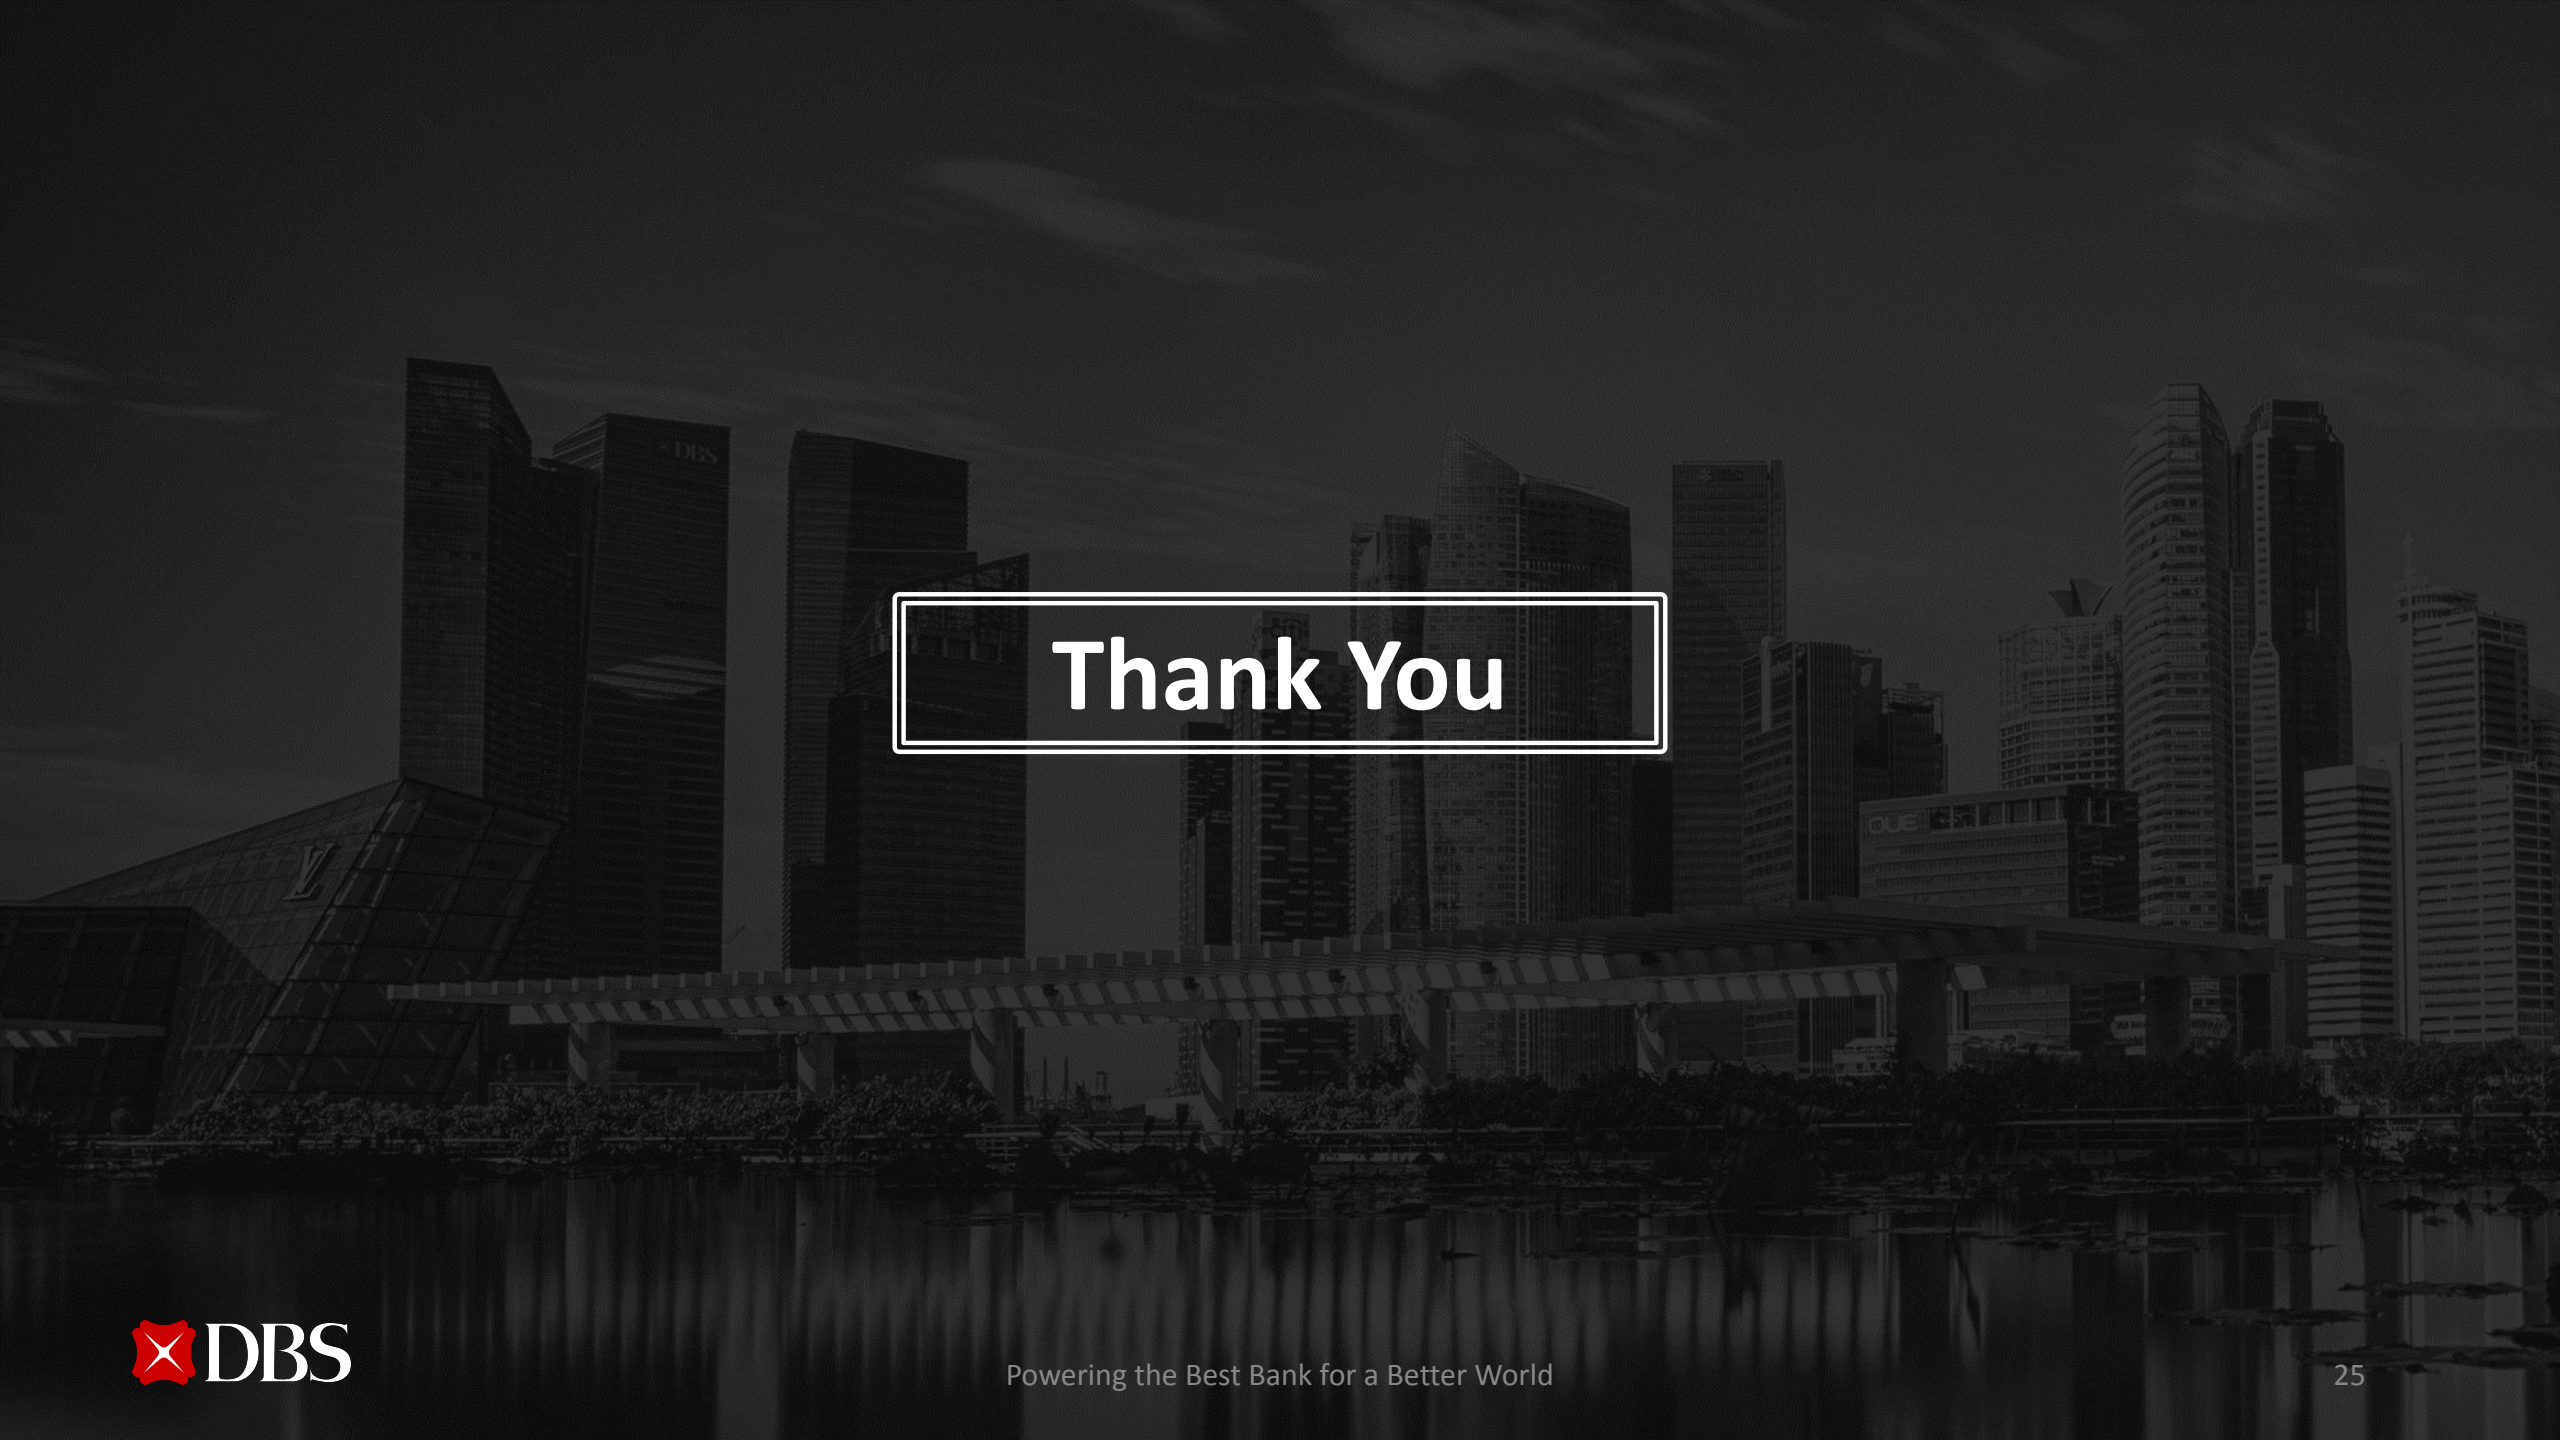
Thank You