

SQL Azure and SqlBulkCopy

Presented by Herve Roggero
Managing Partner, Blue Syntax Consulting
SQL Azure MVP

Email: hroggero@bluesyntax.net
Twitter: @hroggero
LinkedIn: SQL Azure and SQL Server Security
LinkedIn: Azure Florida Association
www.bluesyntax.net



Agenda

- Available tools to load data
- INSERT commands
- SqlDataAdapter
- Analyzing BCP
- Analyzing SqlBulkCopy
- Performance test

Available tools/techniques

- INSERT
- INSERT (2)
- INSERT BULK
- SqlDataAdapter
- Migration Wizard
- SSIS
- BCP / INSERT BULK
- SqlBulkCopy

What makes up load time?

- Connection Time
- Send Query
- Compile Query
- Execute Query
- Send Data

INSERT

- Classic way to insert records
- Logged operation
- Usually single record at a time
- All packets are encrypted => too verbose
- VERY SLOW...


```
INSERT INTO dbo.MyProducts (Name, ListPrice) VALUES ('Helmet', 25.50)  
INSERT INTO dbo.MyProducts (Name, ListPrice) VALUES ('Wheel', 30.00);
```

INSERT (2)

- New INSERT command in SQL Server 2008
 - Table Value Constructor
- Not as verbose
- Operation is still logged
- Faster than regular INSERT, but still SLOW...

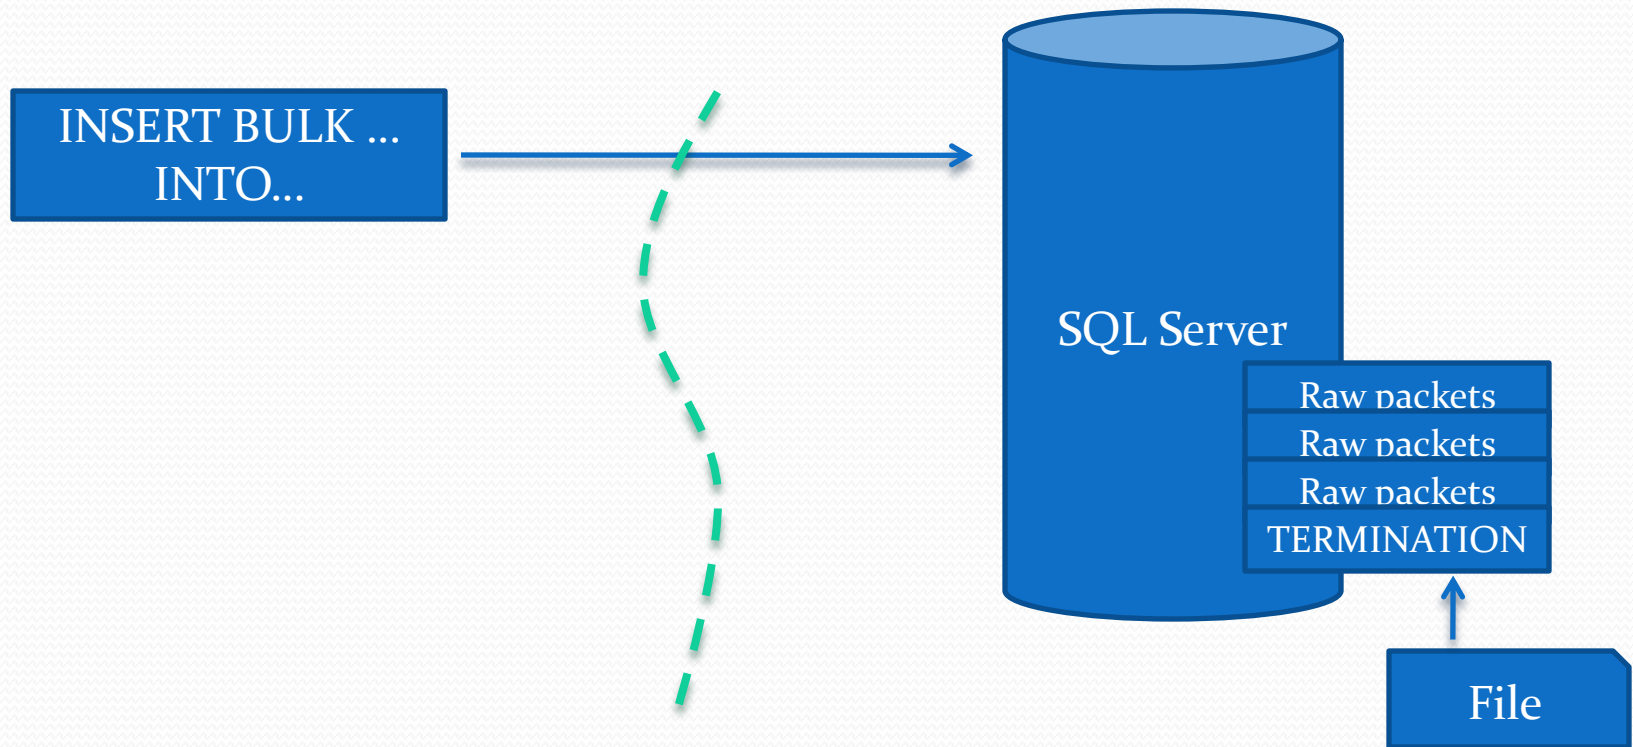
```
INSERT INTO dbo.MyProducts (Name, ListPrice) VALUES  
    ('Helmet', 25.50),  
    ('Wheel', 30.00);
```

INSERT (summary)



BULK INSERT

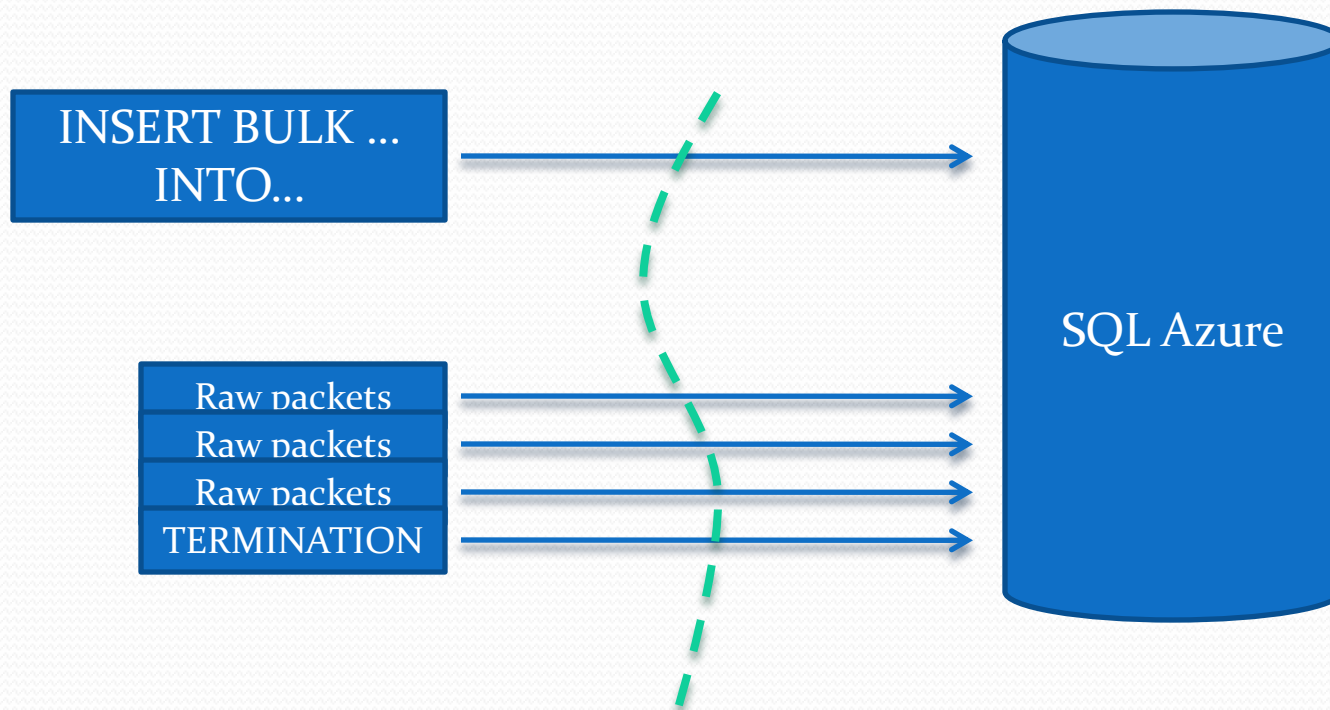
- **NOT SUPPORTED** in SQL AZURE...



INSERT BULK

- Supported... but requires feeding “raw” TDS packets
=> Must be used by a Tool
- Very efficient and lean
 - Sends 1 INSERT statement
 - Everything that follows are raw data packets
 - NO FURTHER COMPILATION/QUERY EXECUTION
 - SQL Azure waits until Session sends “TERMINATION”
- Operation not logged
- FAST...

INSERT BULK (cont.)



BCP

- Uses INSERT BULK
- Performance of BCP depends on INSERT BULK
- However...
 - BCP can also extract data
- FAST

Tools

- Tools that use the BCP program:
 - SSIS: Bulk Insert Task
 - SQL Azure Migration Wizard
 - <http://sqlazuremw.codeplex.com/>
- Tools that use the SQL Bulk Copy API
 - Enzo Data Copy

SqlDataAdapter

- Based on XML Datagrams (differences)
- Very memory intensive
- Can be very slow depending on approach
- Provides a way to work in Batches
- Sends INSERT operations
- SLOW...

SqlBulkCopy

- Uses INSERT BULK operation
- Programmatic access, options similar to BCP
- Can send data in Batch
- FAST

Techniques Summary

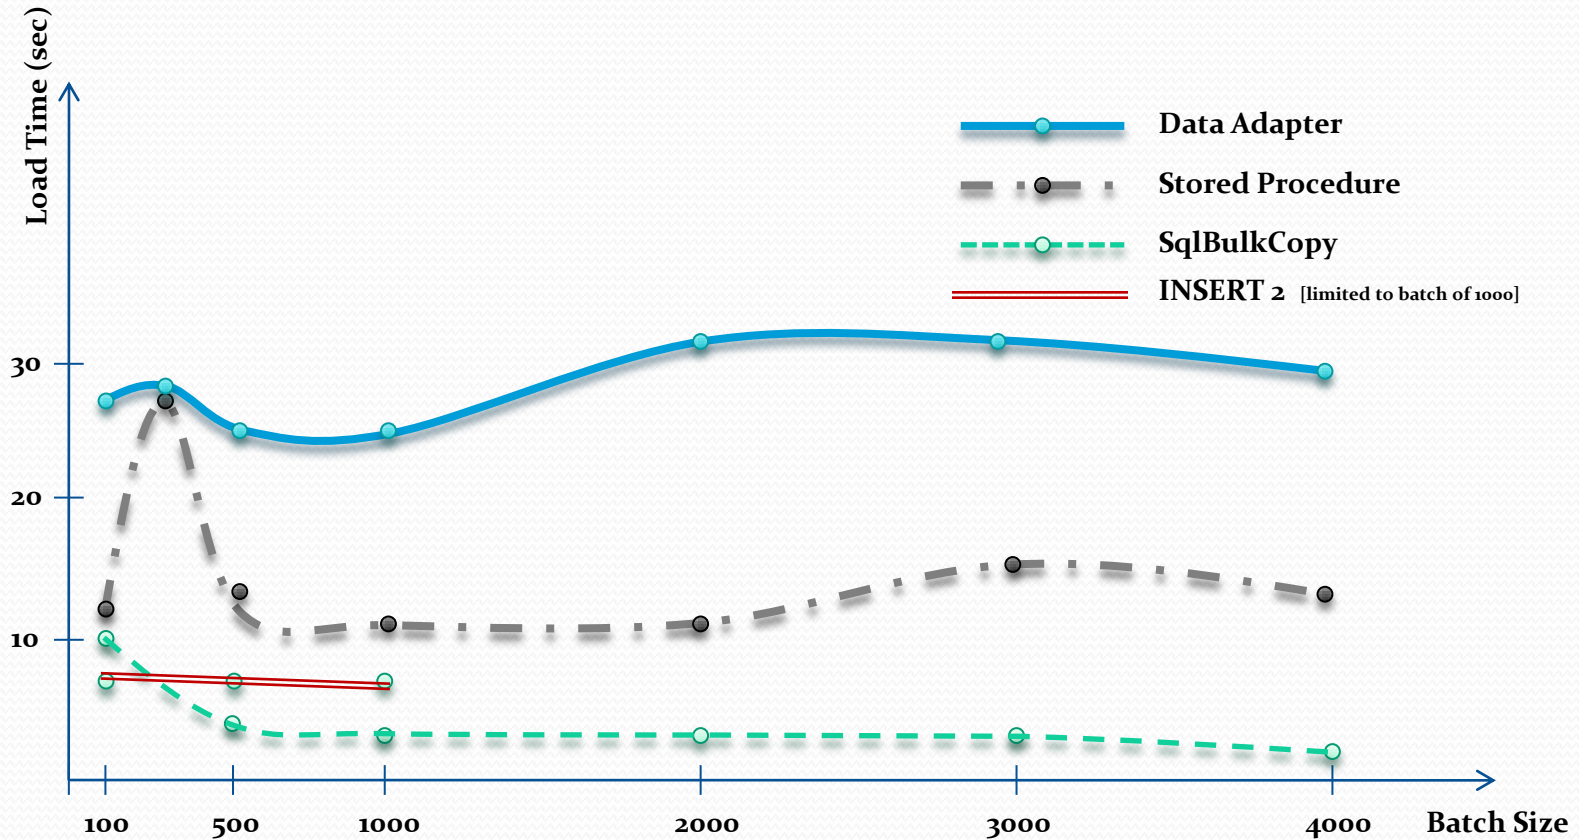
Tool/Technique	Operation	Logged	Speed
INSERT	INSERT	Yes	SLOW
BULK INSERT	BULK INSERT	No	n/a
INSERT BULK	INSERT BULK	No	FAST
BCP	INSERT BULK	No	FAST
SSIS/ Migration Wizard	INSERT BULK	No	FAST
Data Adapter	INSERT	Yes	SLOW
SqlBulkCopy	INSERT BULK	No	FAST
Enzo Data Copy	INSERT BULK	No	FAST

About the performance test

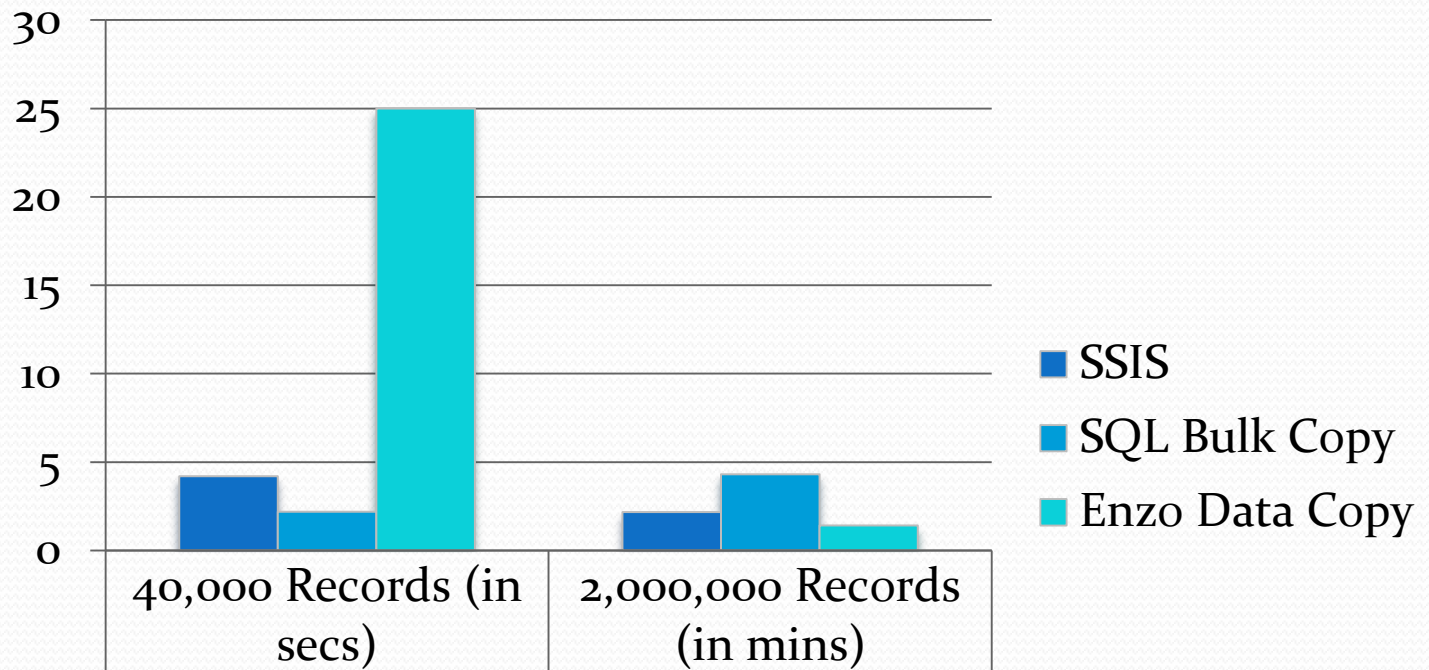
- Load 40,000 records (on-premise -> SQL Azure)
- Uses multithreading (TPL) – 8 threads
- Comparing
 - Data Adapter
 - Stored Procedure
 - SqlBulkCopy API
 - INSERTs

Performance Chart

- X: Batch Size; Y: Load Time {40K records; DOP = 8}



Comparing Tools



	40,000 Records (in secs)	2,000,000 Records (in mins)
SSIS	4.2	2.17
SQL Bulk Copy	2.18	4.3
Enzo Data Copy	25	1.41

Demos

- Sample Code
 - Uses parallel processing on all methods
 - Compares:
 - Data Adapter
 - Inline INSERTs
 - Stored Procedure INSERTs
 - SqlBulkCopy
- SSIS
- Enzo Data Copy (www.bluesyntax.net)

Considerations

- Resilience to Throttling
 - Adapting to throttling conditions
 - Changing workload depending on errors
- Connection Retries
 - Retrying for connection loss
 - Causes include throttling, load balancing, network...

Thank you

Presented by Herve Roggero
Managing Partner, Blue Syntax Consulting
SQL Azure MVP

Email: hroggero@bluesyntax.net
Twitter: [@hroggero](https://twitter.com/hroggero)
LinkedIn: SQL Azure and SQL Server Security
LinkedIn: Azure Florida Association
Blog: www.geekswithblogs.net/hroggero
www.bluesyntax.net

