



*The World's Largest Community
of SQL Server Professionals*

How to Use the SQL Server 2008 Performance Data Collector to Analyze Query Performance

Brad M. McGehee
Microsoft SQL Server MVP
Director of DBA Education
Red Gate Software
www.bradmcgehee.com

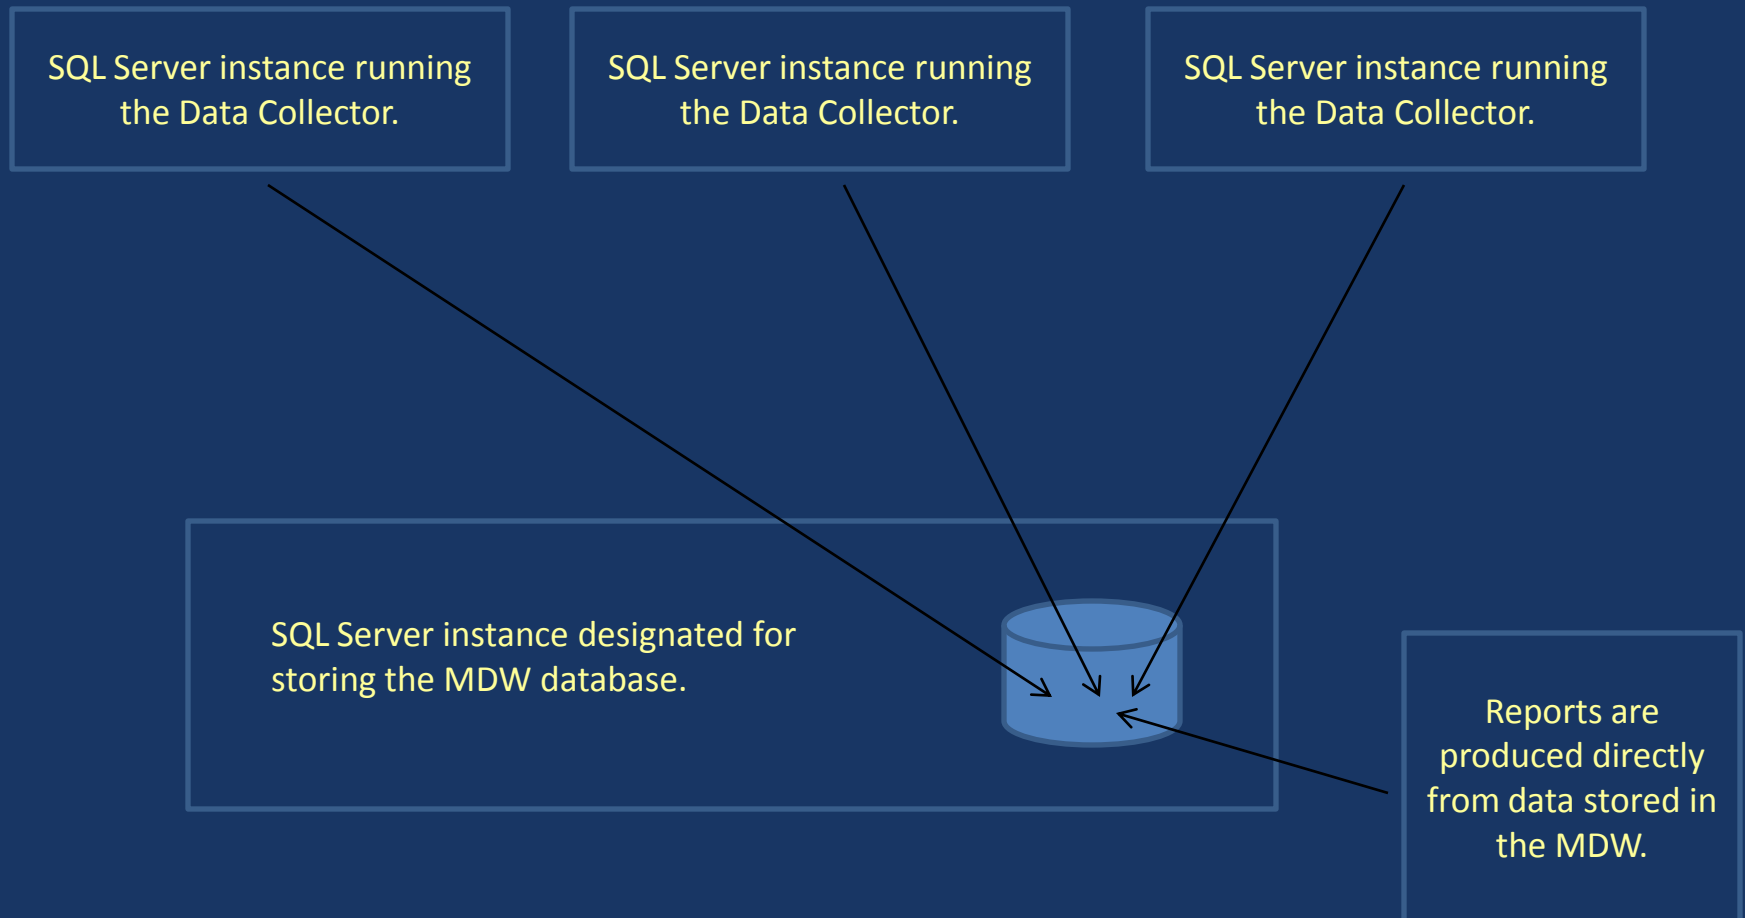
Here is What We are Going to Learn Today

- Introduction to the Data Collector
- Introduction to the Data Collector Reports
- How to Use the Data Collector to Analyze Query Performance

What is the Data Collector

- Data Collector is an optional feature of SQL Server 2008 that **collects and stores** information about SQL Server's 2008's status and performance over time, using **pre-defined data collection sets**.
- Once data is collected and stored in the **MDW database**, built-in reports are available to see what SQL Server is doing.
- Some of the things you can do with the data from the reports include: **baselining, planning, troubleshooting, performance tuning**, and much more.

High Level Overview of Data Collector



Data Collection Sets

- The Data Collector includes **three default data collection sets** it uses to collect various data from SQL Server instances:
 - Disk Usage
 - Server Activity
 - Query Activity
- If you want to collect additional performance data, **you can create your own data collectors.**

Demo

- Quick review of how the Performance Data Collector looks like from the perspective of Management Studio (SSMS).

Benefits of the Data Collector

- Easy to install and configure, and comes with 2008.
- By default, collects **most of the key data** needed by a DBA to resolve many different SQL Server-related problems, especially performance-related ones.
- Allows you to **customize** what data is collected, in case you don't like the default data collected.
- Offers the flexibility for you to **manually query**, or **create Reporting Services reports**, on the historical data stored in the management data warehouse.

Limitations of the Data Collector

- It can only collect data from SQL Server 2008 (not previous versions).
- Memory, CPU and disk I/O resources are consumed on the SQL Server instances being monitored.
- In most cases, a dedicated SQL Server instance is required, if many instances are to be monitored.
- Customizing data collection sets and reports is not particularly easy.
- Once installed, it can't be uninstalled, only disabled.

How Much Overhead is Used by the Data Collector

- The Data Collector will add about 4% to the current CPU load of the SQL Server instance being monitored.
- Between 60MB and 300MB of RAM, depending on what is going on, is also required on the instances being monitored.
- Using default data collection settings, the Data Collector can generate about 250-300 MB of data *every day* per SQL Server instance being monitored, which is stored in the MDW.

How to Read & Interpret Data Collector Reports

- Three key reports (with drill-down):
 - Disk Usage Report
 - Server Activity Report
 - Query Statistics Report
- Demo reports, with special emphasis on how to use the Query Statistics Report to analyze query performance.

Take Homes for Today

- The SQL Server 2008 Data Collector has a lot of potential for helping DBAs track performance data over time.
- To take full advantage of the tool, you will need to take some time and effort to master it, especially if you want to create your own data collection sets.
- We still don't fully understand the amount of resources needed by the Data Collector and how this will affect production SQL Server instances. Don't put into production until you have tested it.

Q & A

- Please ask your questions clearly and loudly.
- If you don't get your questions answered now, see me after the session, or e-mail me.

Find Out More

Free E-Books:

www.sqlservercentral.com/Books

Check these out:

www.SQLServerCentral.com

www.Simple-Talk.com

Contact me at:

bradmcgehee@hotmail.com

Blogs:

www.bradmcgehee.com

www.twitter.com/bradmcgehee

[Click Here for a free 14-day trial of the Red Gate SQL Server Toolbelt](#)