

Microsoft®  
**TechEd**  
Russia 2012



**Microsoft®**

Microsoft®  
**TechEd**  
Russia 2012



# Data Partitioning in VLDB

Dmitri Korotkevitch

<http://aboutsqlserver.com>

# About me

- ④ 10+ years of experience working with SQL Server
- ④ Microsoft SQL Server MVP
- ④ Microsoft Certified Master
- ④ Microsoft Certified Professional Developer
  
- ④ Blog: <http://aboutsqlserver.com>
- ④ Email: [dmitri@aboutsqlserver.com](mailto:dmitri@aboutsqlserver.com)



# Agenda

- ④ Why Partitioning?
- ④ Partitioning by Scaling-Out
- ④ Partitioning through Federations in Windows Azure
- ④ Data Partitioning

# Why Partitioning?

Everything stored in  
the same place

Same schema, indexes  
and compression

Maintenance  
challenges

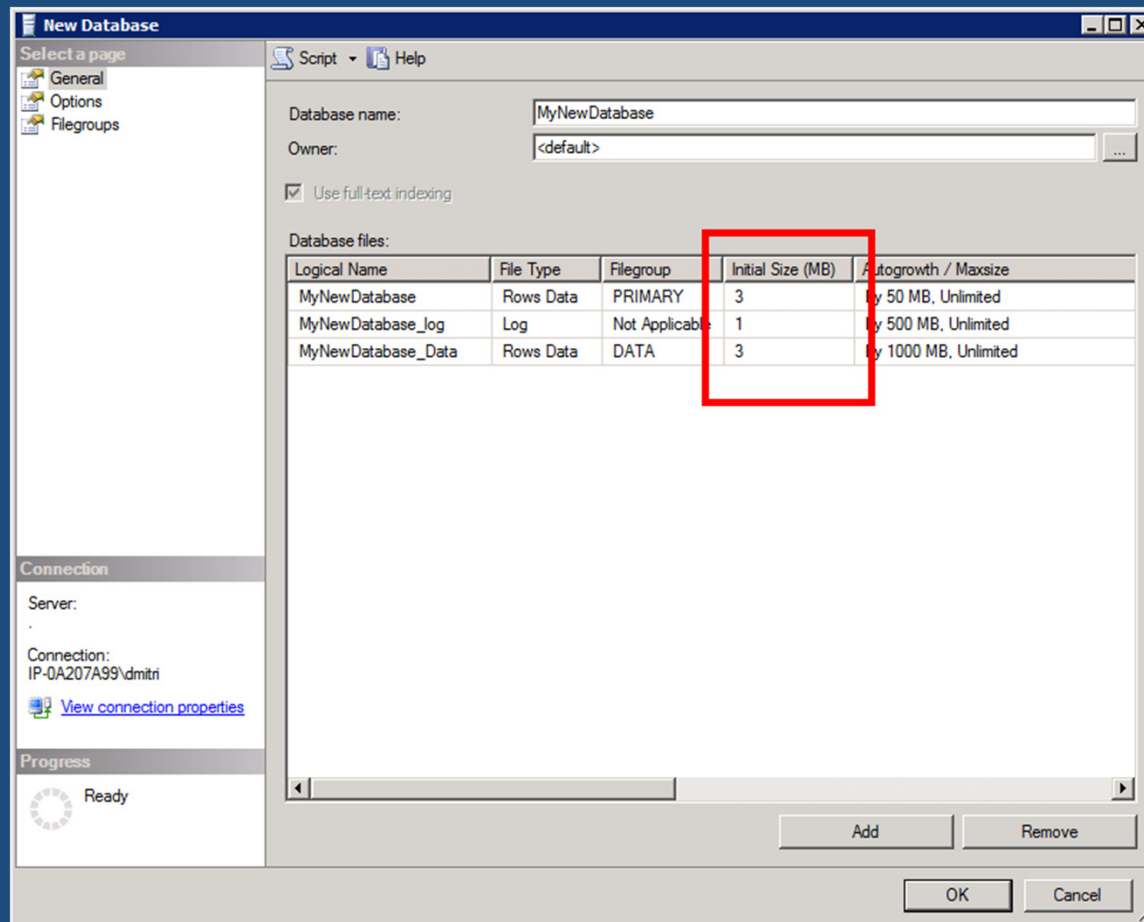


Challenges with  
Backups

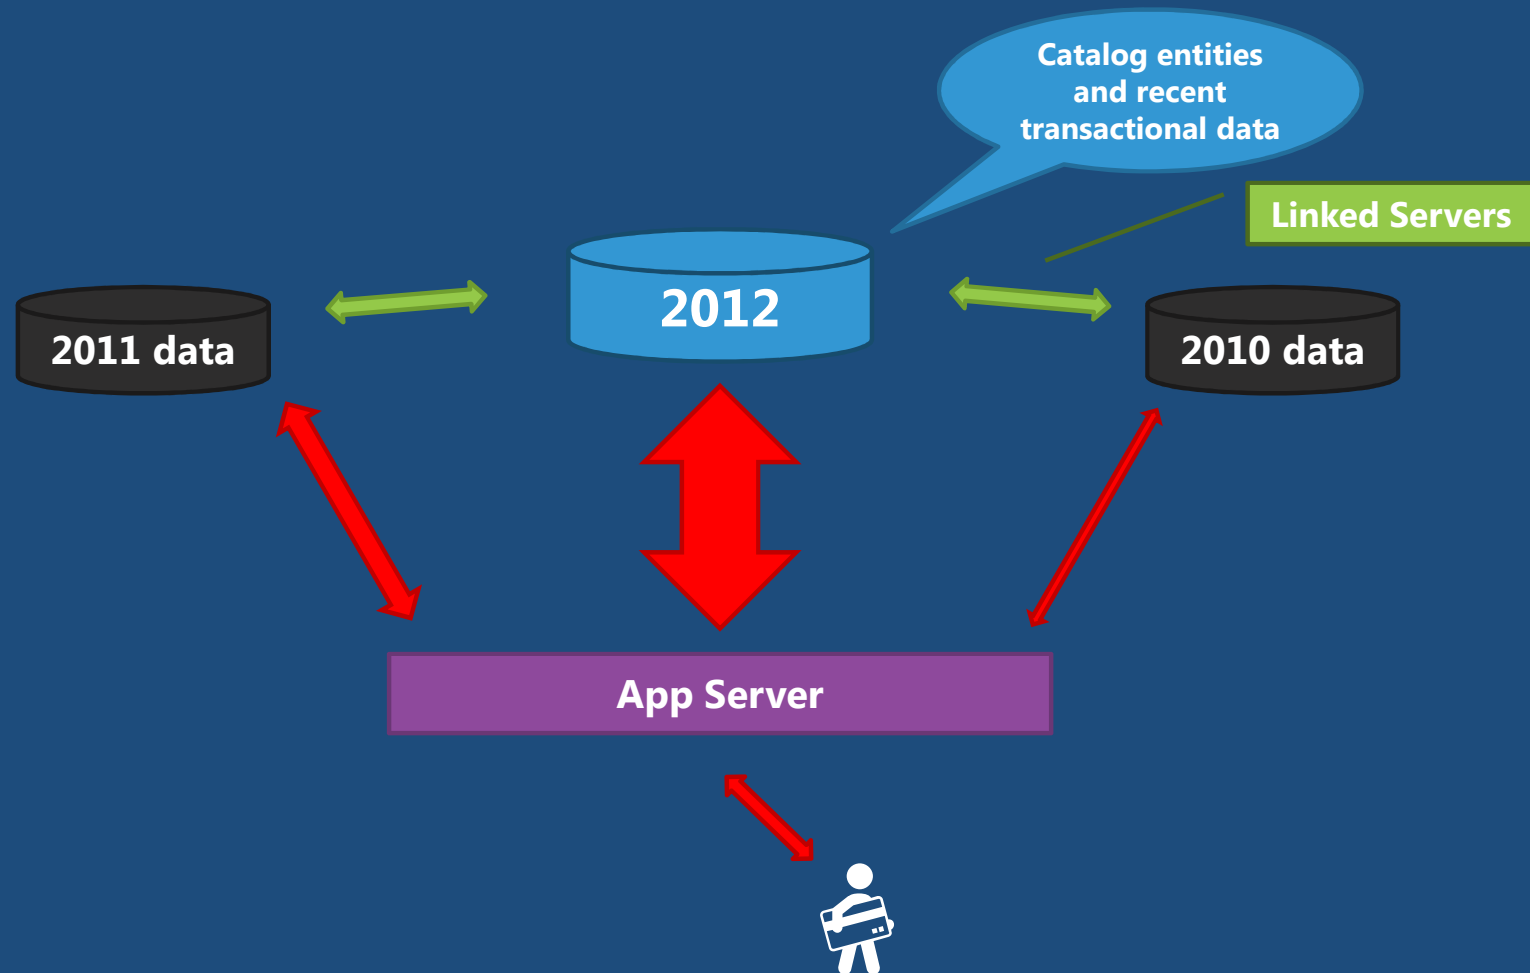
Long Disaster Recovery  
time

Less accurate statistics

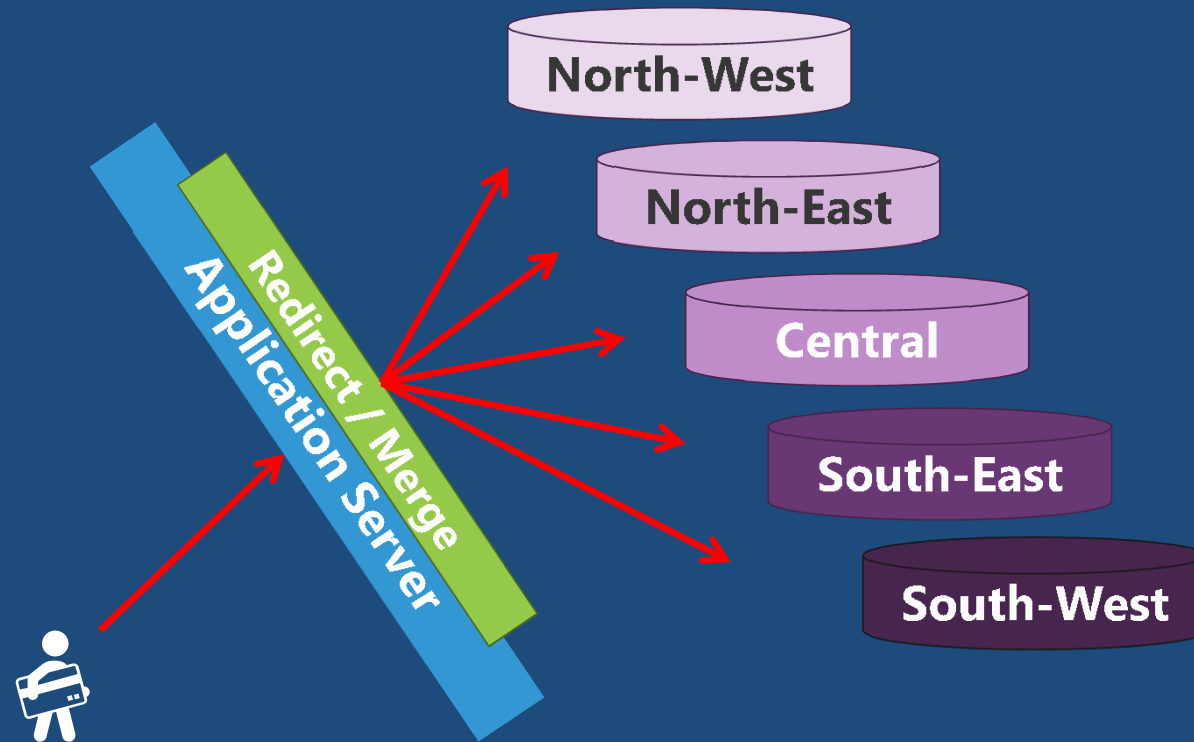
# When to partition data?



# Vertical Partitioning



# Data Sharding (Share-Nothing)





# Things to think about


## Cons

- ④ **Development cost – YMMV**
- ④ Maintenance cost
- ④ Hardware and License Cost (OS, SQL Server)
- ④ Potential performance problems with Linked servers
- ④ Legacy support

## Pros

- ④ Higher **availability** in some cases
  - ④ Share Nothing – when shard is down only subset of the customers is affected
  - ④ Vertical Partitioning – when “historical” data is unavailable, customers can still access current data
- ④ Can help with resource bottlenecks
  - ④ Especially with the Clouds

# Windows Azure SQL Database Federations








```
CREATE FEDERATION ClientFederation (CID int RANGE)
go

ALTER FEDERATION ClientFederation
SPLIT AT (CID=100)
go

ALTER FEDERATION ClientFederation
SPLIT AT (CID=200)
go
```

# Windows Azure SQL Database Federations






# Windows Azure SQL Database Federations

# Windows Azure SQL Database Federations

- ④ Share-Nothing approach
  - ④ No cross-database / cross-federations joins
  - ④ No fan-out queries
    - ④ Herve Roggero's Enzo Shard Library on CodePlex
- ④ Federation members can *technically* have different schemas
- ④ Help to scale-out beyond 150GB DB size limit
- ④ Carefully choose Federation Key
  - ④ Federate multiple related tables
  - ④ Split in the way that spreads data equally across the members

# Partitioned Tables




```
CREATE PARTITION FUNCTION pfOrders(DATETIME)
AS RANGE RIGHT
FOR VALUES('2010-01-01', '2010-02-01',
/*...*/ '2011-01-01', '2011-02-01',
/*...*/ '2012-11-01', '2012-12-12');

CREATE PARTITION SCHEME psOrders
AS PARTITION pfOrders
TO ([Orders2010], [Orders2010],
/*...*/ [Orders2011], [Orders2011],
/*...*/ [Orders2012], [Orders2012]);

CREATE TABLE dbo.Orders
(
    OrderID INT not null,
    OrderDate DATETIME not null,
/*...*/
    CONSTRAINT PK_Orders
    PRIMARY KEY CLUSTERED
    (OrderDate, OrderId)
    ON psOrders(OrderDate)
);
```

# Partitioned Views



```
CREATE TABLE dbo.Orders2010_01
(
    OrderID INT not null,
    OrderDate DATETIME not null,
    /*...*/
    CONSTRAINT PK_Orders2010_01
    PRIMARY KEY CLUSTERED (OrderDate, OrderId),
    CONSTRAINT CHK_Orders2010_01 CHECK
    (OrderDate >= '2010-01-01' and OrderDate < '2010-02-01')
) ON [Orders2010];
/*...*/
CREATE TABLE dbo.Orders2012_12
(
    OrderID INT not null,
    OrderDate DATETIME not null,
    /*...*/
    CONSTRAINT PK_Orders2010_01
    PRIMARY KEY CLUSTERED (OrderDate, OrderId),
    CONSTRAINT CHK_Orders2010_01 CHECK
    (OrderDate >= '2012-12-01' and OrderDate < '2013-10-01')
) ON [FGOnFastStorage]
go

CREATE VIEW dbo.Orders(OrderId, OrderDate /*...*/)
WITH SCHEMABINDING
AS
    SELECT OrderId, OrderDate /*...*/ from dbo.Orders2010_01
    UNION ALL
    SELECT OrderId, OrderDate /*...*/ from dbo.Orders2012_12;
```

# Comparing Approaches

## Partitioned Tables


- ⌚ Enterprise (and Developer) Edition only
- ⌚ 1,000 / 15,000 partitions max
- ⌚ No easy way to move partition to another FG online
- ⌚ Statistics on table level
- ⌚ One schema and indexes
- ⌚ No online partition level index rebuild
- ⌚ Easy maintenance but schema changes are time consuming
- ⌚ Replication friendly

## Partitioned Views

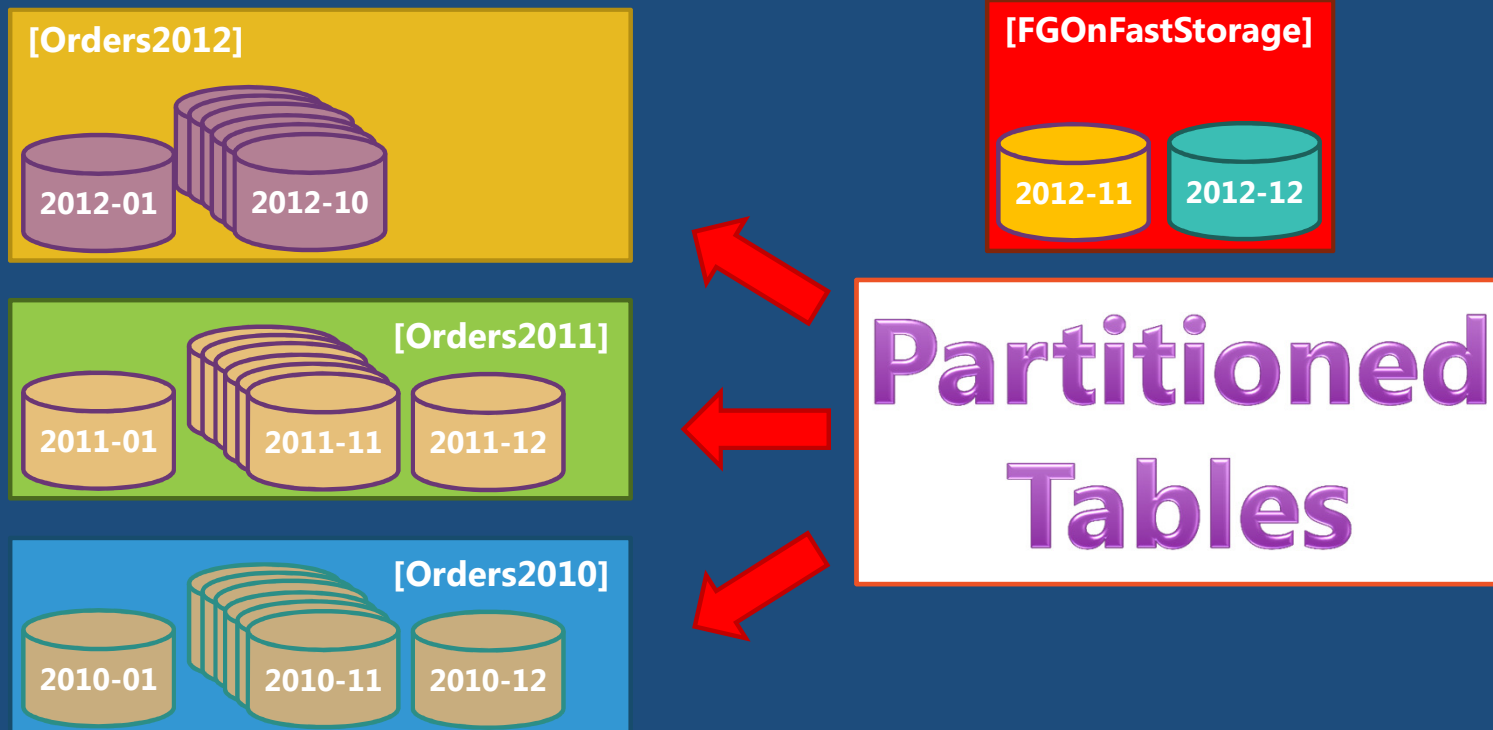
- ⌚ All SQL Server editions
- ⌚ 255 tables max
- ⌚ "Partition" can be moved to another FG online (Enterprise Edition)
- ⌚ Statistics on "partition" level
- ⌚ Schema and indexes can vary per "partition"
- ⌚ Online index rebuild on "partition" level (Enterprise Edition)
- ⌚ Maintenance overhead
  - ⌚ Especially with Replication



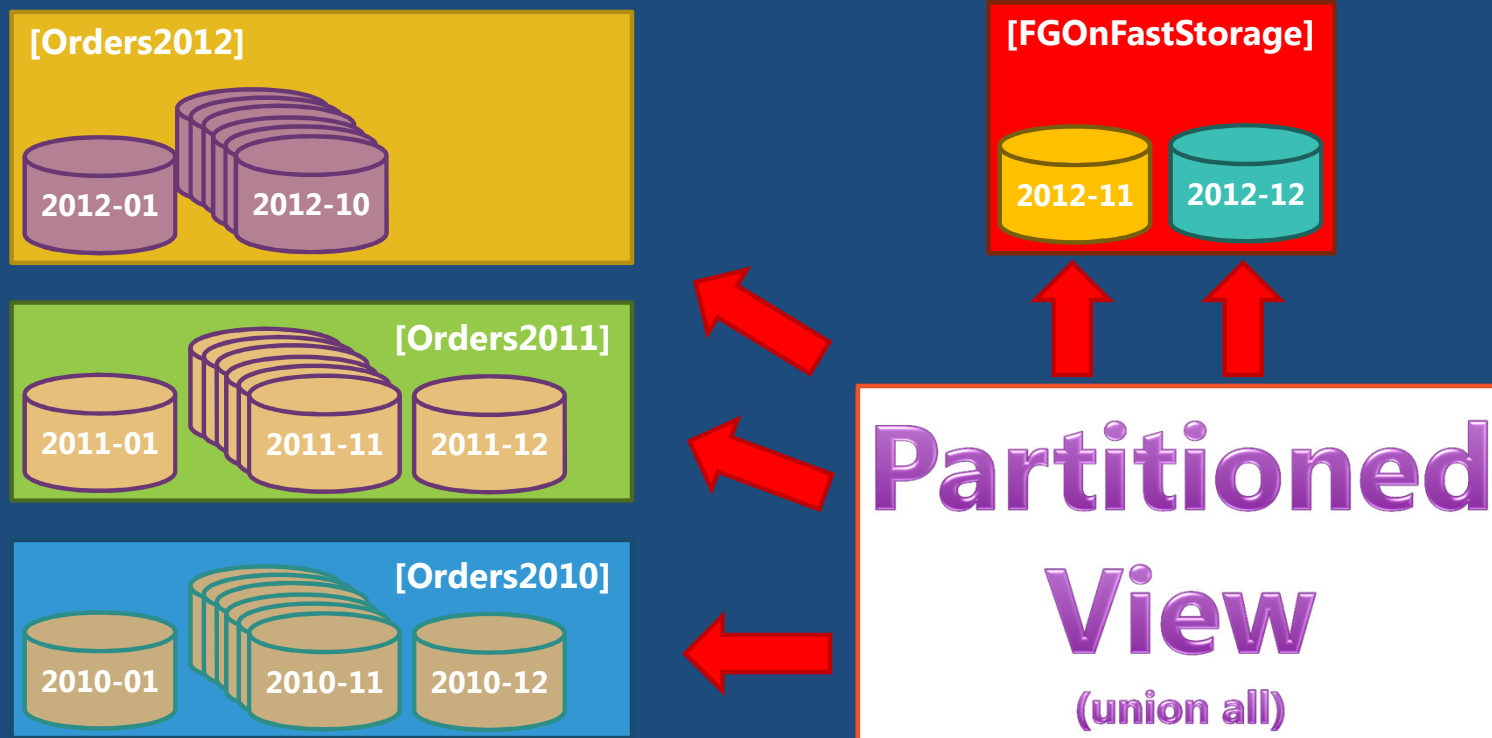
# Mixed Approach



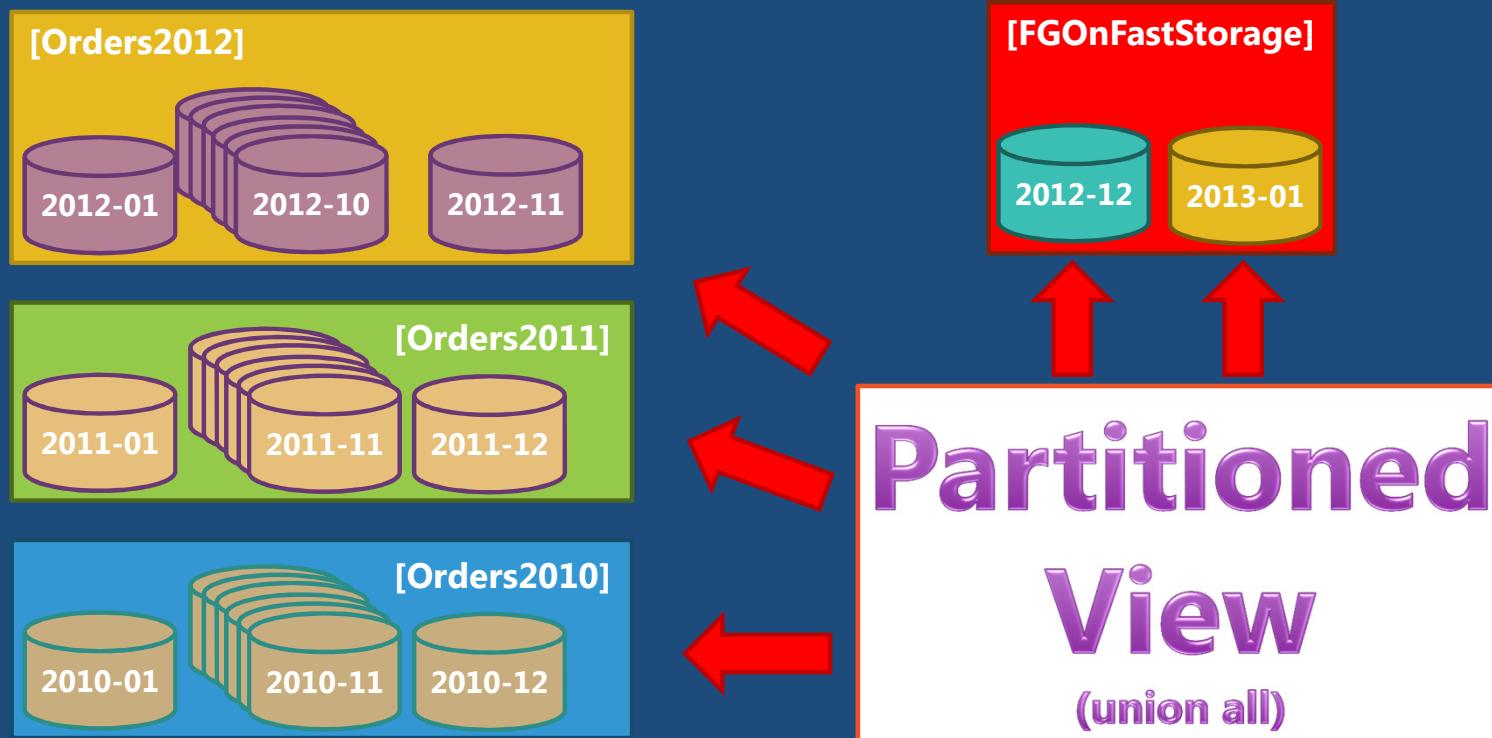
# Mixed Approach



# Mixed Approach




# Mixed Approach



# Partitioning and Data Layout

- ④ Consider backup and DR strategies and storage system constraints and limitations
- ④ Place related objects together
  - ④ Example: *Orders* and *OrderItems*
- ④ Separate read-only and read-write data
  - ④ Make Filegroups read-only when needed
- ④ HA strategies and dynamic files/FG creation
  - ④ Failover cluster – single copy of DB
  - ④ AlwaysOn, Mirroring, Log Shipping – works as long as folder/disk layout is the same

Microsoft®  
**TechEd**  
Russia 2012



## AlwaysOn and Partitioning

# Potential issues with partitioned tables

- ④ Storage size increase
  - ④ Partition column needs to be added to CI and NCI
    - ④ Example: Datetime (8 bytes) \* 10M records per day = ~28GB per NCI per year
- ④ Suboptimal Execution Plans in some cases
- ④ More Info:
  - ④ <http://www.simple-talk.com/content/article.aspx?article=1587>

# We discussed

- ④ Why to partition data
- ④ Different architectural approaches with data partitioning



# Q&A

④ Thank you very much for attending

④ Dmitri Korotkevitch

④ Email: [dmitri@aboutsqlserver.com](mailto:dmitri@aboutsqlserver.com)

④ Blog: <http://aboutsqlserver.com>



© 2012 Microsoft Corporation. All rights reserved. Microsoft, Windows, Windows Vista and other product names are or may be registered trademarks and/or trademarks in the U.S. and/or other countries.  
The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.