

New technologies for BI and Data Warehousing – they're cool alright, but how do they fit together?

Amit Bansal

www.PeoplewareIndia.com

www.SQLServerGeeks.com

www.amitbansal.net

Microsoft®
tech·ed
India | 2011

MARCH 23-25 | BANGALORE



About Amit Bansal

- CTO, eDominator Systems & Peopleware India
- Corporate Trainer/Consultant & Evangelist
- Conducted more than 400 workshops on SQL Server & BI for top notch IT companies world wide
- Microsoft MVP for SQL Server
- Microsoft Certified Trainer Advisory Council member
- Speaker at TechED India, TechED US & TechED Europe
- Technical Reviewer – MSL courses on SQL Server
- SME – SQL Server 2008 certifications
- President – SQLServerGeeks.com

Past

SQL Server
2005
SQL Server
2008
SSAS (UDM),
SSIS, SSRS
PPS
EXCEL
SHAREPOINT

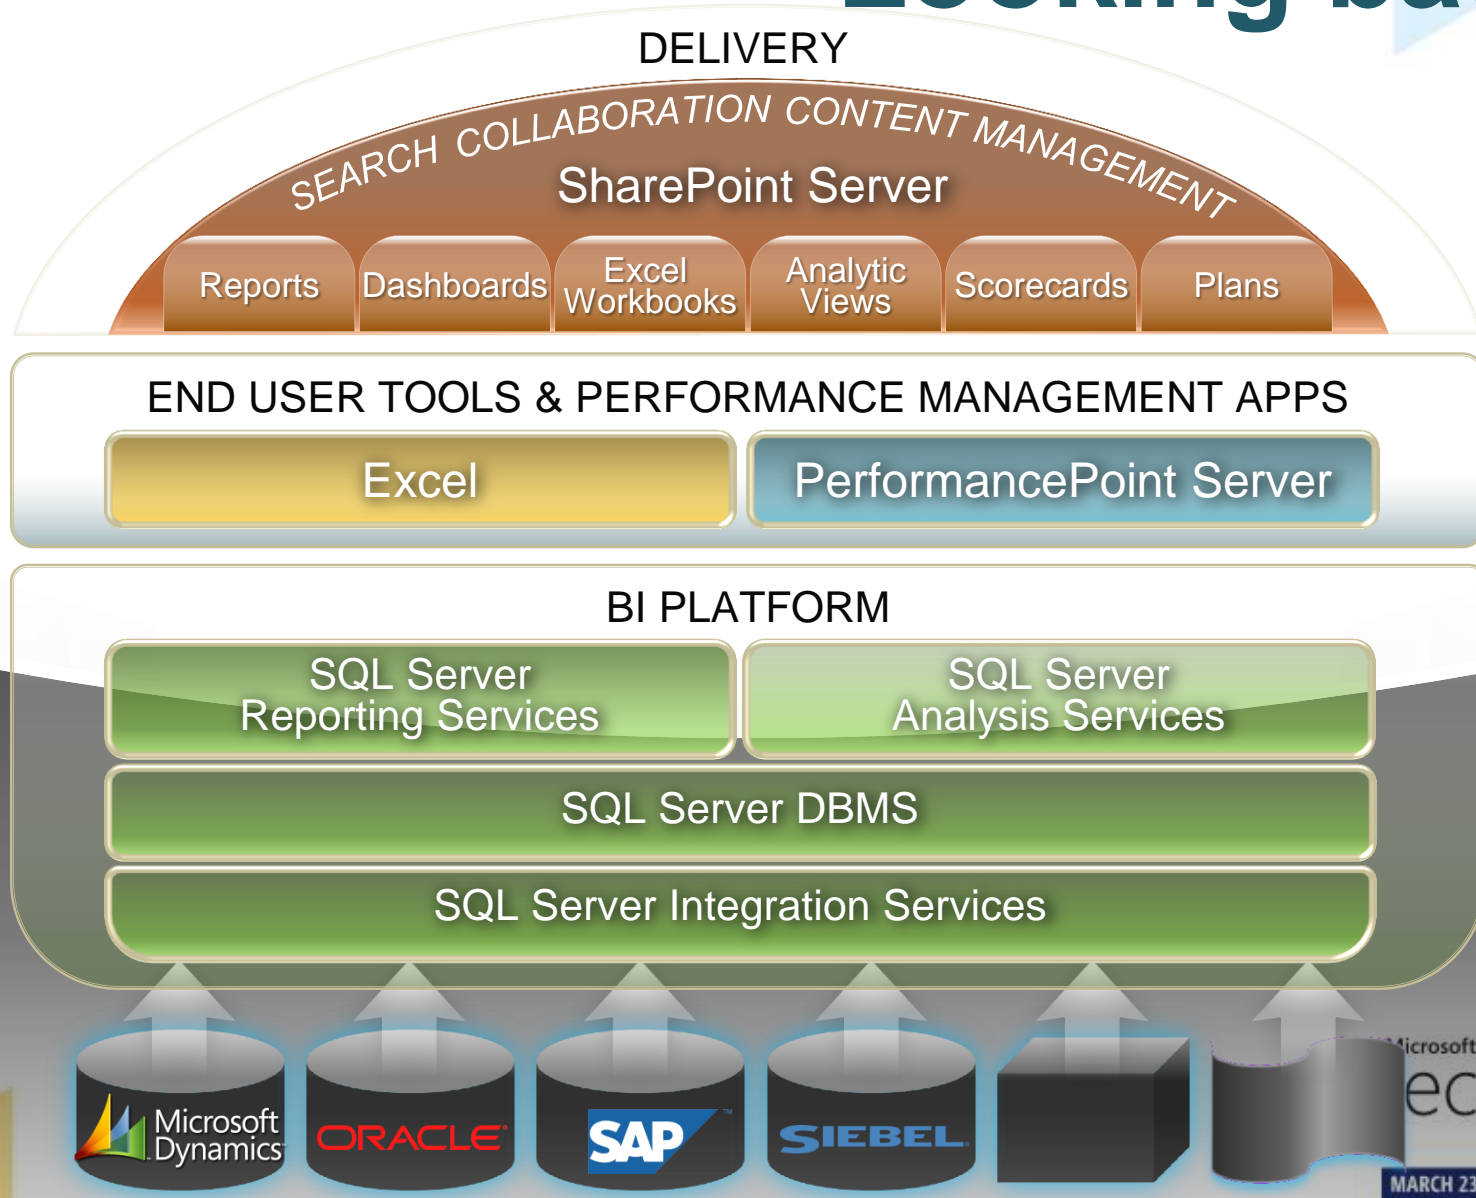
Present

SQL Server
2008 R2
PowerPivot
FastTrack &
Parallel DW
Report Builder
3.0
VertiPaq Engine

Future

‘Denali’
VertiPaq
Column Store
Crescent
BISM

Looking back



Looking back – BI Timeline

Microsoft
SQL Server 2005

Microsoft Office
Business Scorecard
Manager 2005



Microsoft
SQL Server 2005
• SP1

Microsoft
SQL Server 2005
• SP2

Microsoft
Office
• MOSS
• Excel 2007

Microsoft
Office PerformancePoint
Server 2007

Microsoft
SQL Server
Code Name "Katmai"
Microsoft
SQL Server 2008

Microsoft
SQL Server 2008 R2

Microsoft
Office 2010

Microsoft
SharePoint 2010

2005

2006

2007

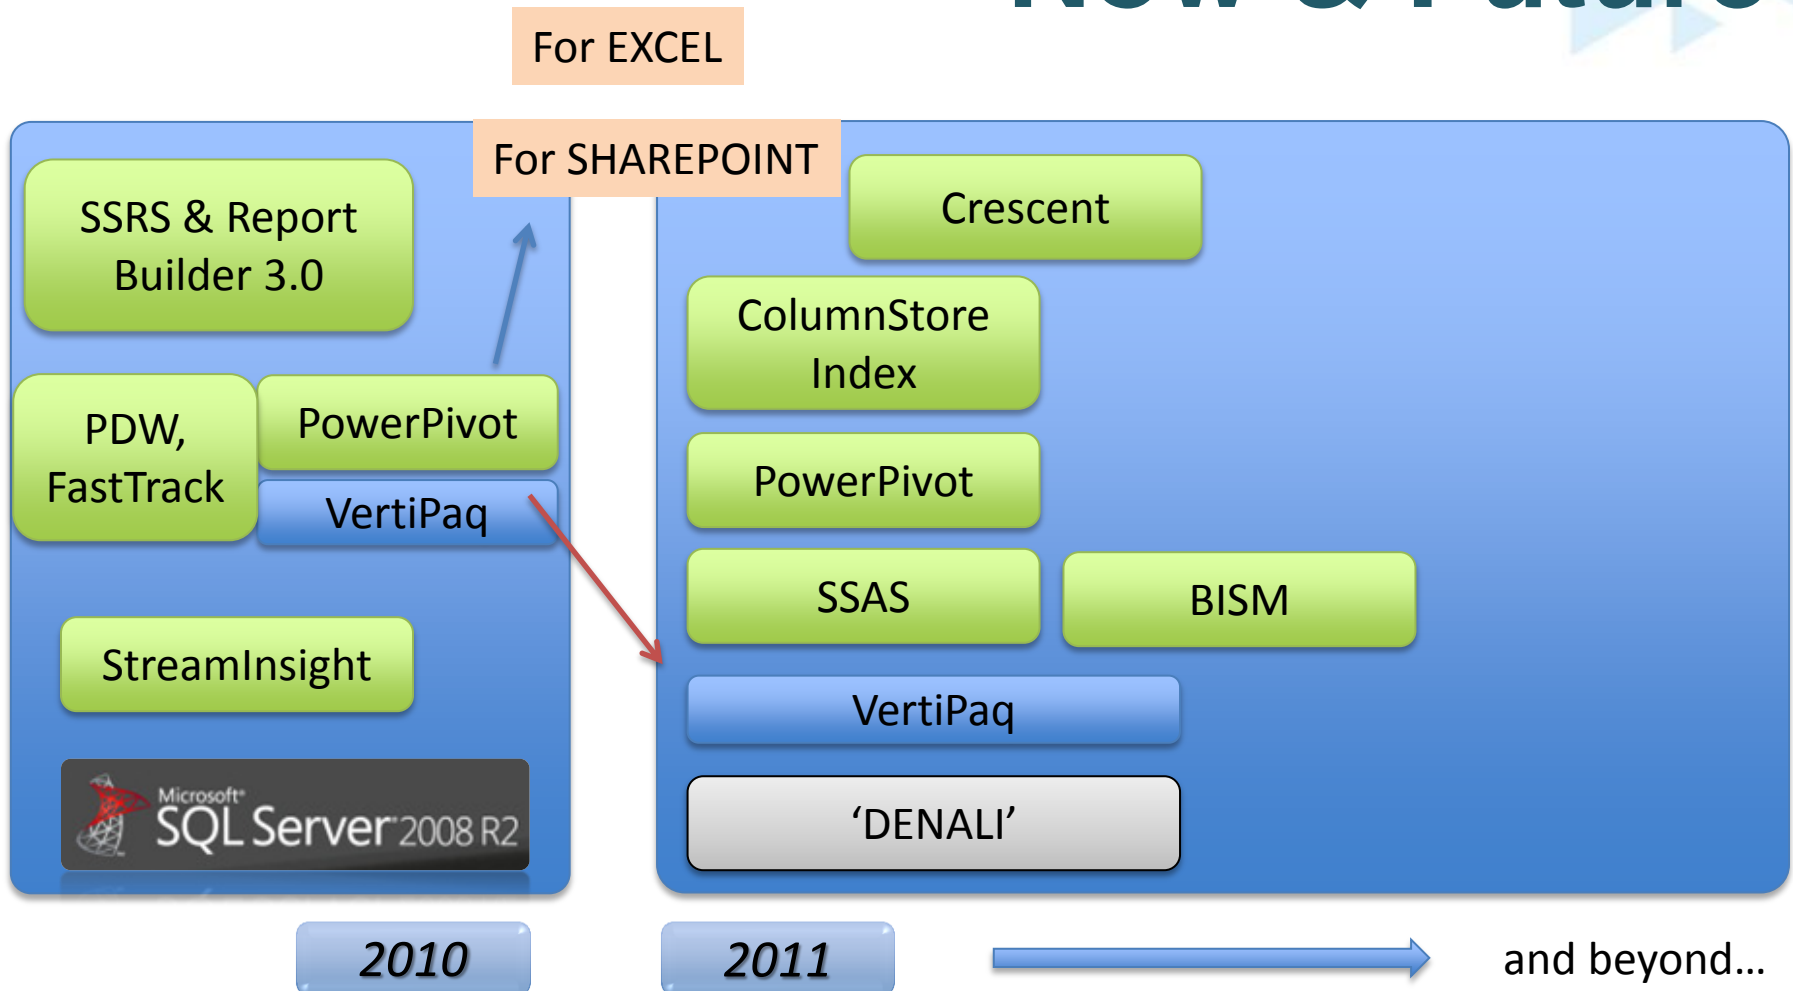
2008

2010

Microsoft
tech.ed
India | 2011

MARCH 23-25 | BANGALORE

Now & Future



Fast Track & PDW - FAQ

- A new edition: SQL Server 2008 R2 Parallel DW edition
- SQL Server equivalent of Windows Server HPC
- Utilizes Massively Parallel Processing (MPP) scale-out architecture to support very large data warehouses
- How is different from other editions?
- Is Fast Track same as PDW?
 - No, Fast Track uses SMP; can run on other editions
- Targeted at 10 TBs to 1+ PB
- Architecture?


StreamInsight - FAQ

- Real-time Business Intelligence systems are event driven, and use Event Stream Processing techniques to enable events to be analysed without being first transformed and stored in a database.
- These in- memory techniques have the advantage that high rates of events can be monitored, and since data does not have to be written into databases data latency can be reduced to milliseconds.

VertiPaq engine


- VertiPaq engine is currently only available to PowerPivot
- Columnar database is its driving principle
 - Column bases storage as opposed to “row store”
 - High level of compression, all in memory – RAM based
- Data retrieval and calculations happen at a much faster rate
- Disk I/O processing is omitted
- Yes, the entire database is in-memory

ColumnStore Index



ColumnStore Index

- Benefits
 - Only columns need to solve the query are fetched from the disk
 - Much better compression
 - Highly improved buffer rates
- Employs VertPaq technology
- Don't have fit in memory
- Pure “column stores”, not a hybrid



ColumnStore Index

- Performance Illustration
 - 1 TB DW
 - Sales fact table with 1.44 billion rows

	Total CPU time (seconds)	Elapsed time (seconds)
Columnstore	31.0	1.10
No columnstore	502	501
Speedup	16X	455X

Project 'Crescent'

- Project “Crescent” expands on the Self-Service BI capabilities delivered with PowerPivot for Excel
- Visualize and interact with the modeled data in a fun and meaningful way using highly interactive visualizations, animations and smart querying
- Present and share insights with others in the organization through rich storyboard presentation capabilities.
- Not available in CTP 1

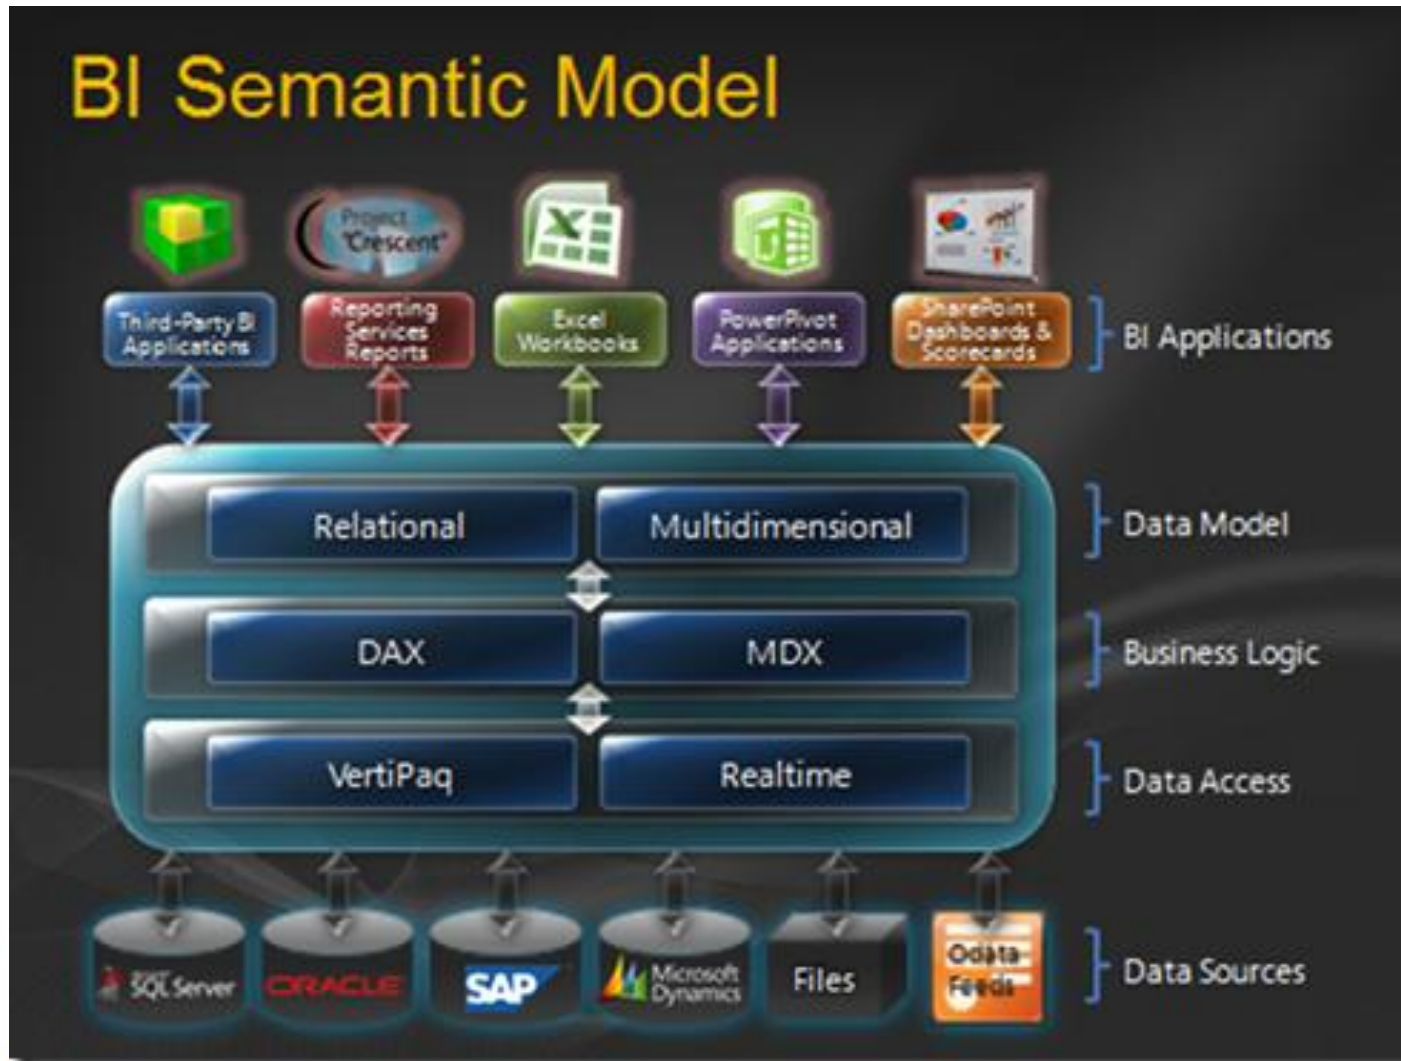
Project 'Crescent'

- Design Experience
 - Interactive, web-based authoring and sharing of information
 - Familiar Microsoft Office design patterns
 - Powerful data layout with banding, callout and small multiples visualizations
- Meta-Data Integration/Interactivity
 - Full interconnectivity with PowerPivot
 - Drive greater insight through smart and powerful querying
 - Zero configuration highlighting and filtering
 - Animated trending and comparisons
- Presentation Ready
 - Presentation and story board turn pervasive information into persuasive information
 - Deliver and collaborate through SharePoint
 - Interactive PowerPoint runtime

BI Semantic Model

- Expresses business logic in business terms
- Model richness scales to complexities of business
- Hides data access and integration
- Accessible by variety of tools
- Very high performance
- Data consumed in familiar business terms
- Single model can be used for reporting, analysis, dashboard, ...

BI Semantic Model



Bringing it all together

BI for Everyone – in “DENALI”

BI for IW (Self Service)

EXCEL, PowerPivot,
BISM, Crescent,
SharePoint

BI for IT Pro

SSAS, UDM, MDX,
SSRS, SSIS

Resources

Software Application Developers



<http://msdn.microsoft.com/>



[msdnindia](https://www.facebook.com/msdnindia)



[@msdnindia](https://twitter.com/msdnindia)

Infrastructure Professionals

Microsoft *TechNet*

<http://technet.microsoft.com/>



[technetindia](https://www.facebook.com/technetindia)



[@technetindia](https://twitter.com/technetindia)

Microsoft®
tech.ed
India | 2011

MARCH 23-25 | BANGALORE



Microsoft[®]

Be what's next.[™]

© 2011 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.



Microsoft[®]
tech.ed
India | 2011

MARCH 23-25 | BANGALORE