Database Platform Selection Tool

Mansi Narula

November 2014

Offline

Online

Volume	Velocity	Variety
5-15 Years of History 4PB Largest Table 16M Analytic Queries 14K Users 4M Batch Queries 900K Ad Hoc Queries	37PB Read 3PB Write 16+TB/day Semi-Structured Data 36 TB/hour x-Platform Data Transfers	3.5PB+ Structured Data 10PB Semi-Structured Data (80% compressed) 10K+ Name/Value Pairs
6PB Consumed 2TB Daily Average 700M Active Items 300M Active Site Users 8K Average Application Connections/DB	200B+ eBay Queries/day 4K eBay Batch Runs/day 25GB/sec Peak Site Traffic	800+ Oracle Instances 300+ MongoDB Nodes 300+ MySql Nodes 200+ Cassandra Nodes

One Size Doesn't Fit All

Database Platforms

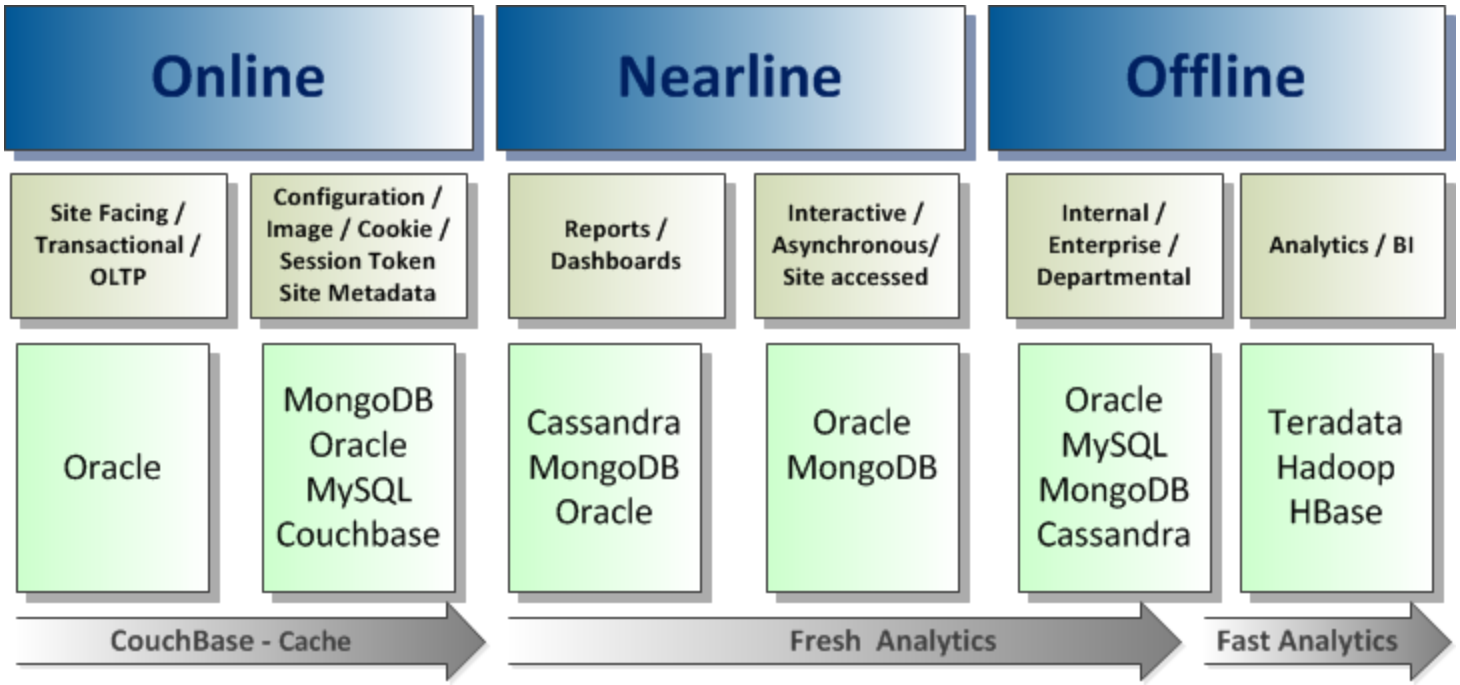

Database Platform Selection program

Objectives

- Simplify Database Platform Selection Process
- Minimize Data Architecture team Involvement
- What to Choose When ?
- Preliminary Qualification
- Followed by Database Platform Selection Card

Scope of database platform selection program

- Database platform selection process
  - What Database platform to use and when
- Database platforms for online, nearline and offline.
  - Oracle
  - MySQL
  - MongoDB
  - Cassandra
  - Hadoop/ Hbase
  - CouchBase
  - Teradata
- Data Architecture Strategy
  - Guidelines and Definitions
  - Database platform selection flowchart and diagrams
  - Database Platform selection scorecard
- Adoption and Communication program
  - Brownbag sessions
  - DA WIKI pages, Tools, Schema
- Database Platform Selection self-service tool (DBPS)
  - Create simple UI based on the flowchart.
  - Publish application link on DA Tools (http://tools) and Schema (http://schema)
- Measurements of success
  - Adoption by PM and PD teams.
  - Decreasing time of DA team involvement into Database Platform Selection process

Oracle Database Platform

- Relational Database
- Fully Compliant to ACID
- Data Concurrency, which ensures that users can access data at the same time
- Data Consistency, ensures that each user sees a consistent view of the data
- Supports Transactions
- Robust and Secure
- Used in Majority of eBay's Site Facing Applications

MySQL (InnoDB)

- Relational Database
- Transactional Support provided by MVCC
- Row-level Locking
- Foreign Key Support
- Indexing using clustered B-tree indexes
- Online Non-blocking Backup

Mongo DB Use Cases & Management

- Suitable Mongo DB use cases:
  - Higher read to write ratio, e.g., in-memory cache
  - Highly available, high throughput, balanced scale-out reads
  - Reasonably complex data models & access patterns
  - Dev friendly ecosystem with speed & agility
  - MongoDB inbuilt sharding currently not recommended due to operational complexity
- Database management
  - Add / Upgrade slave or add backup nodes
  - Upgrade server software
  - Upgrade to external solid state storages

Cassandra

- Higher write to read ratio
- Mix workload if working set fits in memory or leverages SSD for tight read latency
- Always available for both read and write without Single Point of Failure (SPOF)
- Out of the box support for active-active multiple data centers
- Local latency requirement from App server to DB server.
- Column Family oriented data structure
- Need real-time or near real-time aggregations.

Why Hbase ?

- Database platform that can store petabytes of data efficiently in a cost-effective way
- Provides random data access in close to real time
- Scalable Reads and Writes
- No fixed schema. Storage evolves with application
- Full consistency per Operation

Task	RDBMS	Hbase
Data Layout	Row-oriented	Sparse Column Store
Transactions	Yes	Single Row Only
Query Language	SQL	APIs*
Scalable	Not natively	Yes
Max Data Size	TBs per node	PB+
Read/write throughput limits	Thousands of queries/second	Thousands of queries/second/node

Couchbase

- Memcached Protocol enabled highly scalable & available distributed cache
- Recommended as a pure Key-Value Cache only
- Extremely fast and consistent Key-Value pair cache
- Used for Session and Token Store—requires fast read and write
- Automatically Sharded 1024 ways, allowing excellent scale out

Teradata

- RDBMS, based on parallelism & shared nothing architecture
- Primarily Used in data warehousing applications
- Fully compliant to ACID principles
- Handles Massive dataset
- Supports Joins and Complex queries
- Shared nothing can scale it horizontally within one database instance to handle increase in data volume, increase in number of users, increase in number of objects.

Database Platform Selection program

Database Platform Selection process

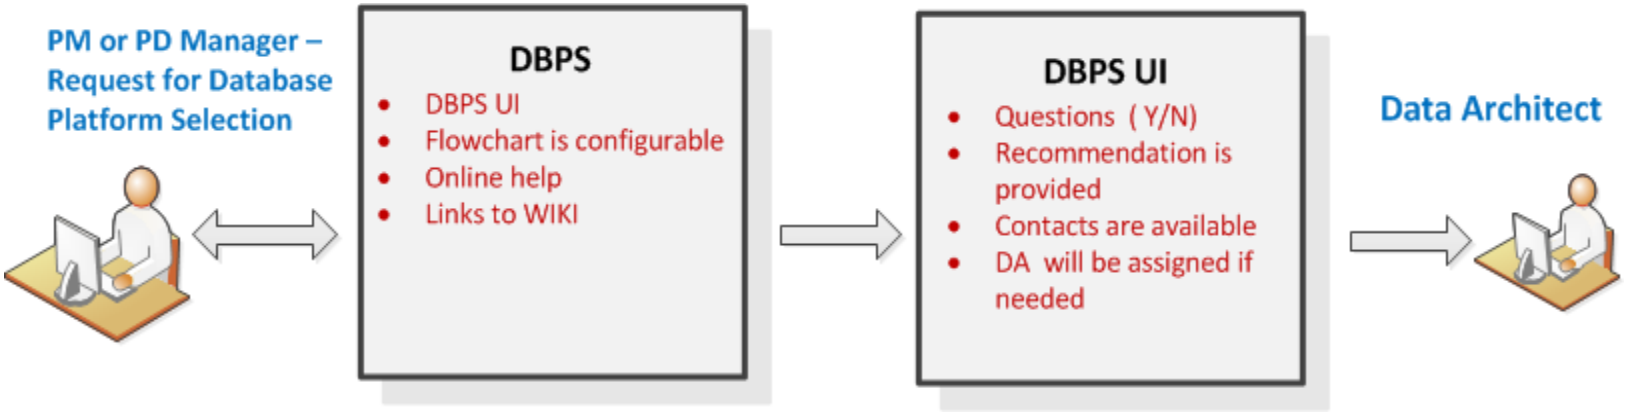

- PD / PM team provide basic information about the application .
- The Database Platform Selection self service tool provides recommendation based on provided initial information and flowchart.
- Online help and links to WIKI pages are available
- Data Architect works with PD team and other stakeholders to finalize the database platform solution for the project, if needed

eBay, Inc. database platform selection flowchart

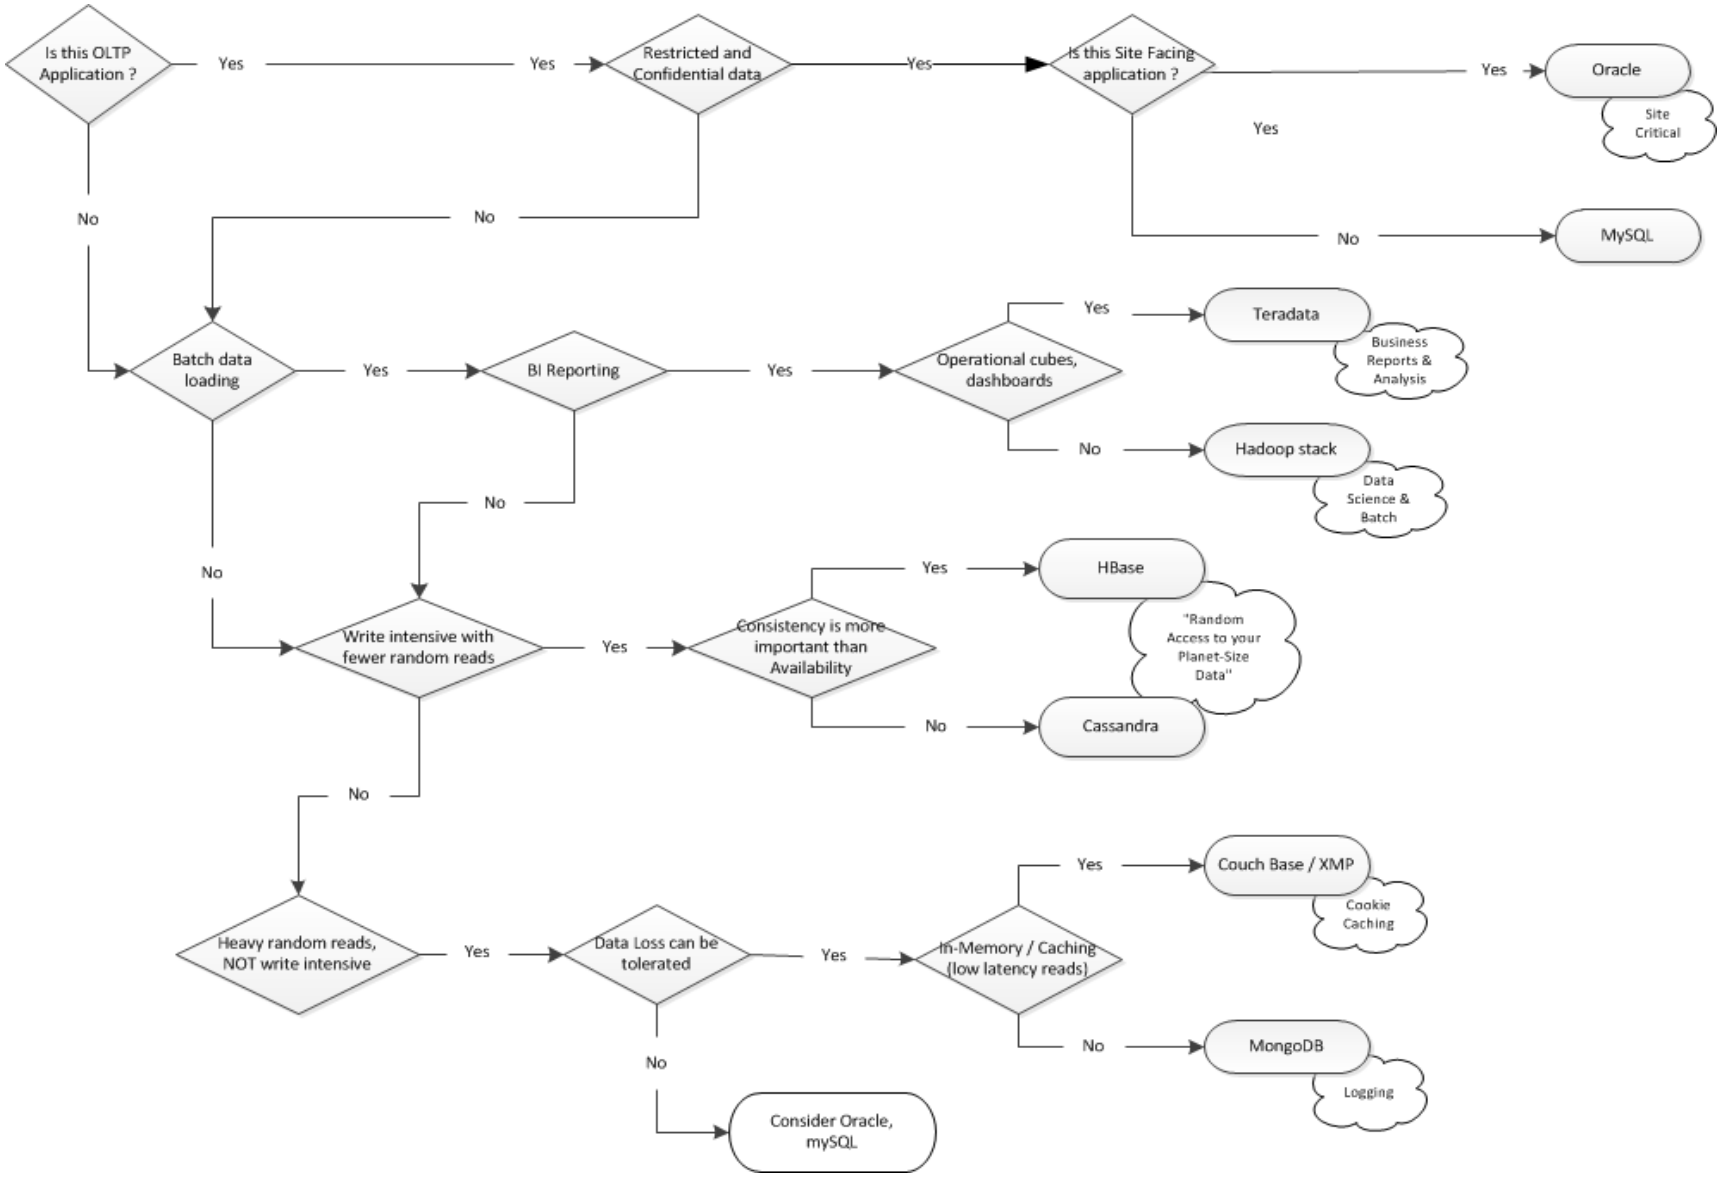

eBay, Inc. database platform selection UI

DA Tools-
https://dmaas.corp.ebay.com/dbselect

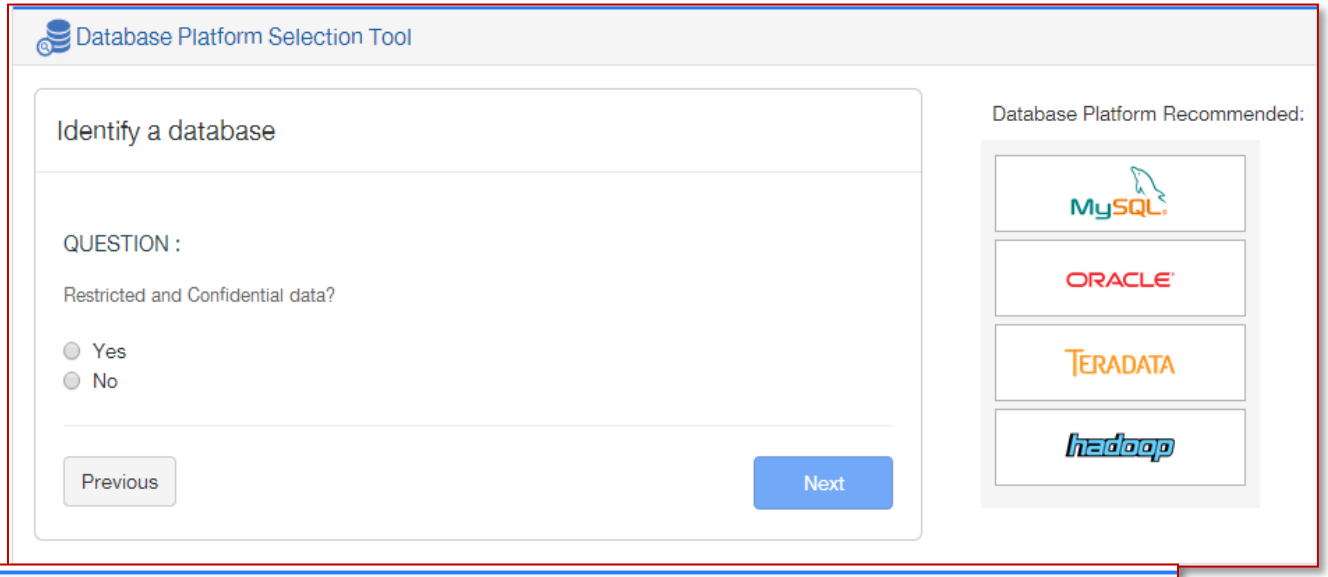

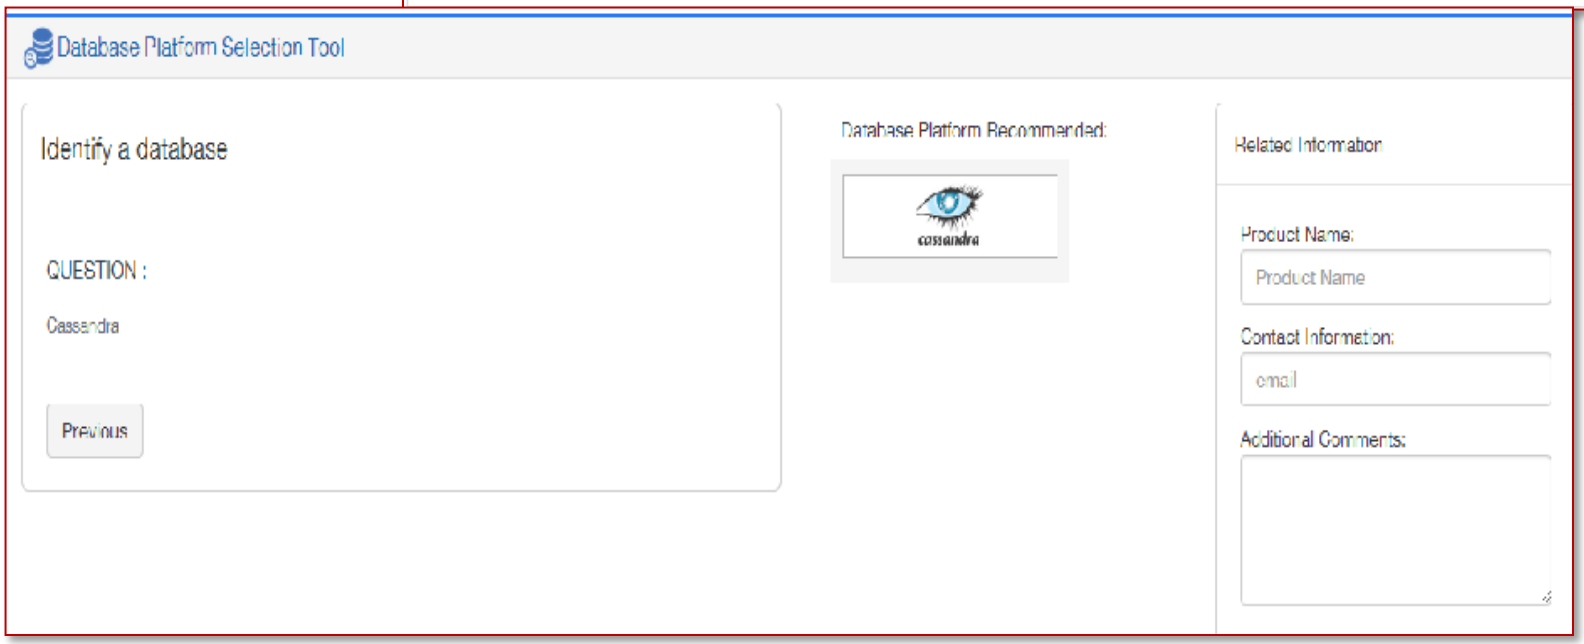

Database Selection Scorecard

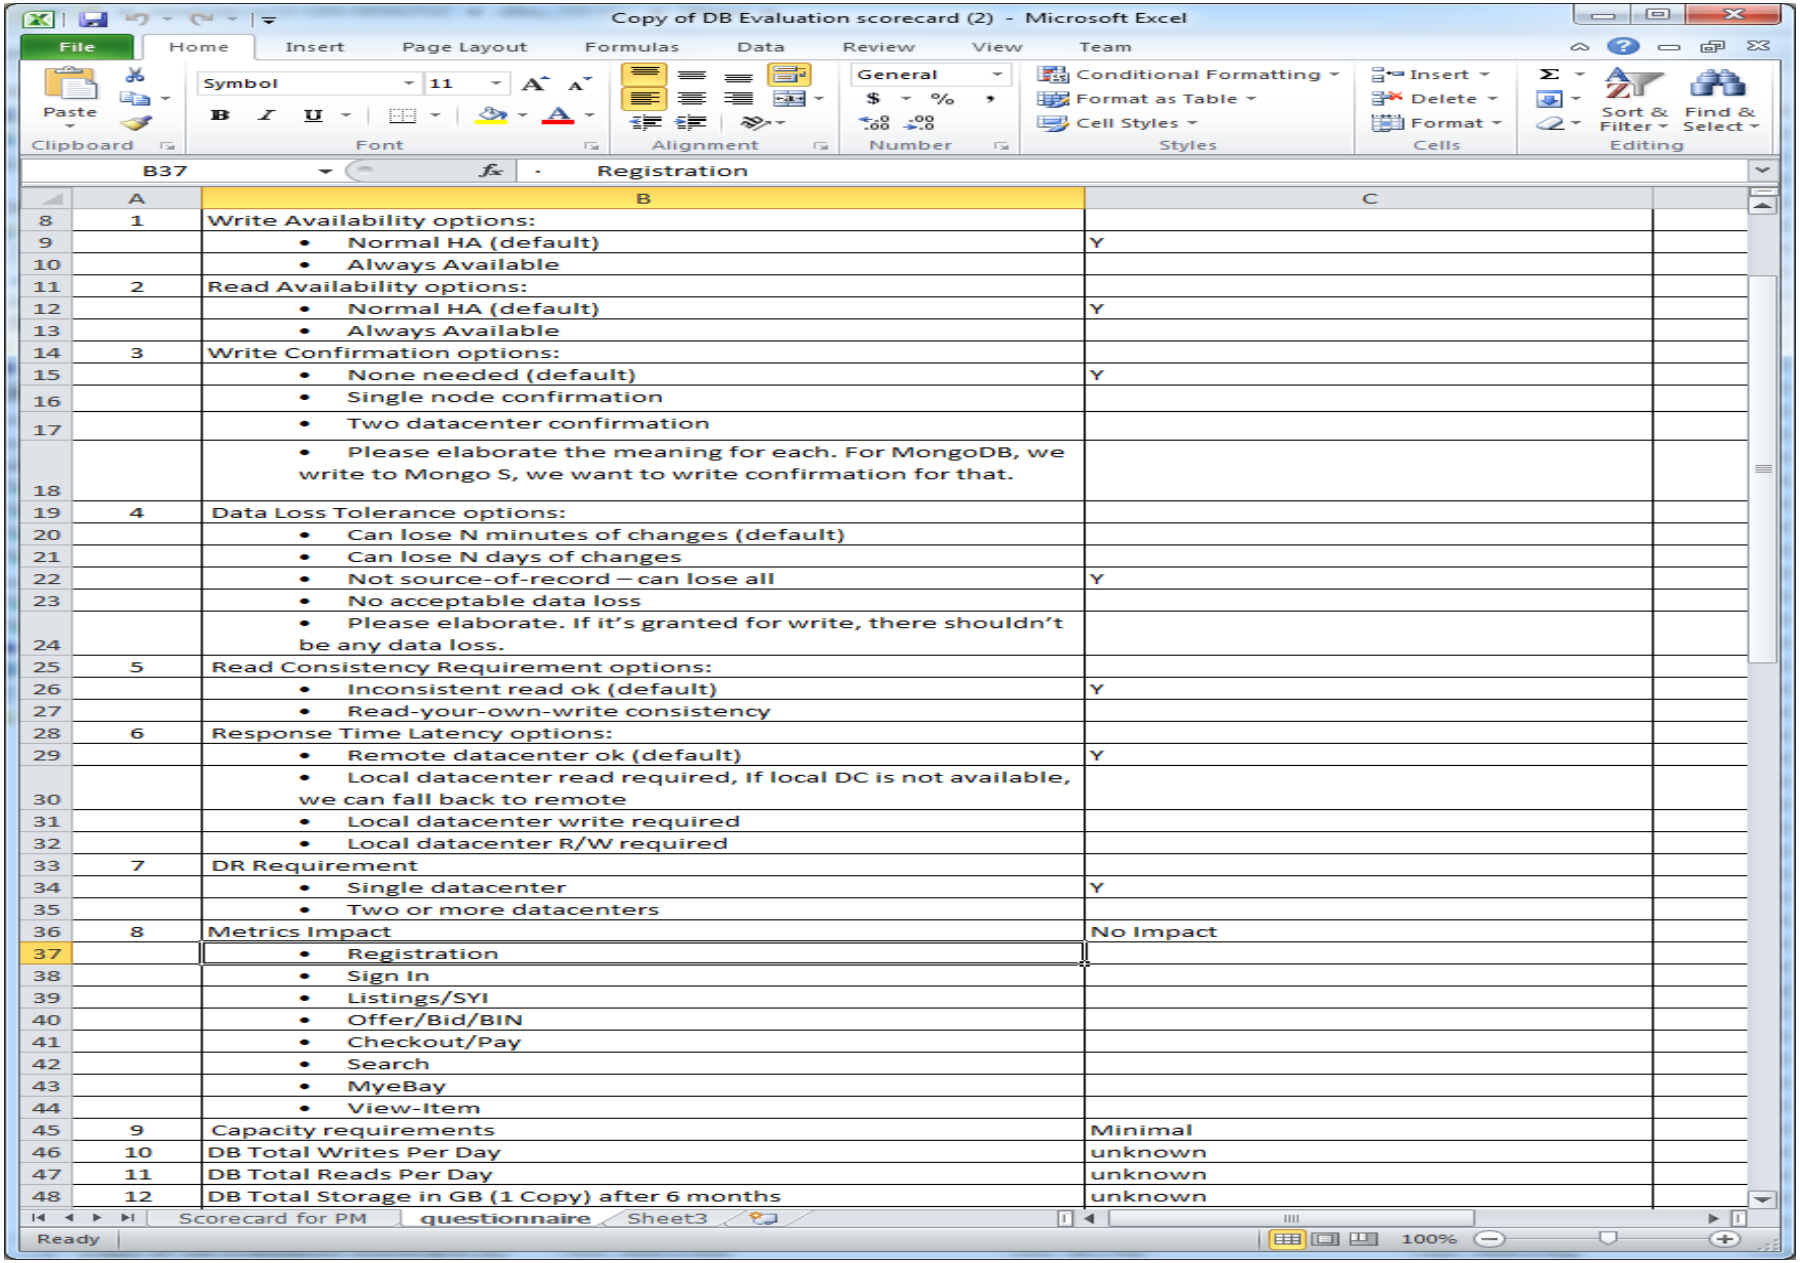

	A	B	C
8	1	Write Availability options:	
9		• Normal HA (default)	Y
10		• Always Available	
11	2	Read Availability options:	
12		• Normal HA (default)	Y
13		• Always Available	
14	3	Write Confirmation options:	
15		• None needed (default)	Y
16		• Single node confirmation	
17		• Two datacenter confirmation	
18		• Please elaborate the meaning for each. For MongoDB, we write to Mongo S, we want to write confirmation for that.	
19	4	Data Loss Tolerance options:	
20		• Can lose N minutes of changes (default)	
21		• Can lose N days of changes	
22		• Not source-of-record – can lose all	Y
23		• No acceptable data loss	
24		• Please elaborate. If it's granted for write, there shouldn't be any data loss.	
25	5	Read Consistency Requirement options:	
26		• Inconsistent read ok (default)	Y
27		• Read-your-own-write consistency	
28	6	Response Time Latency options:	
29		• Remote datacenter ok (default)	Y
30		• Local datacenter read required, If local DC is not available, we can fall back to remote	
31		• Local datacenter write required	
32		• Local datacenter R/W required	
33	7	DR Requirement	
34		• Single datacenter	Y
35		• Two or more datacenters	
36	8	Metrics Impact	No Impact
37		• Registration	
38		• Sign In	
39		• Listings/SYI	
40		• Offer/Bid/BIN	
41		• Checkout/Pay	
42		• Search	
43		• MyeBay	
44		• View-Item	
45	9	Capacity requirements	Minimal
46	10	DB Total Writes Per Day	unknown
47	11	DB Total Reads Per Day	unknown
48	12	DB Total Storage in GB (1 Copy) after 6 months	unknown

Questions ???