



May 11-13 | Orlando, Florida

2011

Orders of magnitude: Scale-Out your SQL Server Data

Mark Broadbent

@retracement

A bit more about me!

- More than 20 years in IT and more than 14 years using SQL Server. Worked at many large global corporations and SMEs such as Microsoft, Nokia, Hewlett Packard and Encyclopaedia Britannica.
- Presented at SQLBits 7 and 8
- MCITP Database Development SQL 2008
- MCITP Database Administrator SQL 2008/ 2005 and MCDBA SQL 2000
- Microsoft Certified Application Developer (C# .net)
- Microsoft Certified Systems Engineer + Internet
- Participate on #sqlhelp, MSDN Forums, Stackoverflow & Serverfault.
- Used to be active in the MS newsgroups until their demise :(
- Run the LinkedIn groups
- Blog: tenbulls.co.uk
- Linux Mint User Group <http://www.linkedin.com/groups?gid=2989801>
- SQL Server Scripting <http://www.linkedin.com/groups?gid=3033621>

Agenda

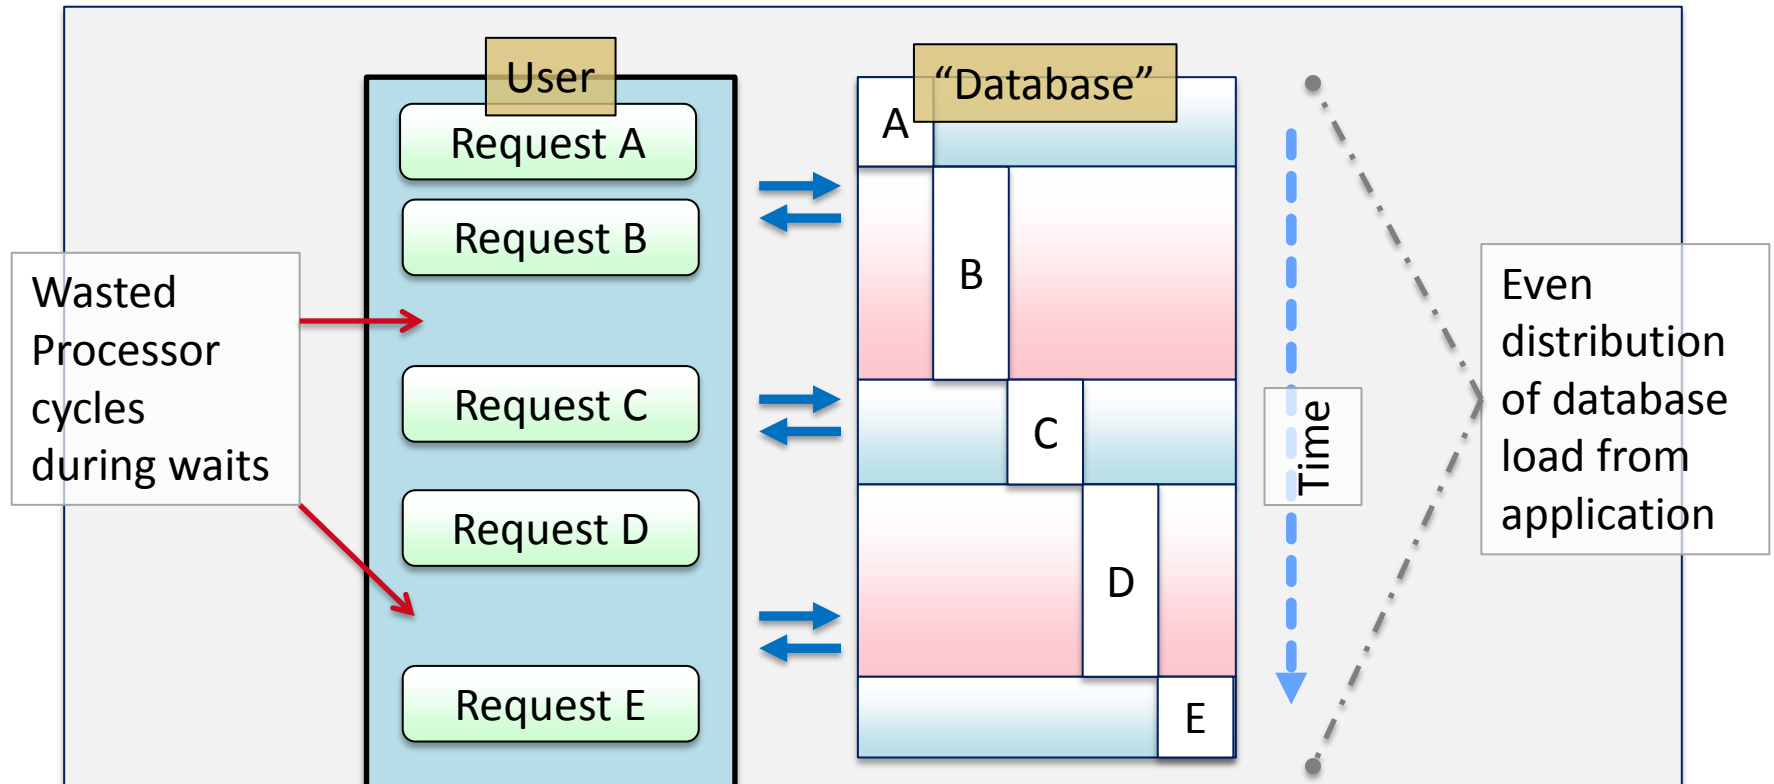
- Why Scale-Out and what should we scale?
- Benefits of Scale-Out?
- Wait!!! Before you Scale-Out, Scale-In...
- Lets look at some strategies and Scale-Out!
- Hybrid Scale-Out
- Taming the Beast

Why Scale-Out

“Gartner Group study, for example, predicted that the amount of data generated by enterprises will grow by a staggering 650 percent over the next five years. Another study sponsored by IBM found that 83 percent of the global CIOs surveyed believe that analyzing and leveraging enterprise data is critical to their companies’ long-term competitiveness.” - *Divining the Future of ERP Software*
<http://dell.to/jsoJik>

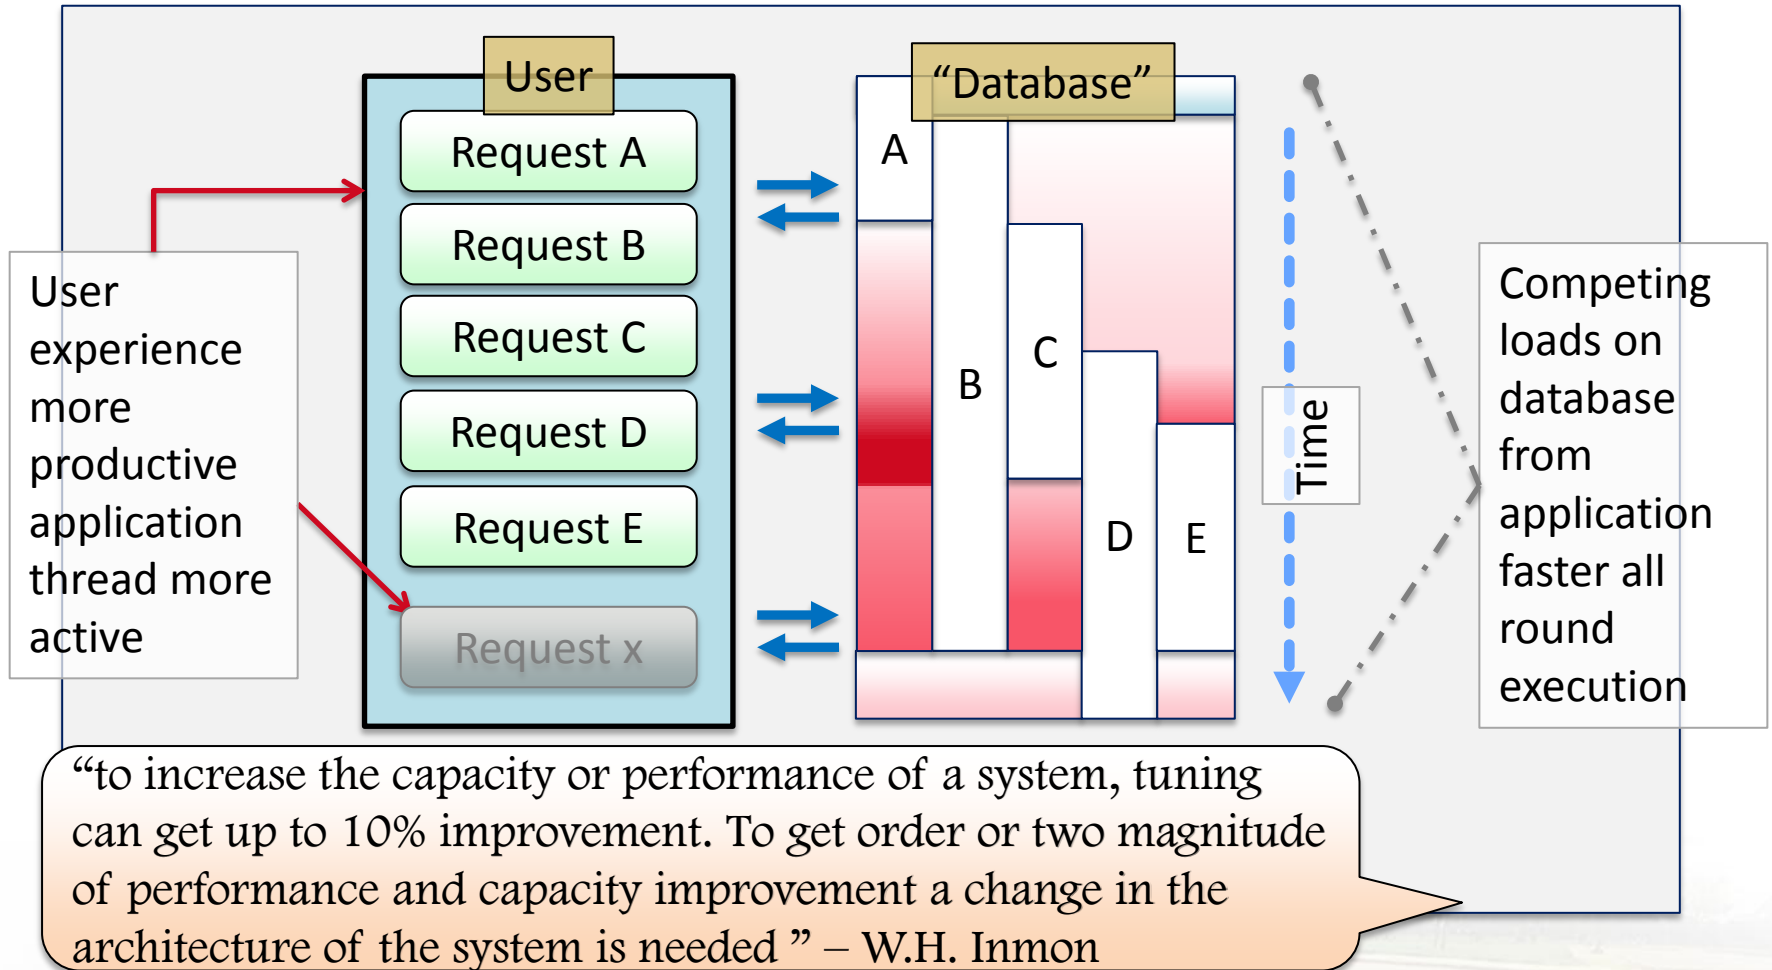
Reference: <http://bit.ly/k2kpwW> (2009 Gartner IT Infrastructure, Operations & Management Summit)

Traditional Synchronous Activity

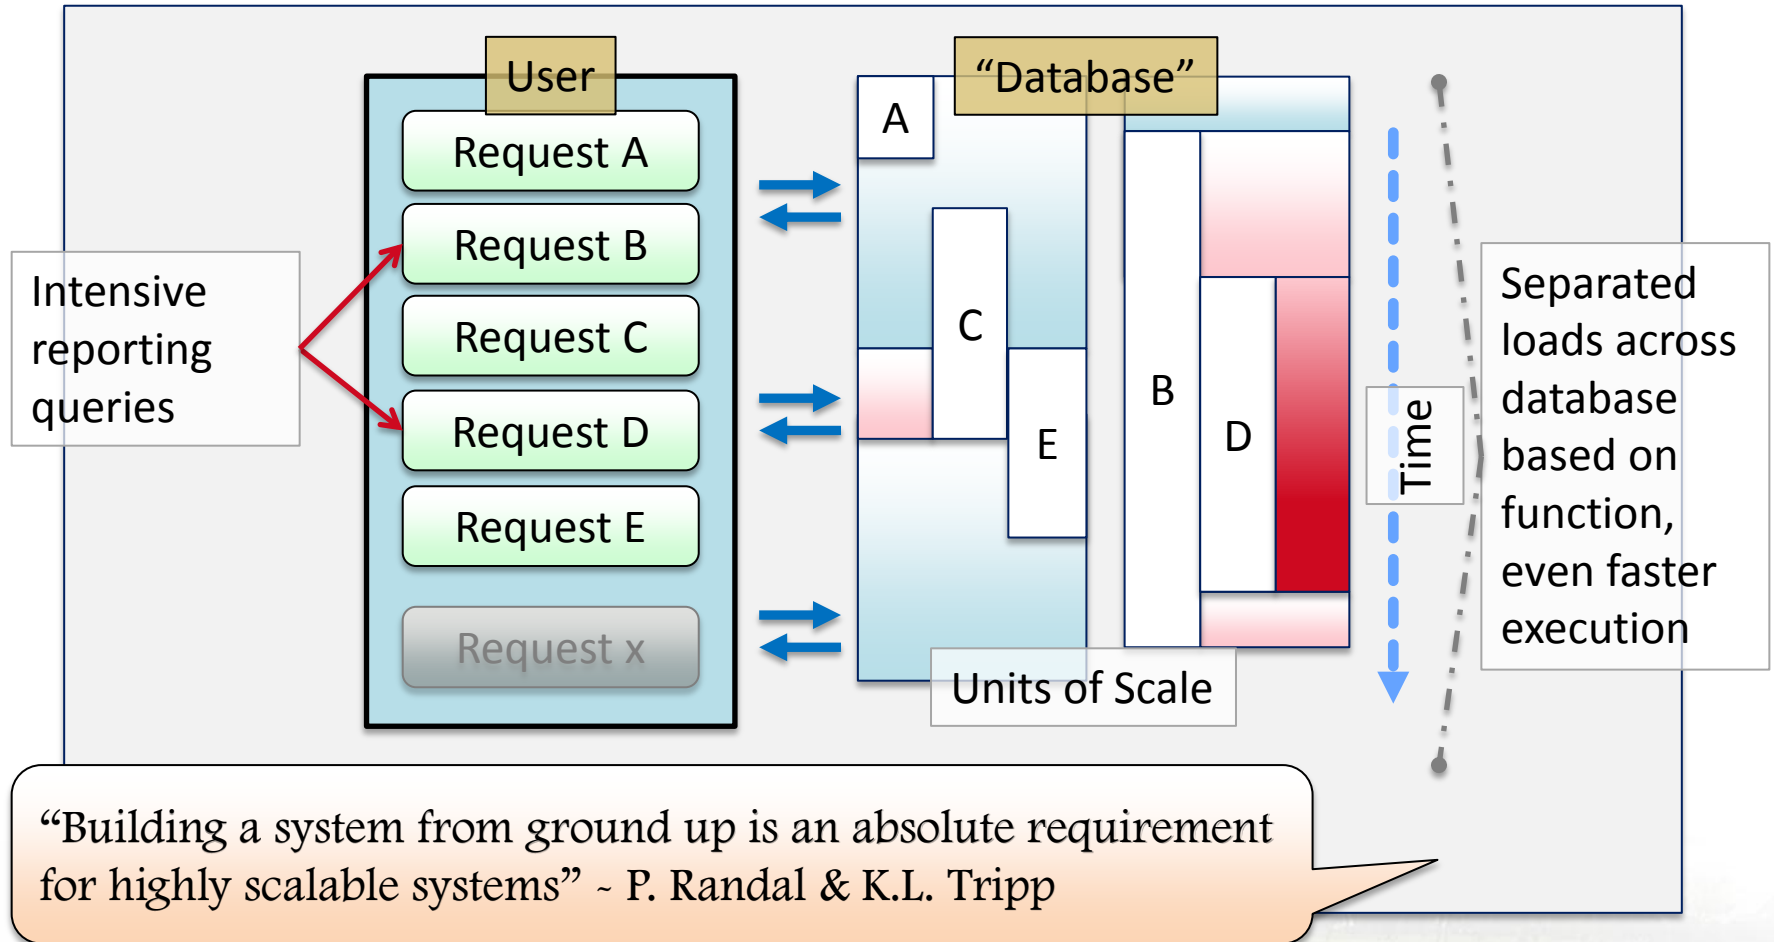


“A common error in client/server development is to prototype an application in a small two tier environment, and then scale up by simply adding more users to the server” – Lloyd Taylor

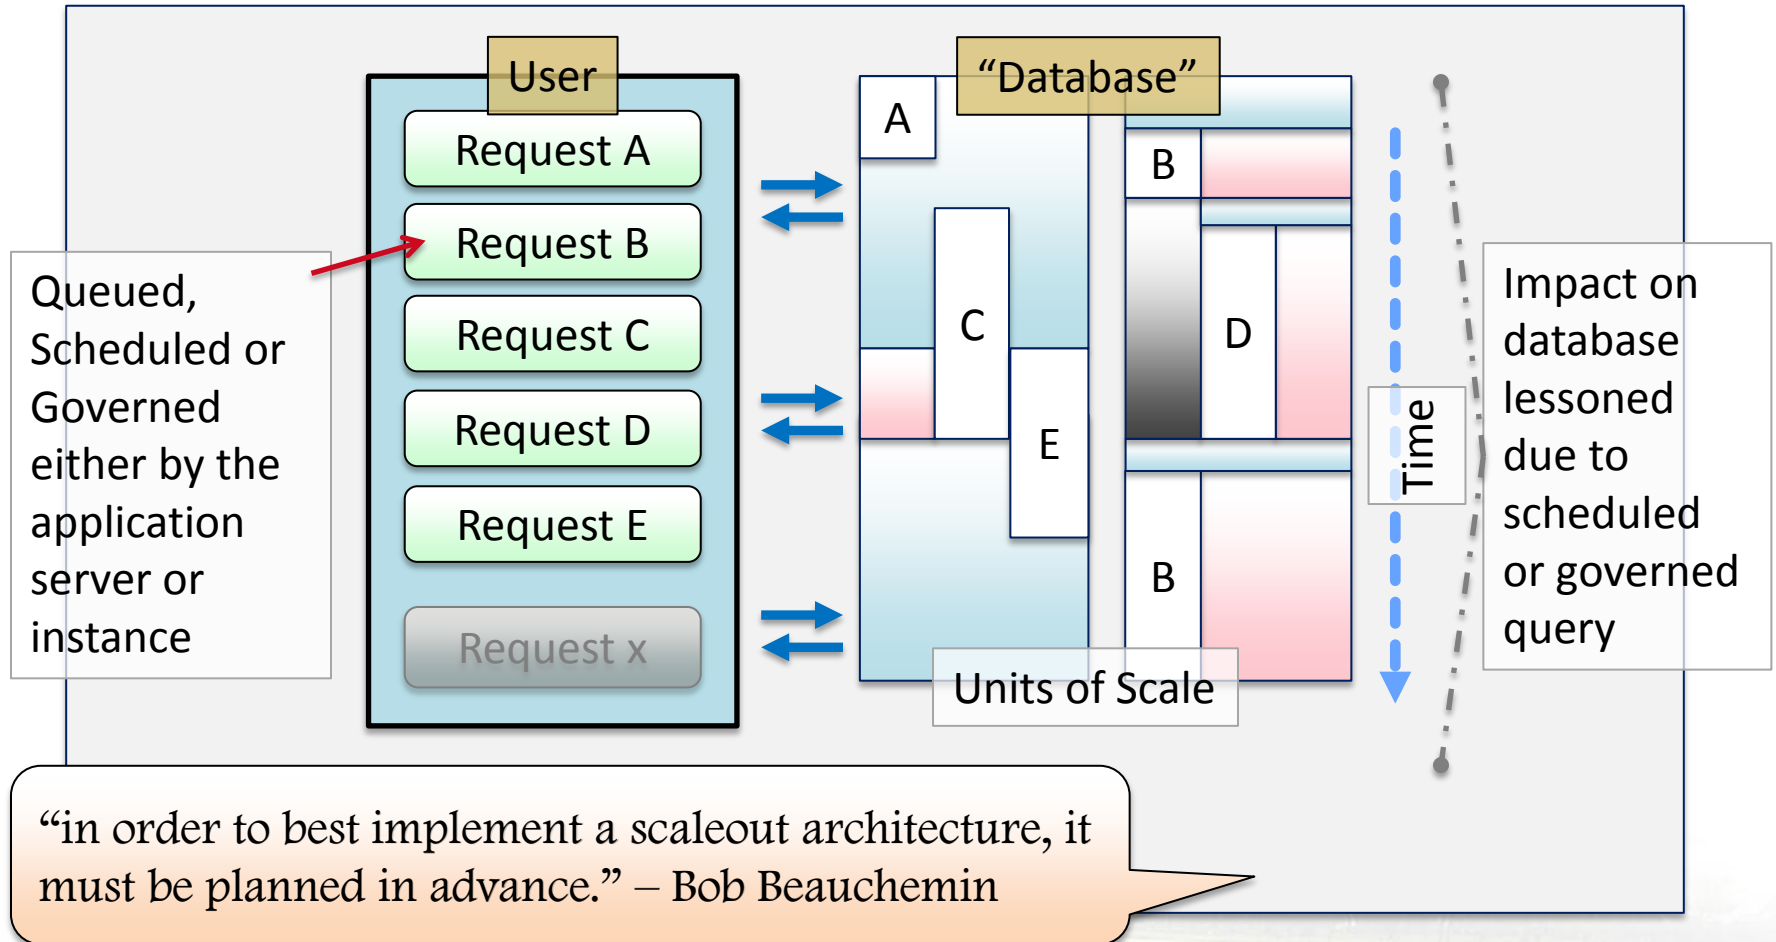
Asynchronous Activity?



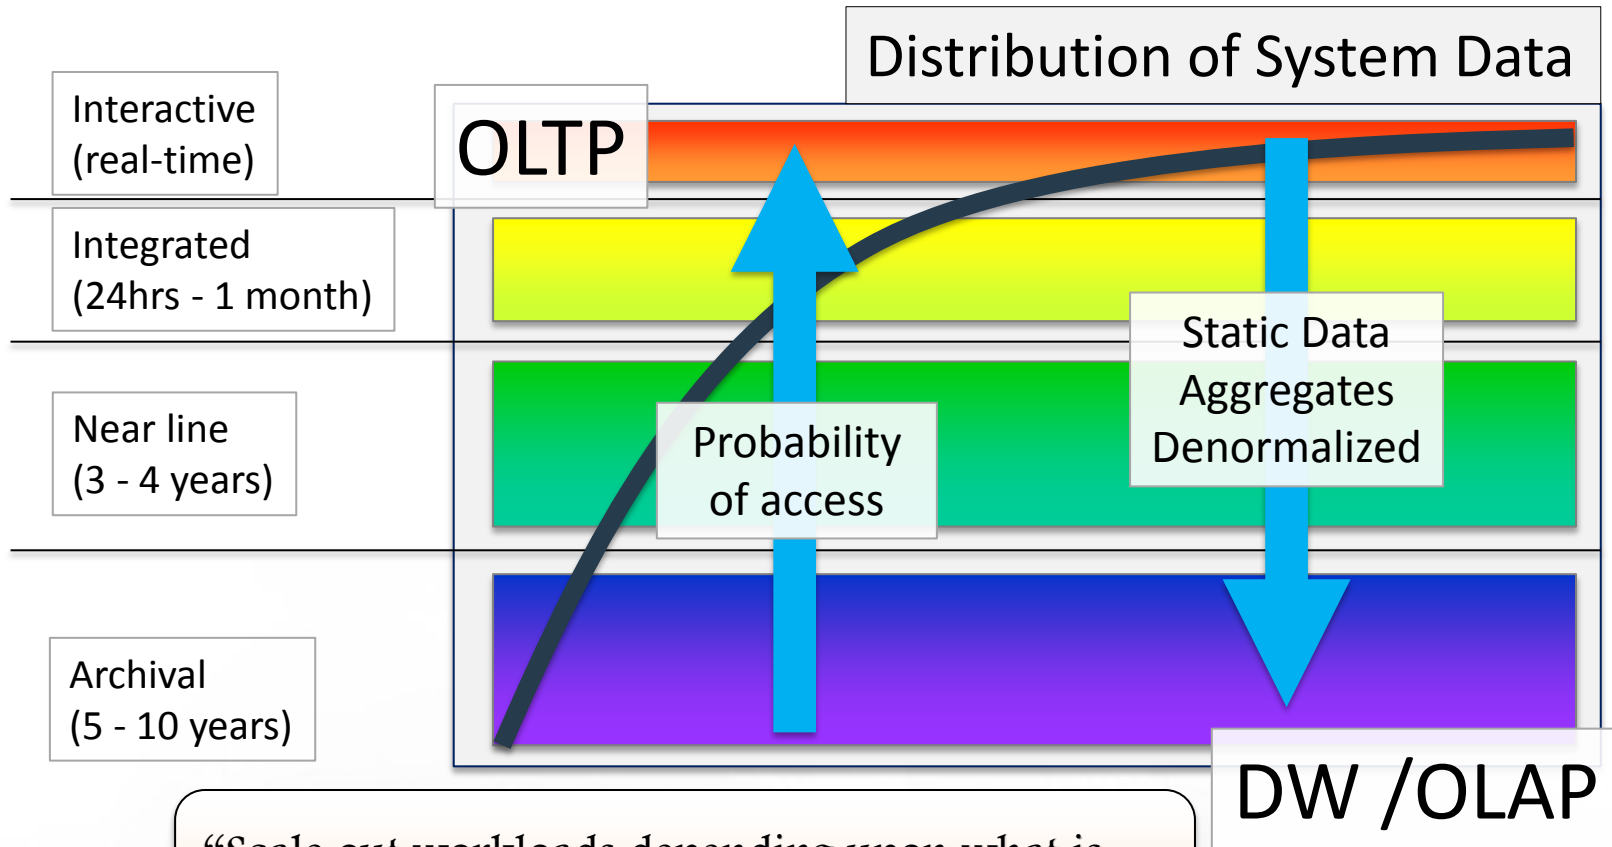
Separation of Function...



Separation of Requirements...



What is truly Mission Critical?



“Scale out workloads depending upon what is truly mission critical.” - Larry Chestnut

DEMO: Scaling for READS

Snapshots

Rolling Snapshots

Rolling Snapshots on Database Mirror

Database Mirror Failover?

How to make use of the Rolling Snapshots

What can Scale-Out give us? # 1

- Availability
 - Portions of data can go offline but doesn't effect the whole
- Disaster
 - Recovery time (reduce time to restore - reduced when less to recover)
 - Large Disk Partitions can take long time to fix
 - Limit the impact of total disaster (i.e. when your DR strategy does not work)
- Cost
 - Reuse commodity hardware for less important data
 - Higher we Scale-Up, more expensive it becomes. Scale-Out can be cheaper

What can Scale-Out give us? #2

- Performance
 - Load balance
 - Mix and match (LOW consumers and BIG consumers)
 - Separation of workload types (OLAP and OLTP)
 - Parallelise a System (separate system requests across multiple hardware)
 - Overcome contentious parts of the DB server such as TempDB
- Capacity
 - Backup time (reduce time to backup)
 - Limited resources


Before you Scale-Out...Scale-In #1

- Keep statistics accurate and up to date
 - Avoid big temp vars, they have no indexes or stats
 - Stats can get skewed, ensure they are maintained
- Query Tuning
 - Avoid table scans
 - Use indexes correctly, and remove duplicates
 - Is parallelism right for your query (OLAP vs OLTP)
 - Reduce size of the result set
 - Always use a WHERE clause
 - DON'T use SELECT * replace with precise column list
 - Sensible clustered key, avoid large covered index and prefer include option

Before you Scale-Out...Scale-In #2


- Reduce PageIO
 - Filtered indexes
 - Sparse columns
 - Use correct data types
 - Use table/ row and page compression
 - Remove your LOBs from tables
 - Use other technologies such as FILESTREAM
 - Vertically partition
- Reduce contention on shared resources
 - Denormalize
 - Filegroups

Scalable Shared Database



Search: sqlhelp

 @GlennAlanBerry None of our customers implemented it. #sqlhelp. Polyserve actually had a better impl of it. 1 r/w node, many read nodes
heigesr2, [+] Mon 14 Mar 03:02 via TweetDeck

 @GlennAlanBerry No customer of mine has implemented it #sqlhelp

SQLHA, [+] Mon 14 Mar 01:35 via web

 @GlennAlanBerry #sqlhelp But I don't believe it's used widely. Never heard anyone use it in clients since leaving MS 3.5 yrs ago
paulrandal, [+] Mon 14 Mar 01:31 via TweetDeck

 @GlennAlanBerry #sqlhelp Apparently a few customers have, including inside MS. Used to be owned by the Storage Engine (i.e. me).
paulrandal, [+] Mon 14 Mar 01:30 via TweetDeck

 @GlennAlanBerry yeah, we looked at it and quickly abandoned it, opting for log shipping instead. #sqlhelp

aaronbertrand, [+] Mon 14 Mar 00:08 via TweetDeck

 @GlennAlanBerry No, everybody I know who has those needs ended up using SAN snapshots & load balancers instead. #sqlhelp
BrentO, [+] Sun 13 Mar 23:49 via TweetDeck

 Has anyone ever actually used the Scalable Shared Databases feature in SQL Server? It seems pretty limited, with many restrictions. #sqlhelp
GlennAlanBerry, [+] Sun 13 Mar 23:48 via TweetDeck

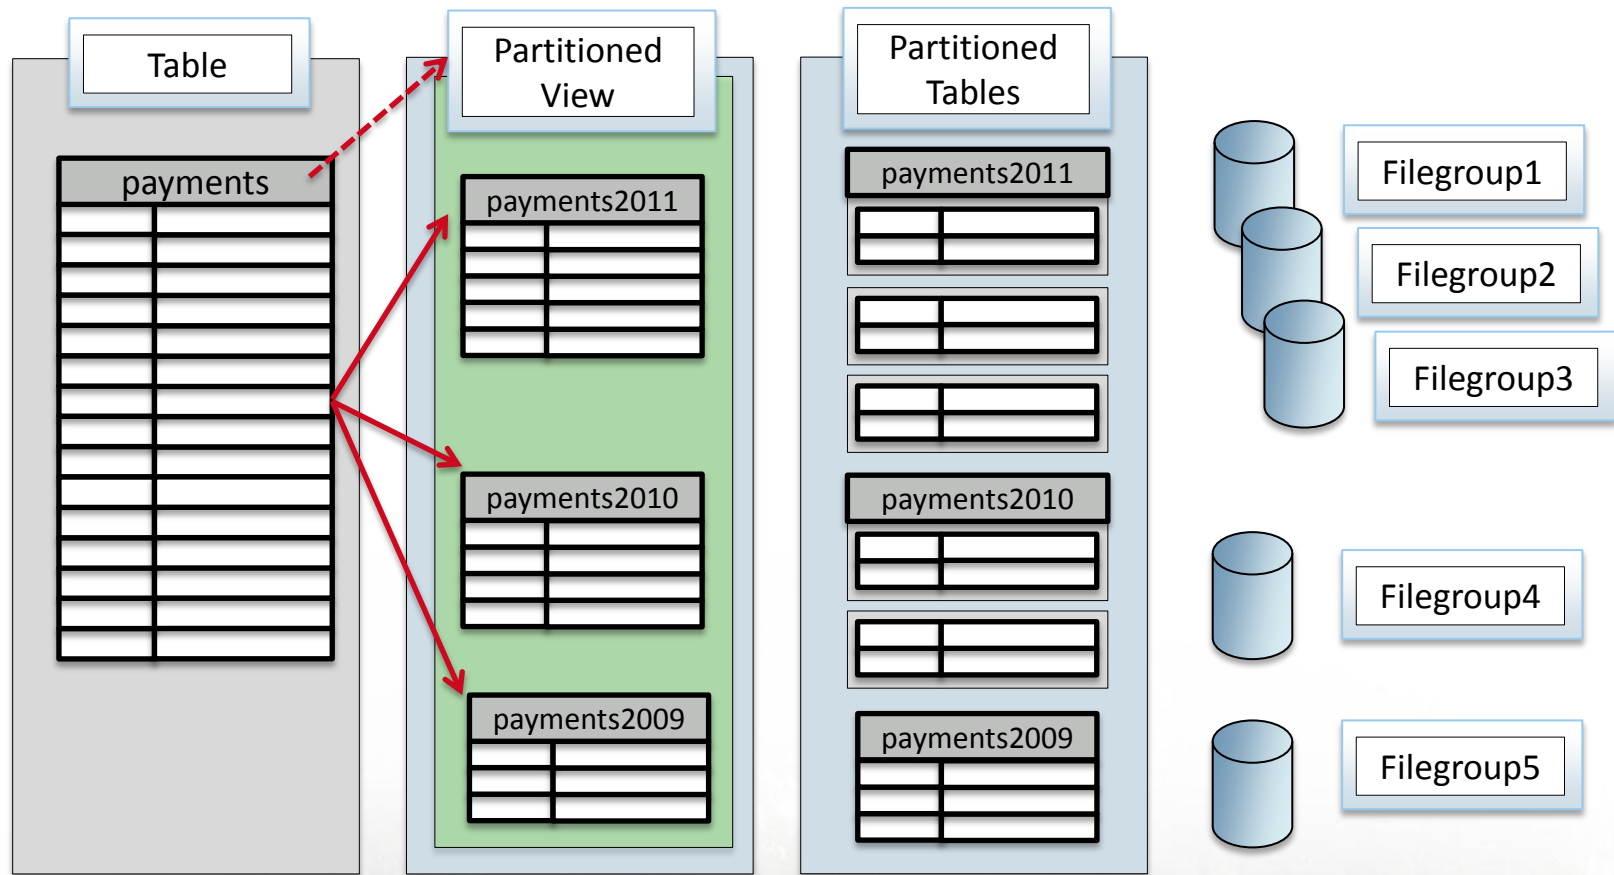
DEMO: Reporting Services Scale-Out

Reporting from Snapshots

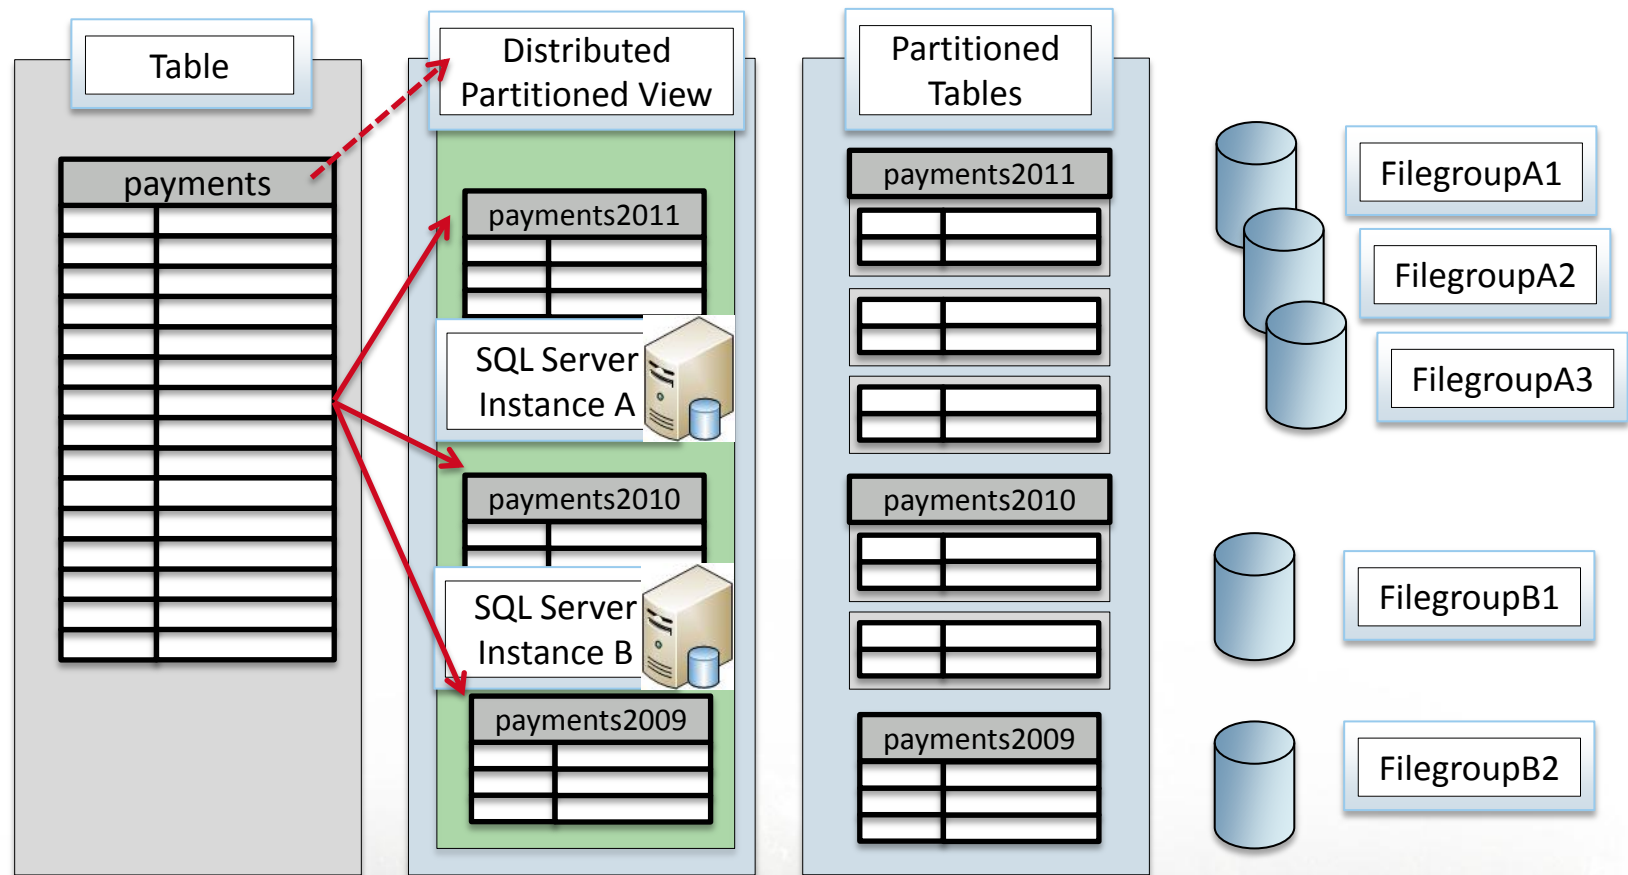
Reporting Services Scale-Out deployment

Scalable Shared Database

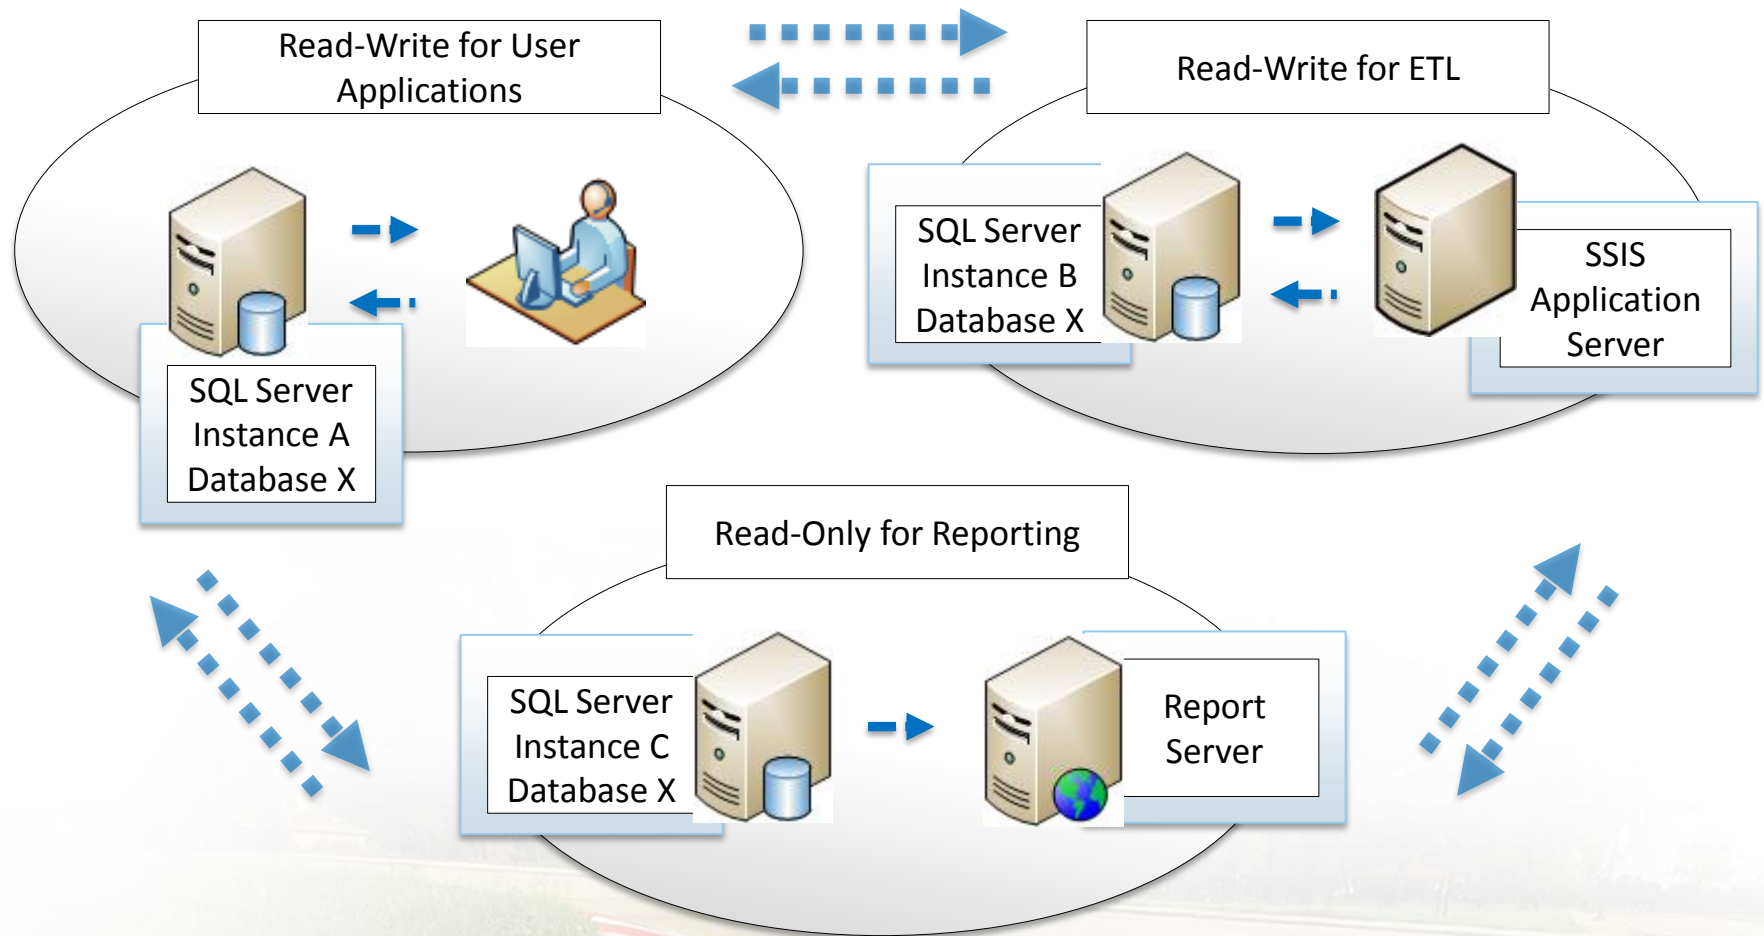
Partitioned tables, views and filegroups



With Distributed Partitioned views



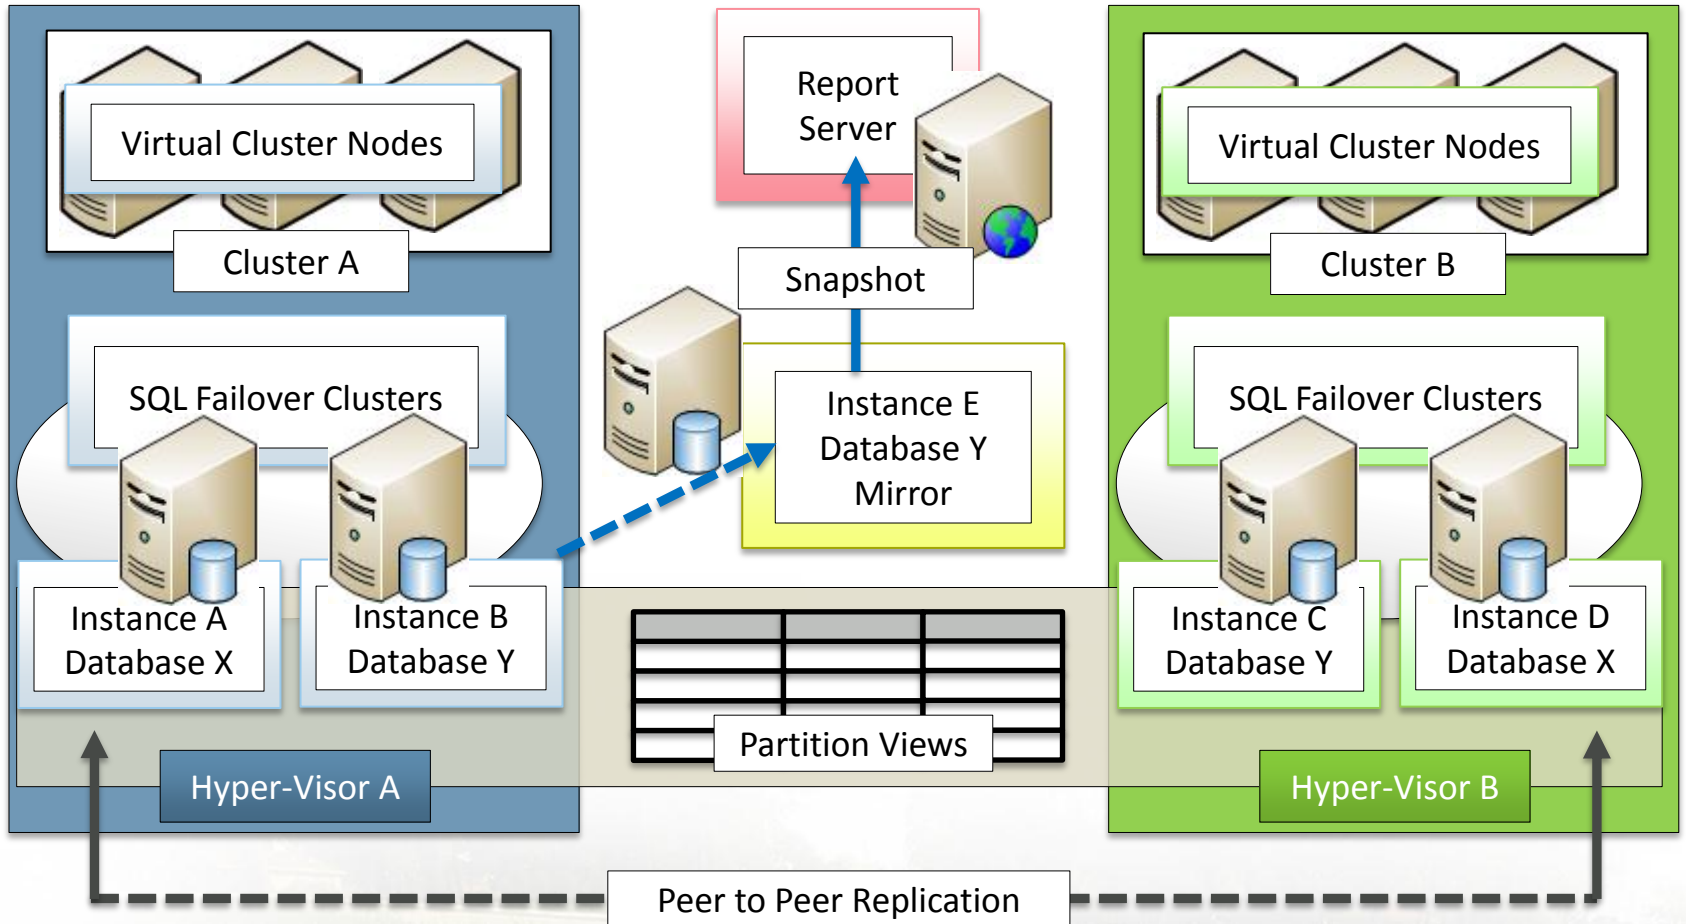
Peer to Peer Replication



Hybrid Scale-Out

- Database Mirroring with rolling snapshots
- SQL Failover Cluster using over-provisioned failover node “hot-swap Scale-Out/Up”.
- Use Hyper-Visor, migrate to over-provisioned host server.
- Clusters, Peer to Peer, Mirroring on Hyper-Visor

Bringing it all together




The SQL Server Scale-Out Toolkit

- Service Broker
- Integration Services
- Replication
- Horizontally Partitioned Views
- Federated Databases
- Partitioned Views
- Log Shipping
- Scalable Shared Database
- Scalable Shared Database for Analysis Services
- Reporting Services Scale-out Deployment
- Synonyms
- Schemas
- Query Notifications
- Full Text Indexing
- Vertical Partitioning
- Powershell
- CLR
- Linked Servers
- Filegroups
- Files
- Clustering
- Backup and Restore
- Mirroring
- Database Snapshots
- Processor Affinity
- Triggers
- Analysis Services Load Balancing

Taming the Beast

- Governance
 - Policy Based Management
 - Resource Governor or WRSM
 - Source Control
- Monitoring
 - MDW and Data Collection
 - Performance condition alerts
 - Extended Events
 - Profiler
 - DMVs
- Naming
 - SQL Client Aliases w/ GP
 - DNS



In Summary

- We discussed
 - Why we should start thinking about Scaling-Out?
 - Benefits from Scale-Out
 - Scaling in before you Scale-Out
 - Scale-Out strategies
 - Hybrid Scale-Out strategies
 - Keeping your scaled environment under control

Further References

- Books

- Apress - Pro SQL Server 2008 Service broker – Klaus Aschenbrenner
- Apress - Pro SQL Server 2008 Replication - Sujoy Paul
- Morgan Kaufman - DW 2.0 - The Architecture for the Next Generation of Data Warehousing – William Inmon, Derek Strauss and Genia Neushloss
- MS Press - Improving .NET Application Performance and Scalability

- Blogs/ Websites

- Partitioned Table & Index Strategies Using SQL Server 2008 <http://bit.ly/g28zQa>
- VoltDB.NET: Synchronous vs. Asynchronous Request Processing <http://bit.ly/k3rY2N>
- Data Warehousing 2.0 and SQL Server: Architecture and Vision <http://bit.ly/4tRXB4>
- Performance Considerations of Data Types – Michelle Ufford <http://bit.ly/aq9Wyr>

- Video/ Webcasts

- PASS Summit 2010: AD270S Database Design Fundamentals - Louis Davidson
- MCM #17 SQL Server Partitioning – SQLSkills <http://bit.ly/ea3C6e>

Please Complete the Evaluation Form

Pick up your evaluation form:

- In each presentation room

Drop off your completed form

- Near the exit of each presentation room
- At the registration area

Presented by Dell



GET IN TOUCH!

Contact Me:

Blog: tenbulls.co.uk

Email: mark.broadbent@sqlcloud.co.uk

Twitter: @retracement

LinkedIn: <http://linkd.in/gylT4I>



May 11-13 | Orlando, Florida

2011

THANK YOU!

For attending this session and
PASS SQLRally Orlando, Florida

Presented by Dell

