Cardinality Estimates in SQL Server 2014

Thomas LaRock
Head Geek
Solarwinds

Why Are You Here?

- You are thinking about migrating/upgrading to SQL 2014 at some point

Why Are You Here?

- You are thinking about migrating/upgrading to SQL 2014 at some point
- You've been *told* to migrate to SQL 2014

Why Are You Here?

- You are thinking about migrating/upgrading to SQL 2014 at some point
- You've been *told* to migrate to SQL 2014
- You want to understand more about specific changes to cardinality estimates

A Little About Me…

Microsoft CERTIFIED Master | SQL Server 2008

MVP Microsoft Most Valuable Professional

sqlrockstar@thomaslarock.com

PASS

linkedin.com/in/sqlrockstar

@sqlrockstar

CONFIO SOFTWARE
part of the SolarWinds family

solarwinds

Agenda

- Why cardinality matters
- What's new in SQL 2014
- See the new CE in action
- Questions

solarwinds

Cardinality: Why It Matters

solarwinds

How most people see SQL Server

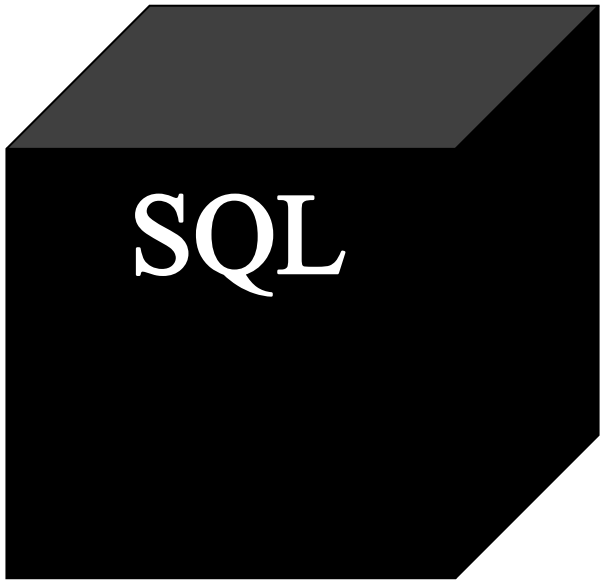

How most people see SQL Server

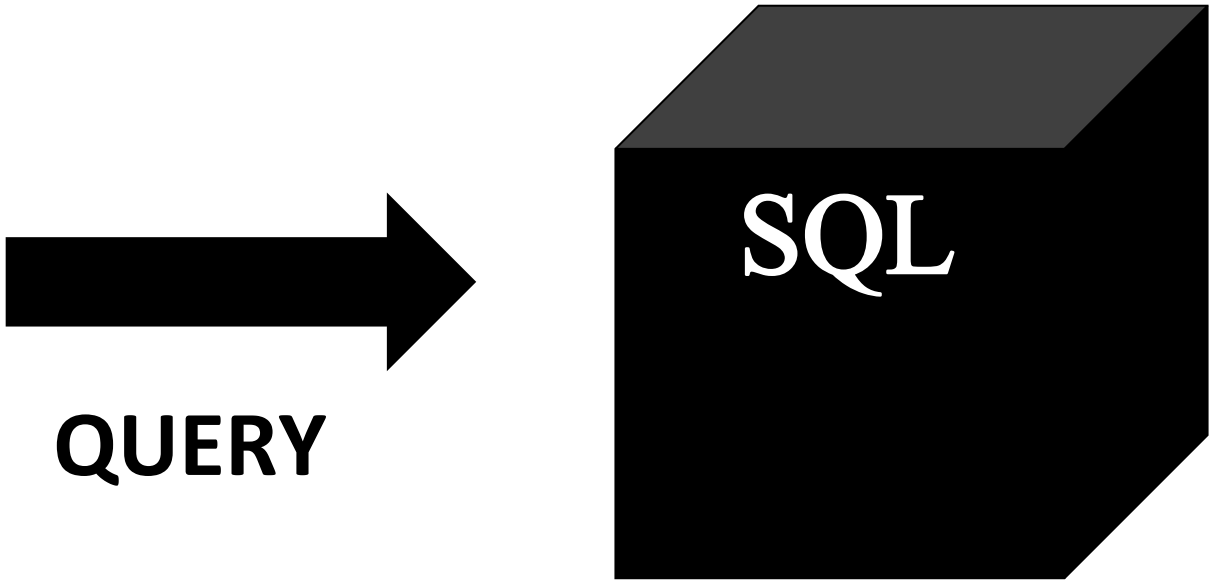

How most people see SQL Server

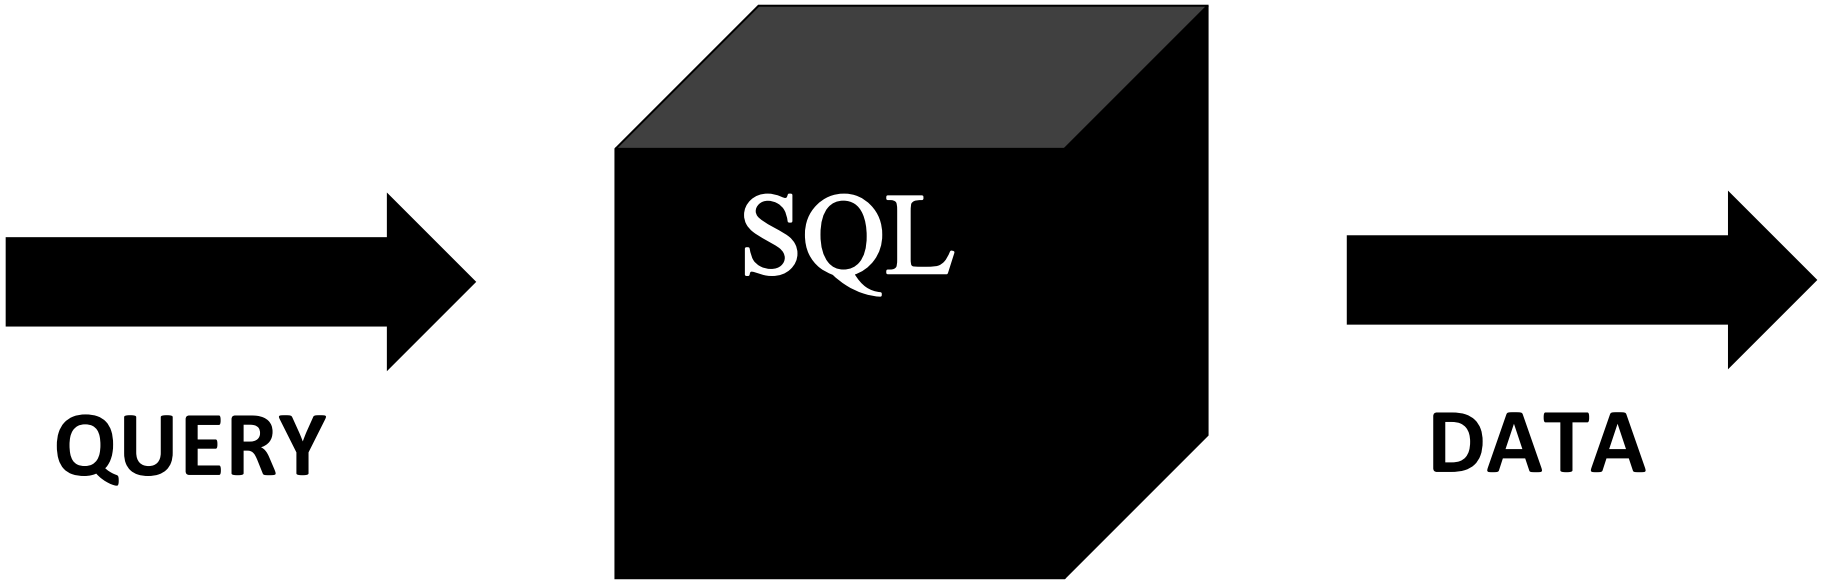

What's inside the box?

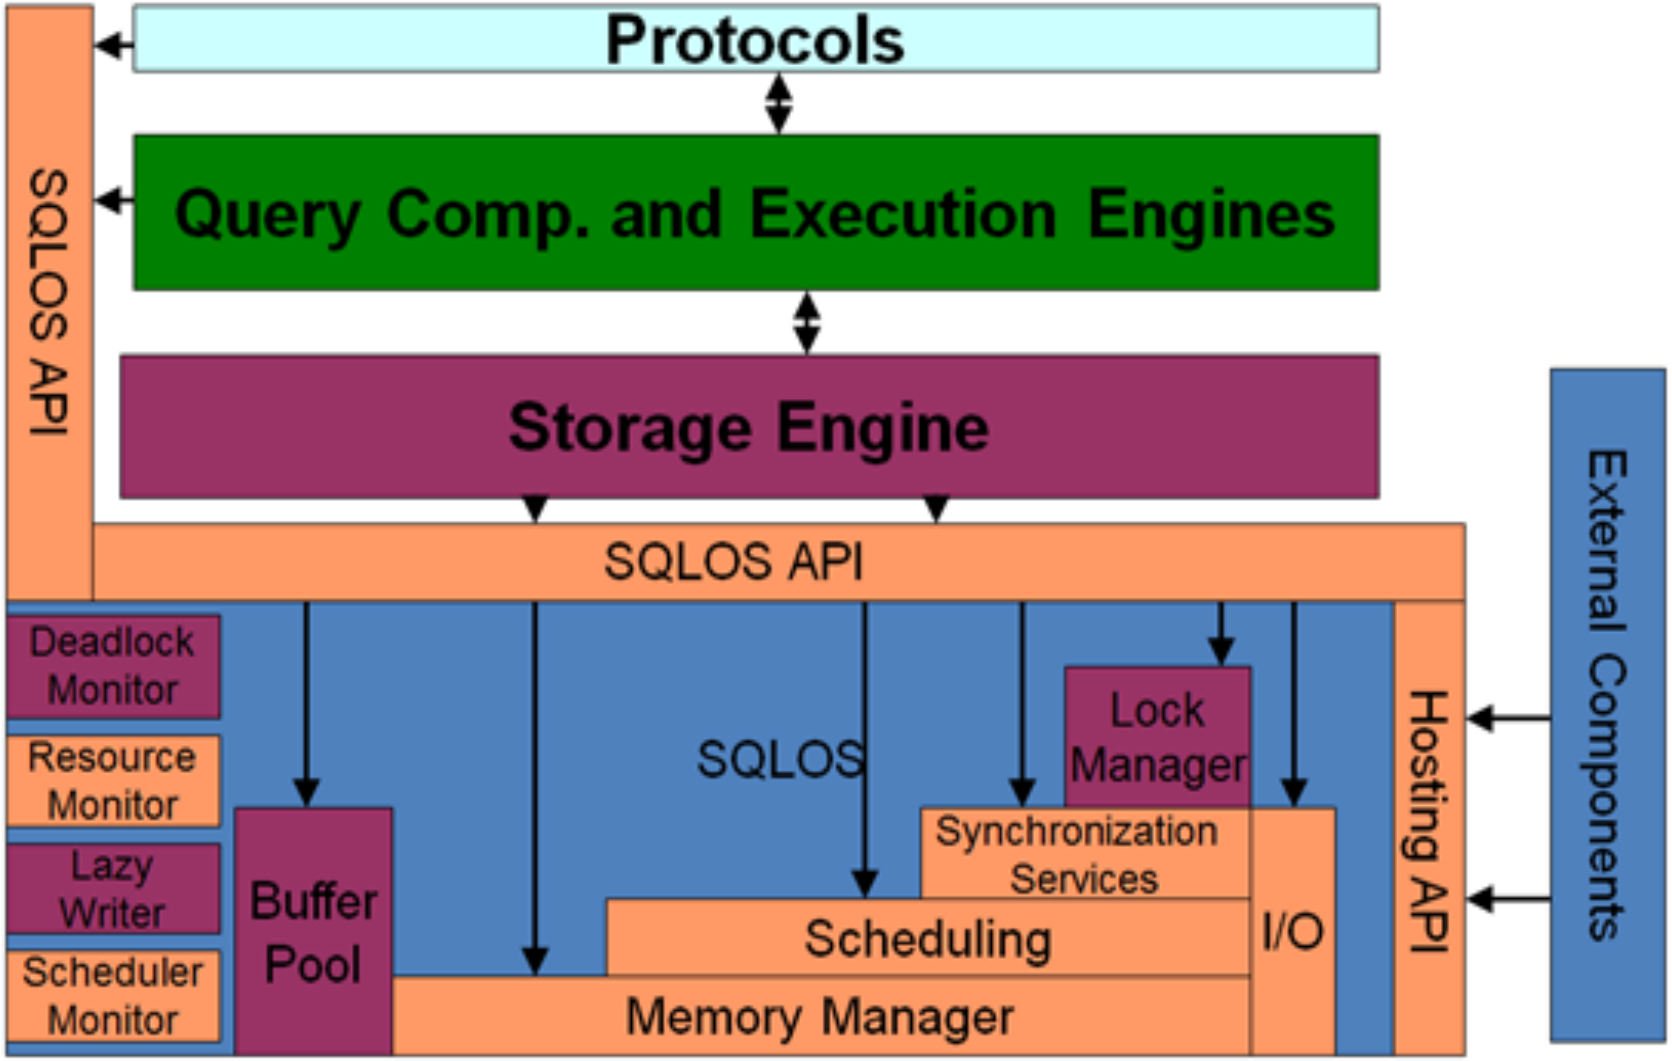

solarwinds

What Happens When a Connection Is Made to SQL Server

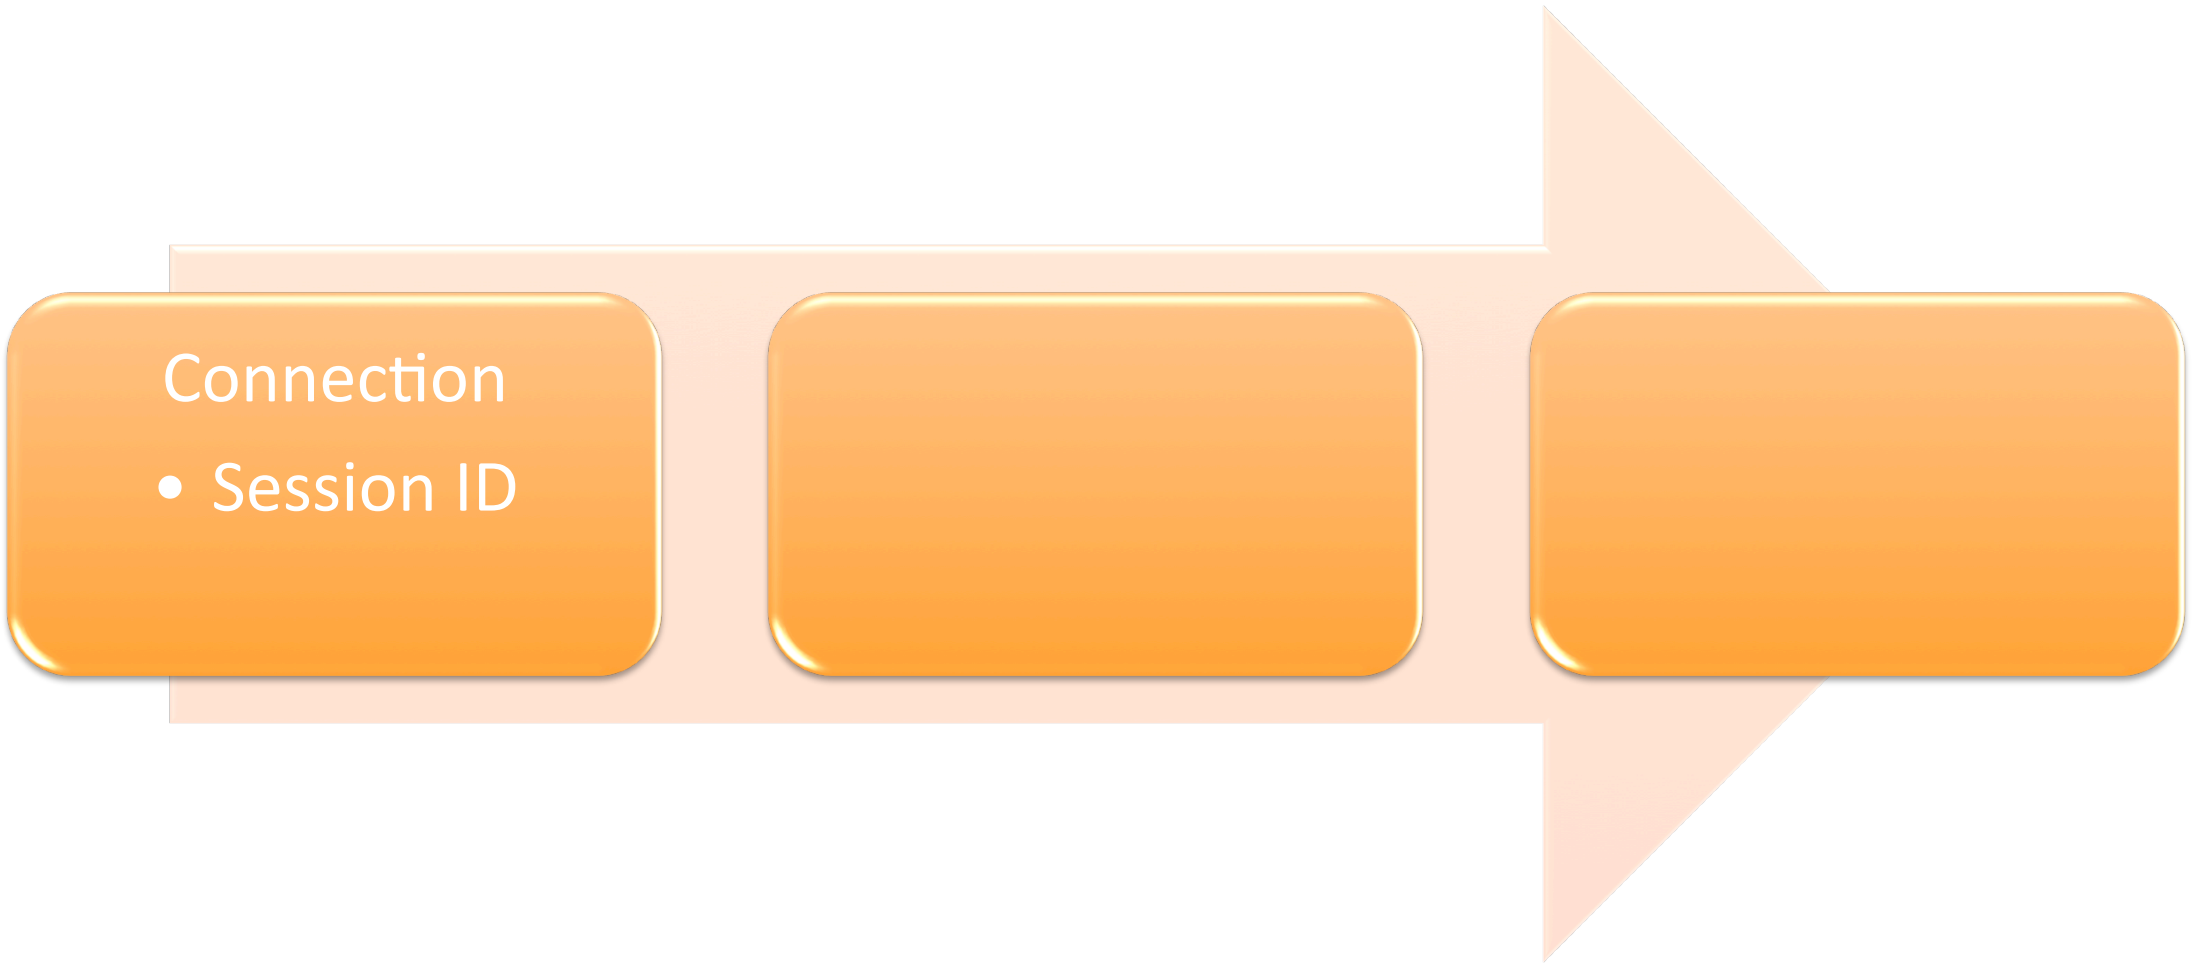

What Happens When a Connection Is Made to SQL Server

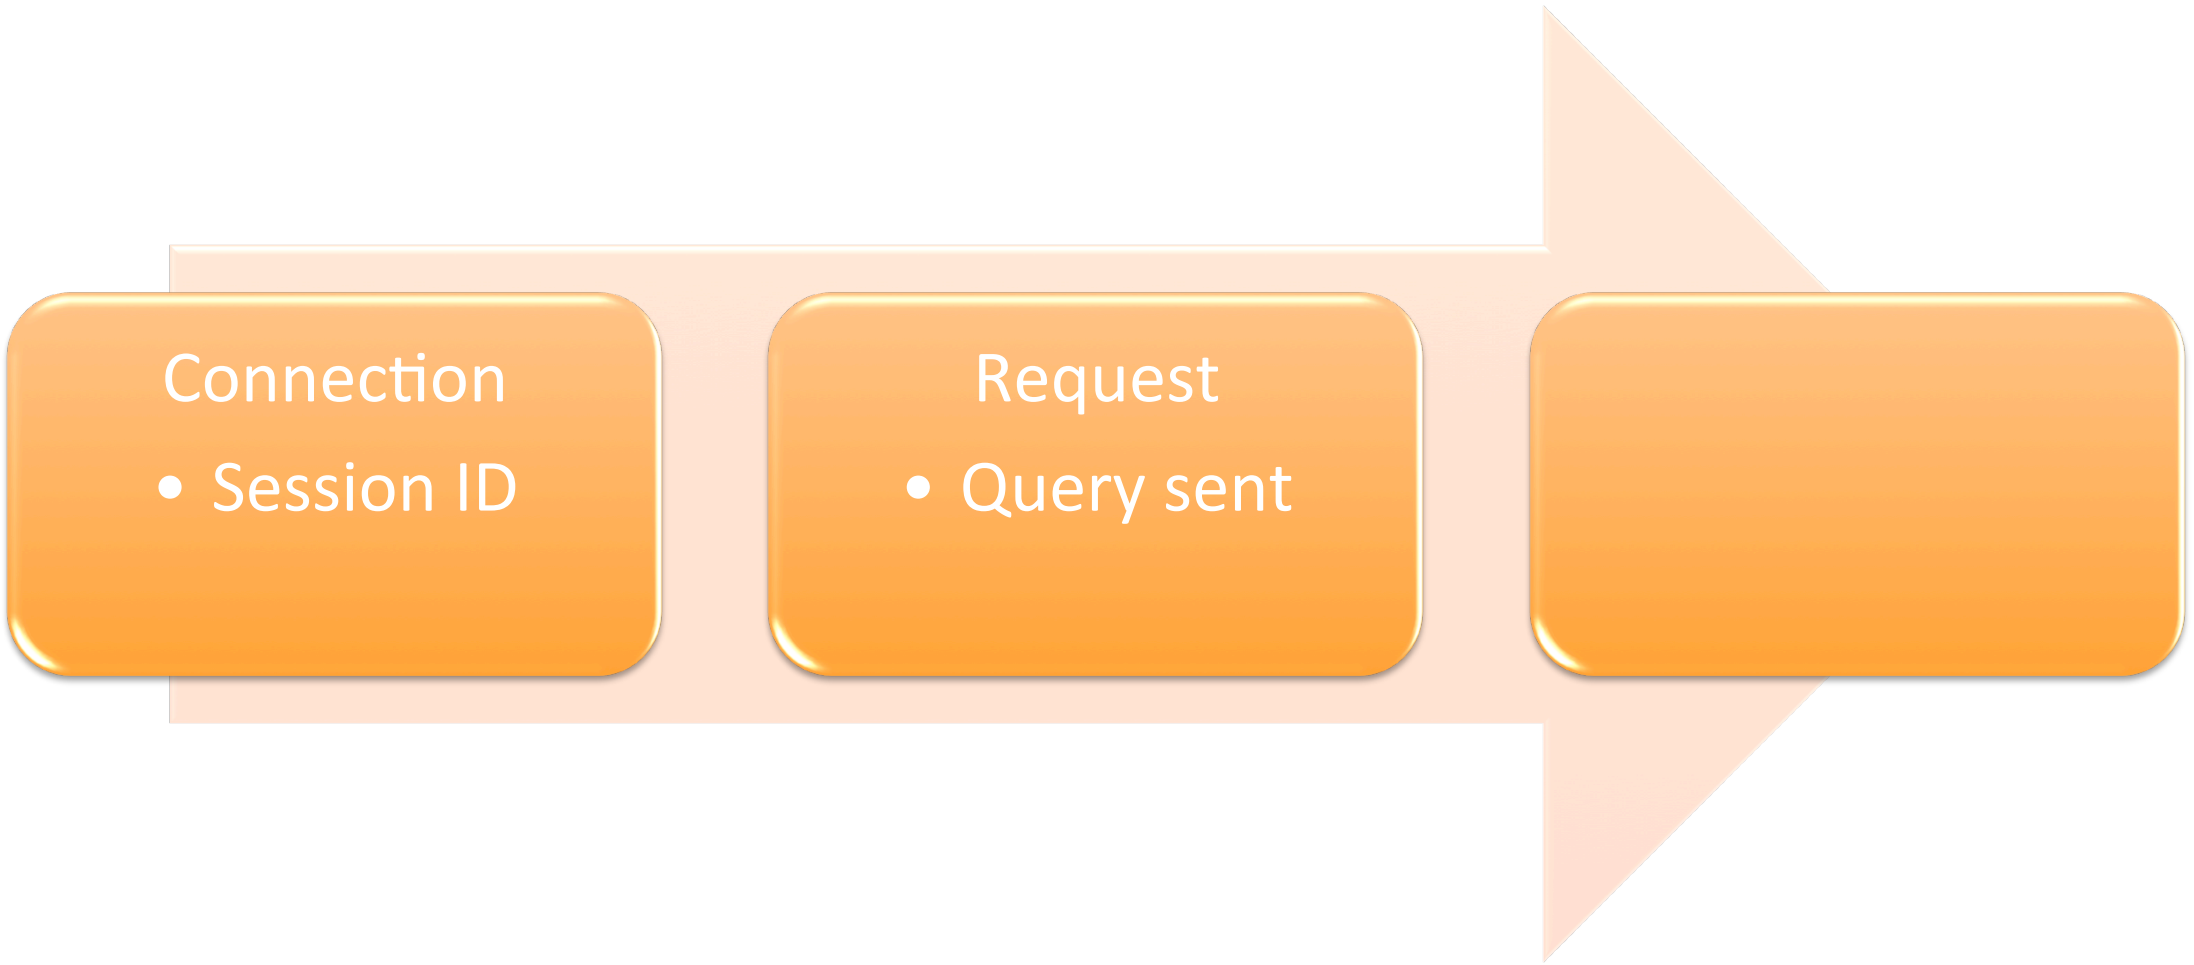

What Happens When a Request Is Sent To SQL Server

Parse

Syntax check; reserved keywords, column and table names

What Happens When a Request Is Sent To SQL Server

Parse

Syntax check; reserved keywords, column and table names

Bind

Name reservation, type derivation, aggregates, grouping

What Happens When a Request Is Sent To SQL Server

Parse

Syntax check; reserved keywords, column and table names

Bind

Name reservation, type derivation, aggregates, grouping

Optimize

Find “good enough” plan

What Happens When a Request Is Sent To SQL Server

Parse

Syntax check; reserved keywords, column and table names

Bind

Name reservation, type derivation, aggregates, grouping

Optimize

Find “good enough” plan

Execute

Storage Engine

What Happens When a Request Is Sent To SQL Server

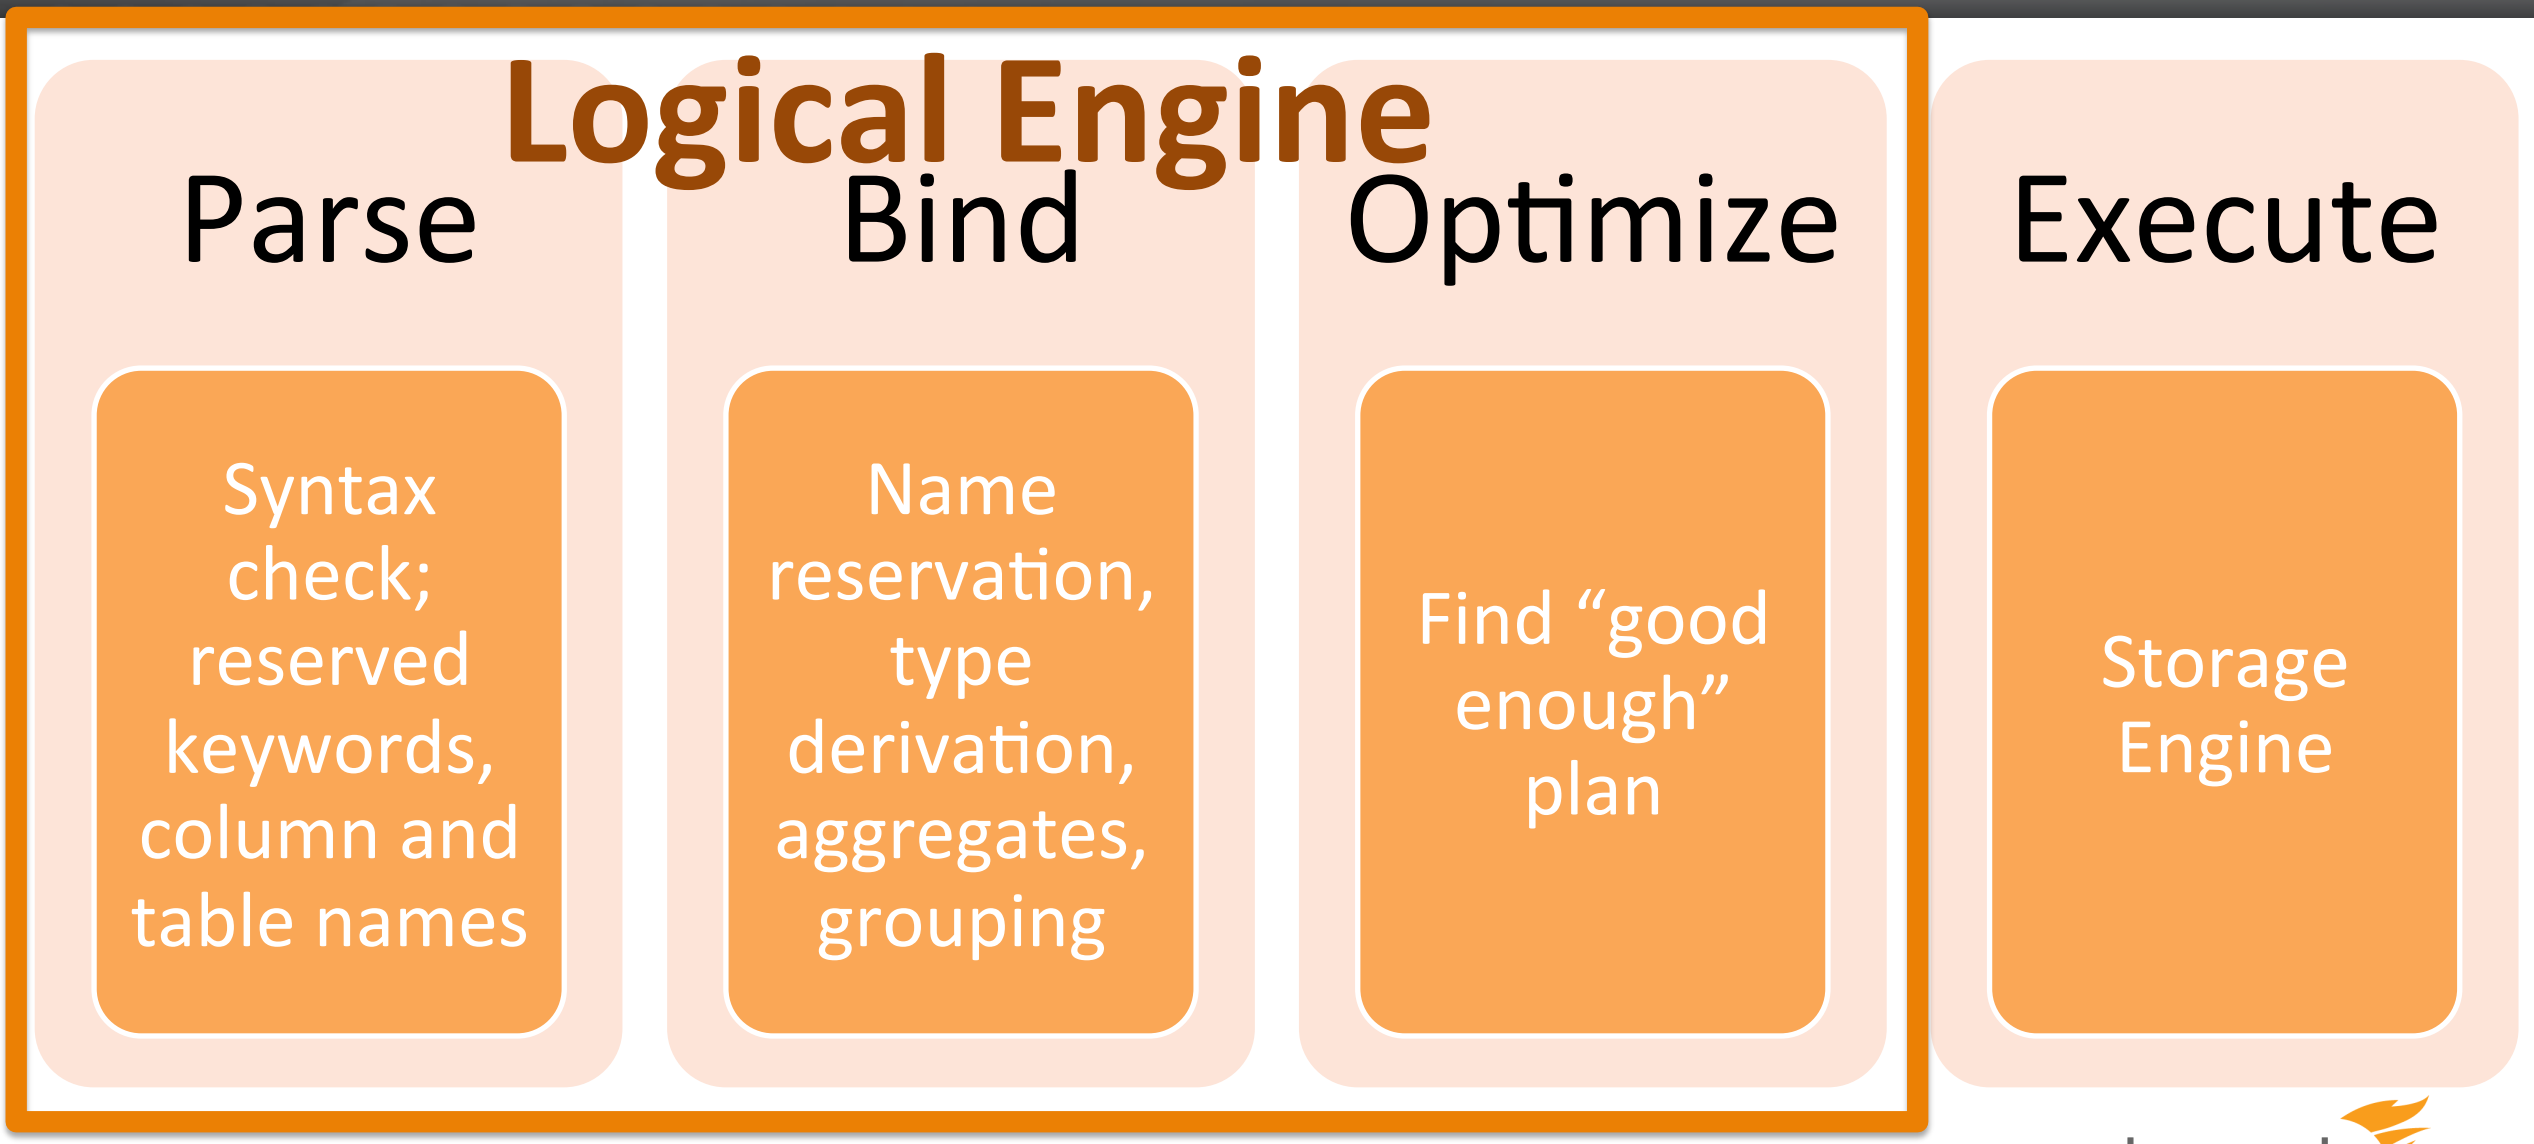

solarwinds

What Happens When a Request Is Sent To SQL Server

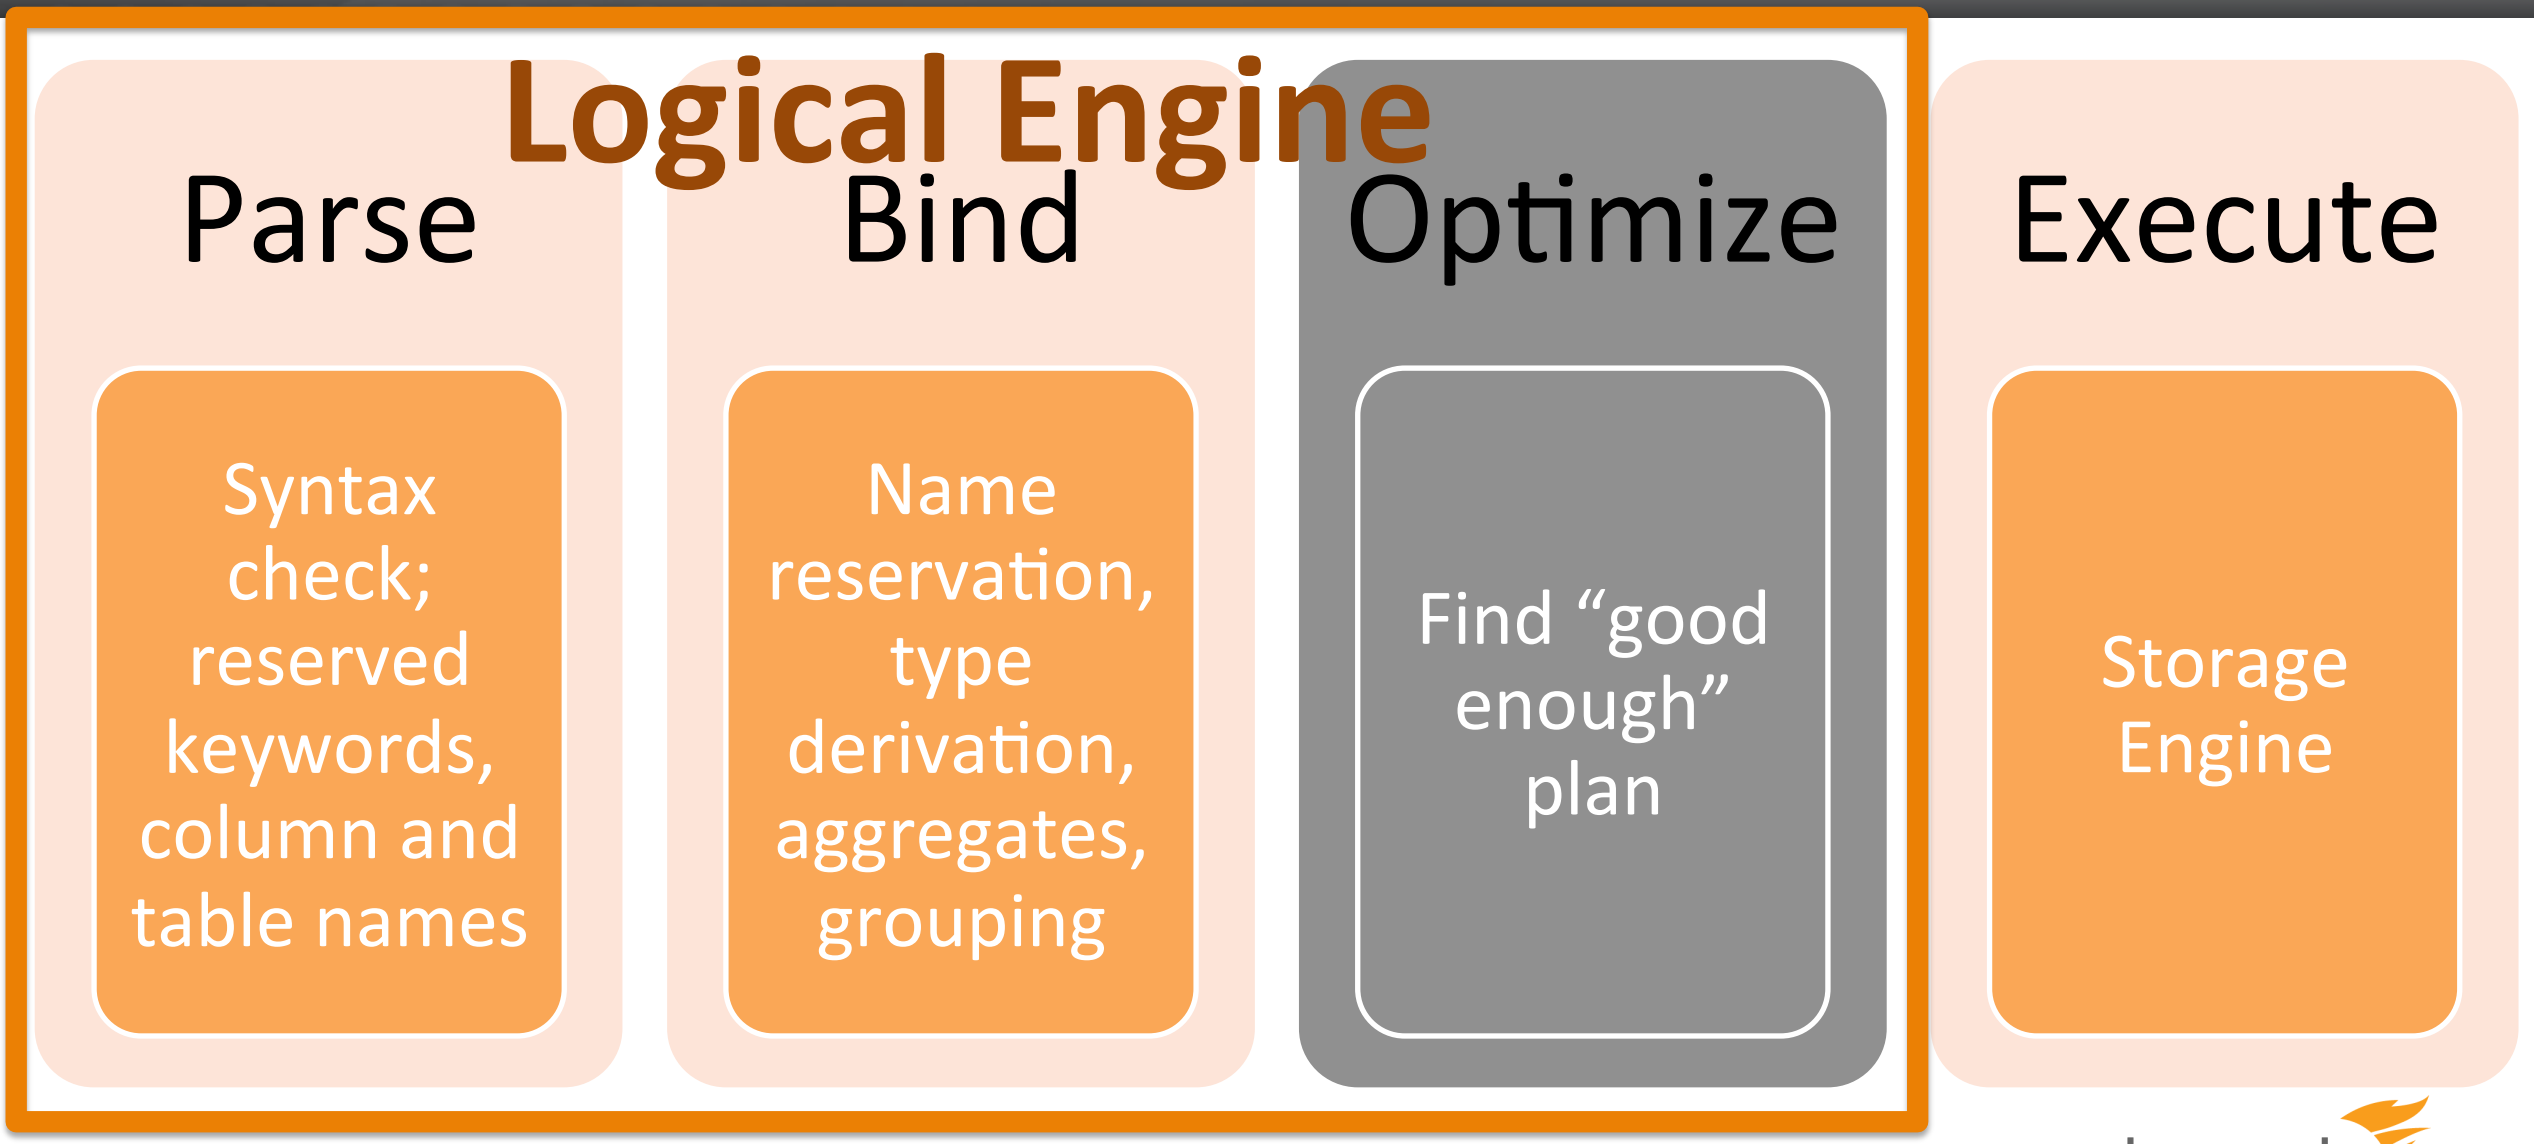

solarwinds

What Happens When a Connection Is Made to SQL Server

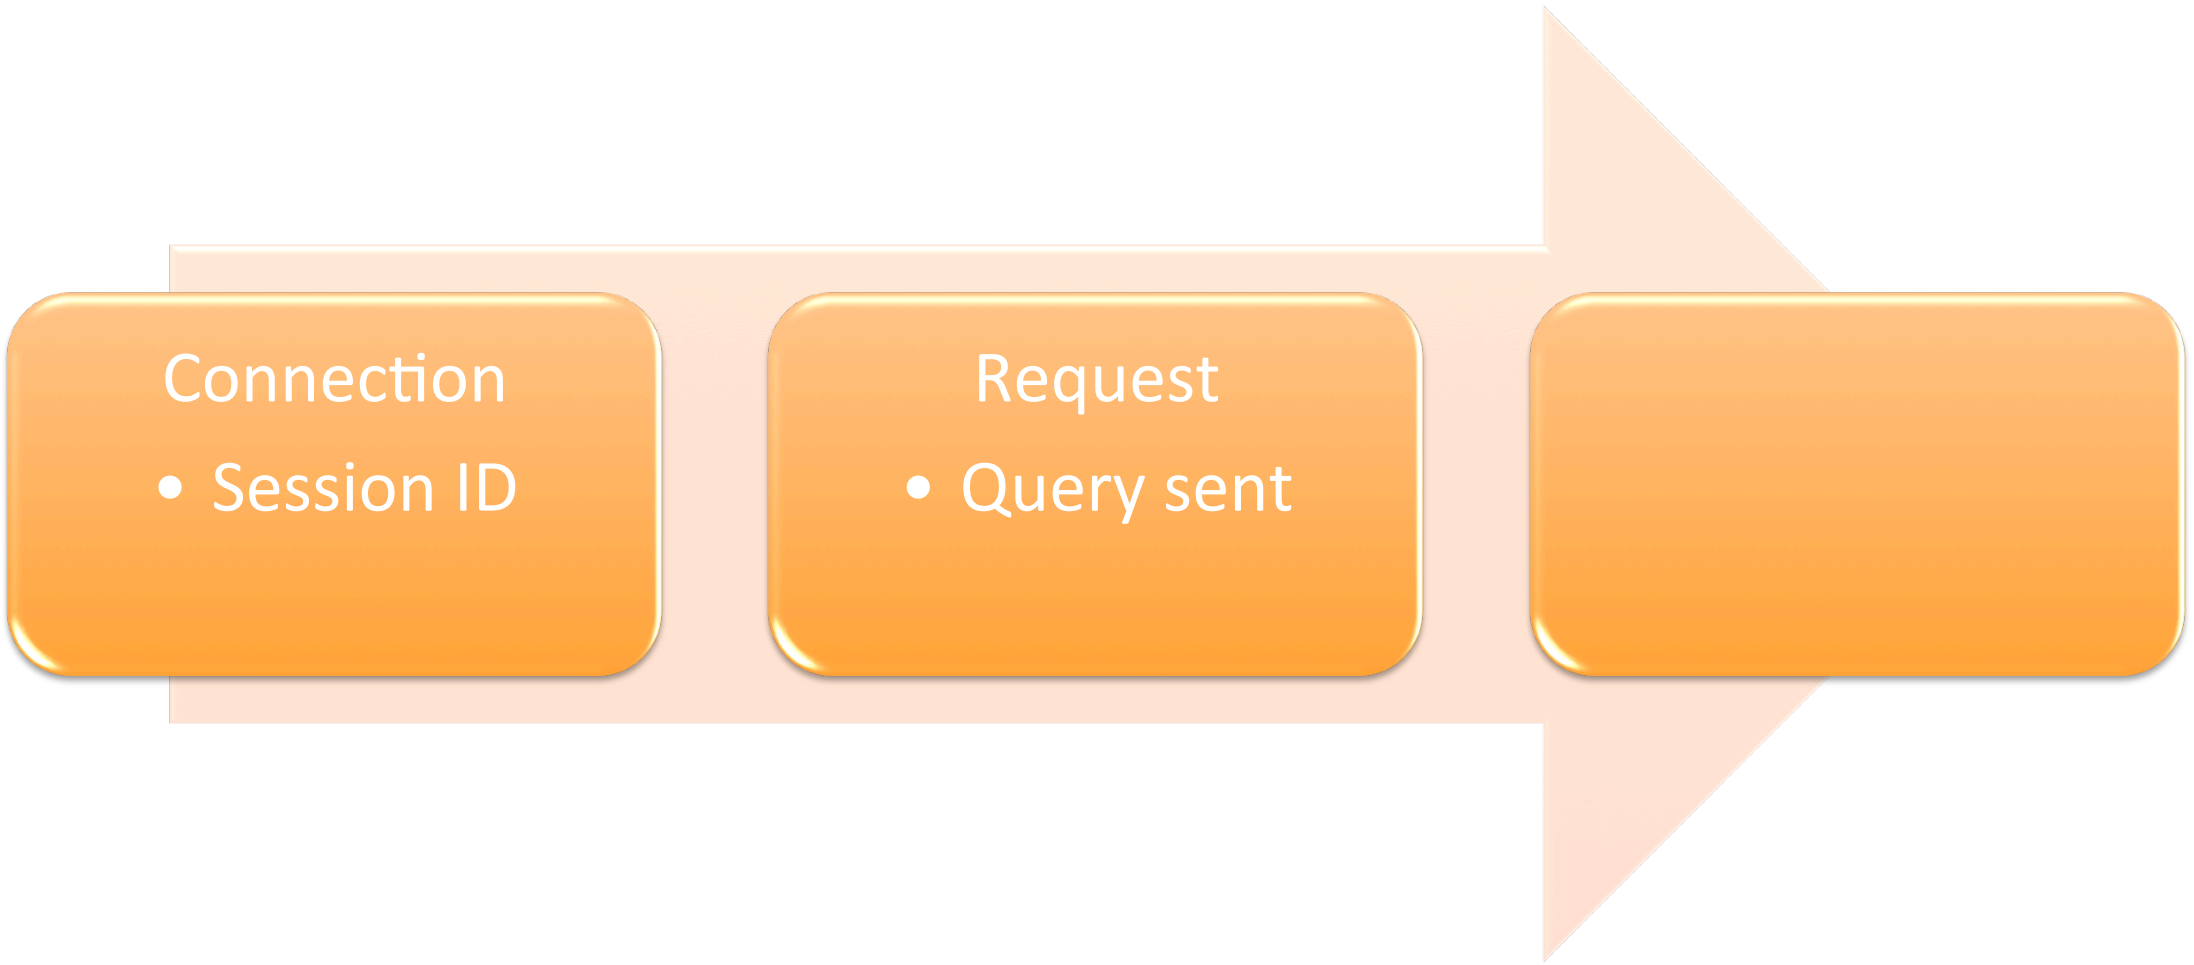

What Happens When a Connection Is Made to SQL Server

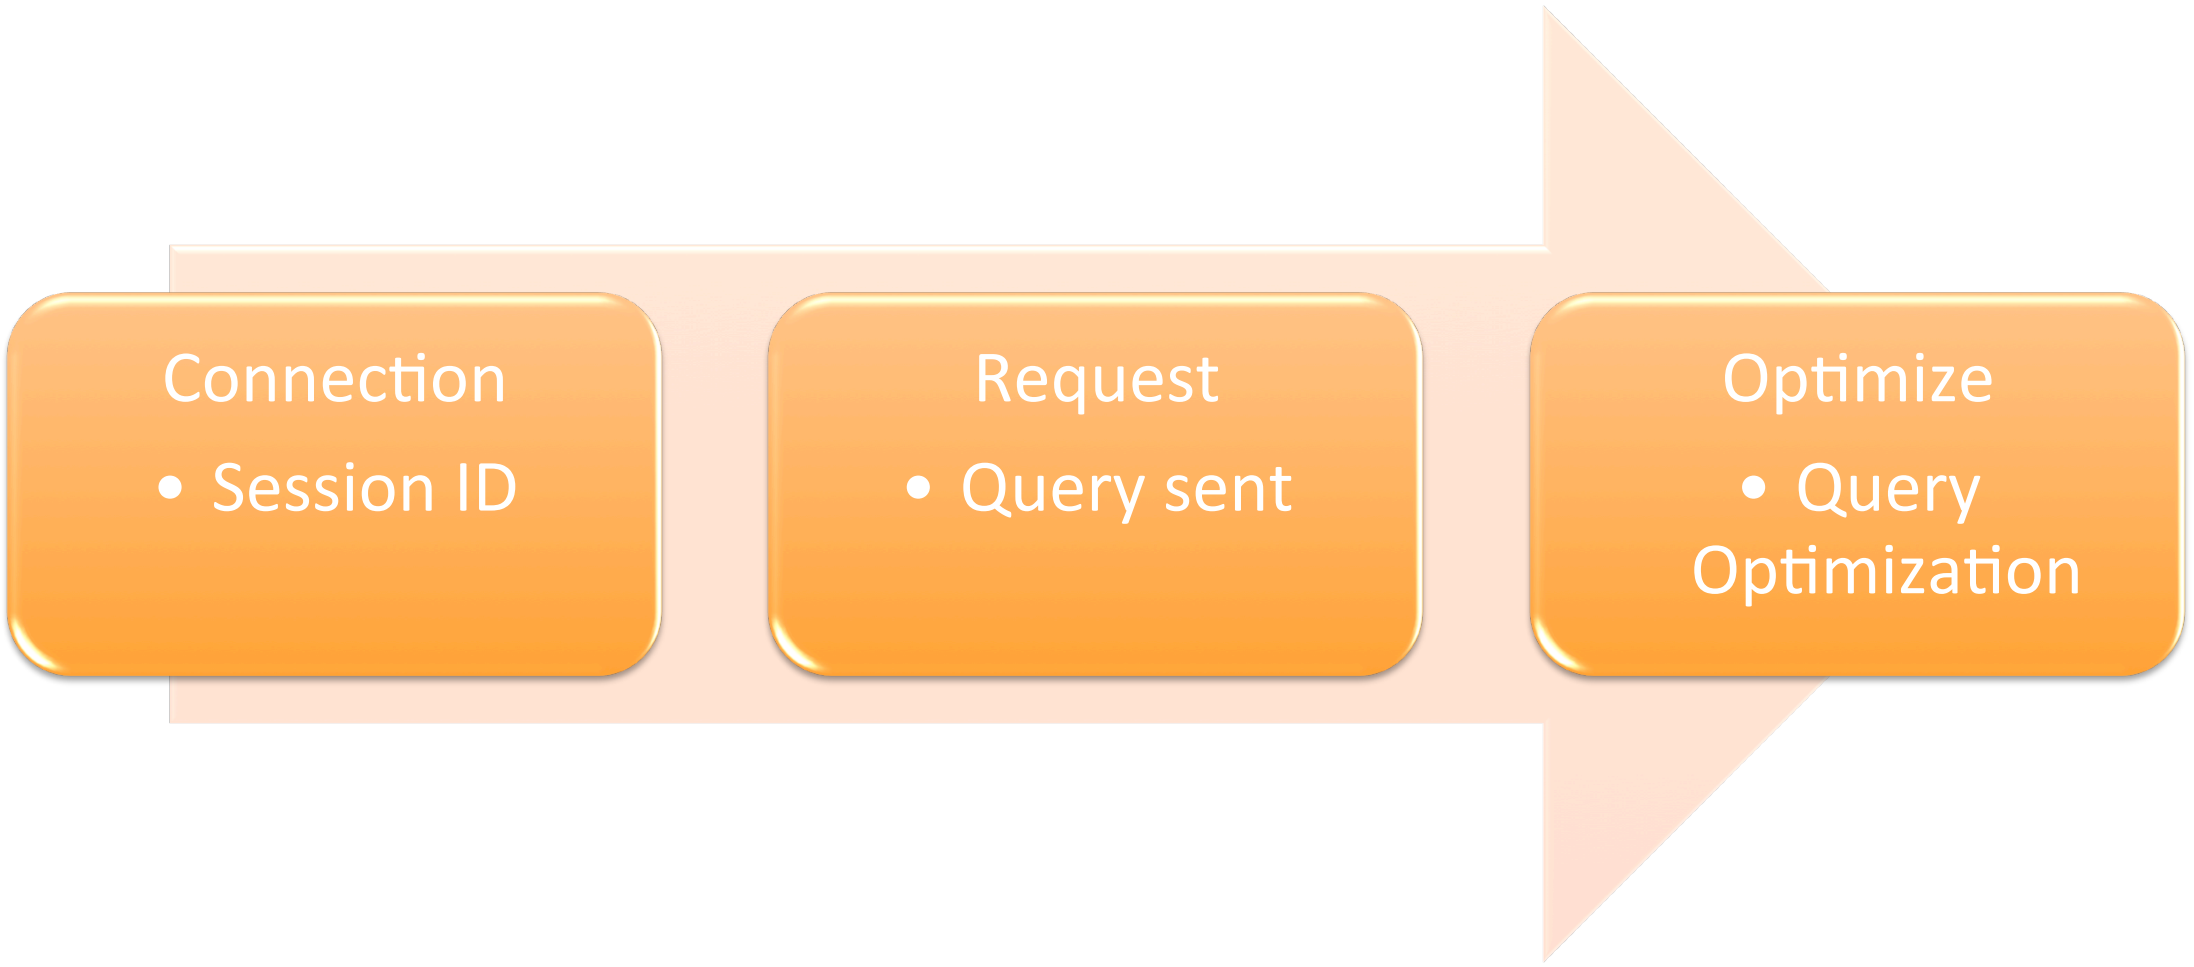

What Happens During Query Optimization in SQL Server

- Primary goal: To find plan with “least cost”

What Happens During Query Optimization in SQL Server

- Primary goal: To find plan with “least cost”
- Will stop when “good enough” plan is found

What Happens During Query Optimization in SQL Server

- Primary goal: To find plan with “least cost”
- Will stop when “good enough” plan is found
- Uses statistics(!) to estimate rows needed for operators in the execution plan

What Happens During Query Optimization in SQL Server

- Primary goal: To find plan with “least cost”
- Will stop when “good enough” plan is found
- Uses statistics(!) to estimate rows needed for operators in the execution plan
- The cost is an abstraction, no correlation to performance (ie., CPU ticks)

What Happens During Query Optimization in SQL Server

- Primary goal: To find plan with “least cost”
- Will stop when “good enough” plan is found
- Uses statistics(!) to estimate rows needed for operators in the execution plan
- The cost is an abstraction, no correlation to performance (ie., CPU ticks)

Query Optimization Model

Contents are a mystery

solarwinds

Query Optimization Model

Contents are a mystery (except to Conor and his team)

solarwinds

Query Optimization Model

Contents are a mystery (except to Conor and his team)

You don’t need the details

solarwinds

Query Optimization Model

Contents are a mystery (except to Conor and his team)

You don’t need the details (no, you don’t)

solarwinds

Query Optimization Model

Contents are a mystery (except to Conor and his team)

You don’t need the details (no, you don’t)

You need the bigger picture

solarwinds

The bigger picture

1. Does plan already exist? In not, continue to next step.

The bigger picture

1. Does plan already exist? In not, continue to next step.
2. Is this a trivial plan? If not, continue.

The bigger picture

1. Does plan already exist? In not, continue to next step.
2. Is this a trivial plan? If not, continue.
3. Apply simplification.

The bigger picture

1. Does plan already exist? In not, continue to next step.
2. Is this a trivial plan? If not, continue.
3. Apply simplification.
4. Is this plan cheap enough? If not, continue.

The bigger picture

1. Does plan already exist? In not, continue to next step.
2. Is this a trivial plan? If not, continue.
3. Apply simplification.
4. Is this plan cheap enough? If not, continue.
5. Start cost-based optimization

The bigger picture

1. Does plan already exist? In not, continue to next step.
2. Is this a trivial plan? If not, continue.
3. Apply simplification.
4. Is this plan cheap enough? If not, continue.
5. **Start cost-based optimization**

Cost-based optimization

- Explore basic rules

Cost-based optimization

- Explore basic rules
- Does plan have a cost less than 0.2?

Cost-based optimization

- Explore basic rules
- Does plan have a cost less than 0.2?
  - Explore more rules, alternate join ordering. If the best plan is less than 1.0, use that plan. If not, but MAXDOP is > 0 and this is an SMP system and the min cost > cost for parallelism, then use a parallel plan. Compare the cost of parallel plan to best serial plan and pass the cheaper of the plans along.

Cost-based optimization

- Explore basic rules
- Does plan have a cost less than 0.2?
  - Explore more rules, alternate join ordering. If the best plan is less than 1.0, use that plan. If not, but MAXDOP is > 0 and this is an SMP system and the min cost > cost for parallelism, then use a parallel plan. Compare the cost of parallel plan to best serial plan and pass the cheaper of the plans along.
- Opt for the cheapest plan after a limited number of explorations

Why This Matters

Query optimizer contains the Cardinality Estimator (CE)

$$\sum_{0 \le i \le m} \; \prod_{0 < j < n} P(i,j)$$

Why This Matters

Query optimizer contains the Cardinality Estimator (CE)

The CE is dependent upon statistics being current and accurate

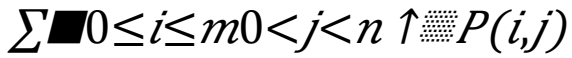

solarwinds

Why This Matters

Query optimizer contains the Cardinality Estimator (CE)

The CE is dependent upon statistics being current and accurate

If stats are missing, incorrect, or invalid then you are likely to get a sub-optimal plan

$$\sum_{0 \le i \le m} \prod_{0 < j < n} P(i,j)$$

solarwinds

What's inside the box?

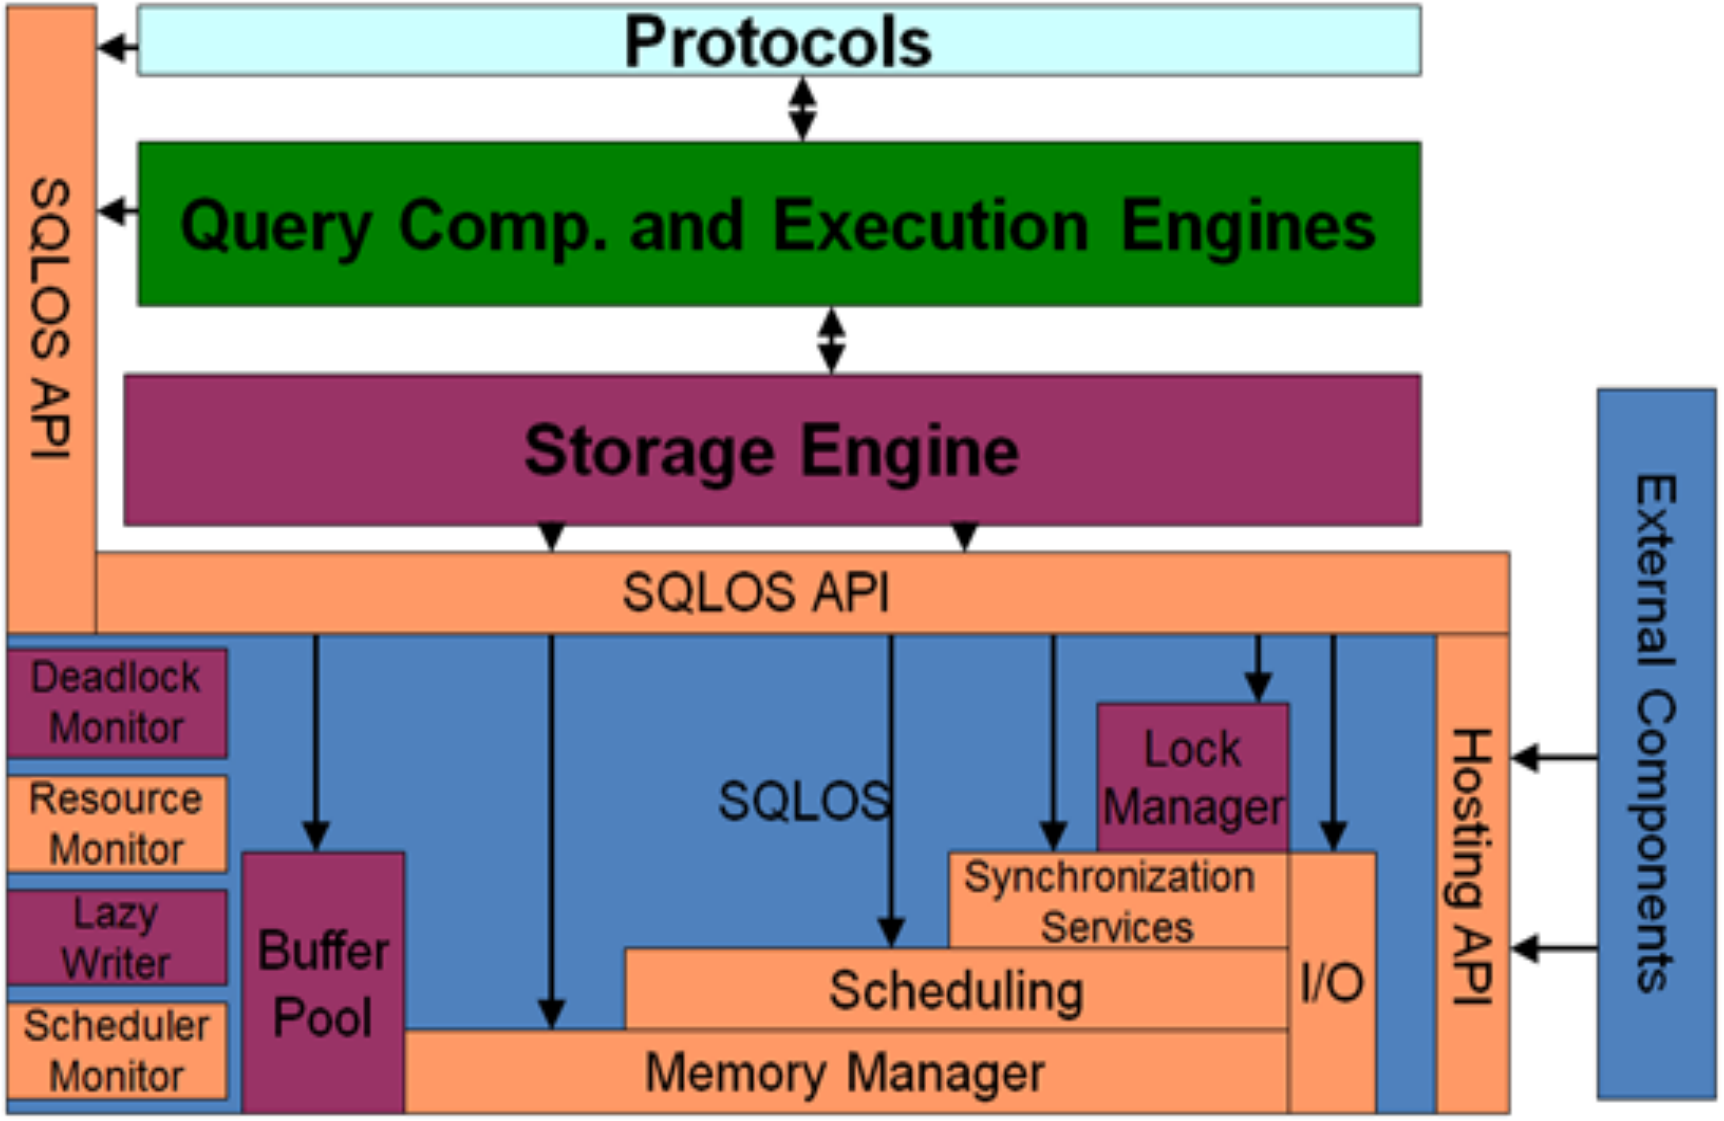

rwinds

The CE is the DNA for your queries!

- DNA is what makes you you
- DNA not easily seen or understood
- CE is similar to DNA
- Tiny but BIG role
- Neither DNA nor CE are trivial pieces of code

What's New In SQL Server 2014

solarwinds

The CE has been unchanged since 7.0!

- There have been slight enhancements over the years
  - Hotfixes, patches, trace flags

The CE has been unchanged since 7.0!

- There have been slight enhancements over the years
  - Hotfixes, patches, trace flags
- Goal for SQL2014 was not to avoid regression

The CE has been unchanged since 7.0!

- There have been slight enhancements over the years
  - Hotfixes, patches, trace flags
- Goal for SQL2014 was not to avoid regression
- It is a new CE
  - Will work better for most queries
  - May not work well for some
  - Testing is vital!

Legacy CE assumptions

Uniformity

Independence

Containment

Inclusion

solarwinds

Legacy CE assumptions

Uniformity

Independence

Containment

Inclusion

solarwinds

Legacy CE assumptions

Uniformity

Independence

Containment

Inclusion

solarwinds

Legacy CE assumptions

- Uniformity
- Independence
- Containment
- Inclusion

solarwinds

See the New CE in Action

DEMO

solarwinds

Likely common scenarios

- ASC/DESC Key Estimates
- Multiple Column Dependency
- Join Estimates

Likely common scenarios

- ASC/DESC Key Estimates
- Multiple Column Dependency
- Join Estimates

Likely common scenarios

- ASC/DESC Key Estimates
- Multiple Column Dependency
- Join Estimates

New CE Changes That Require Action

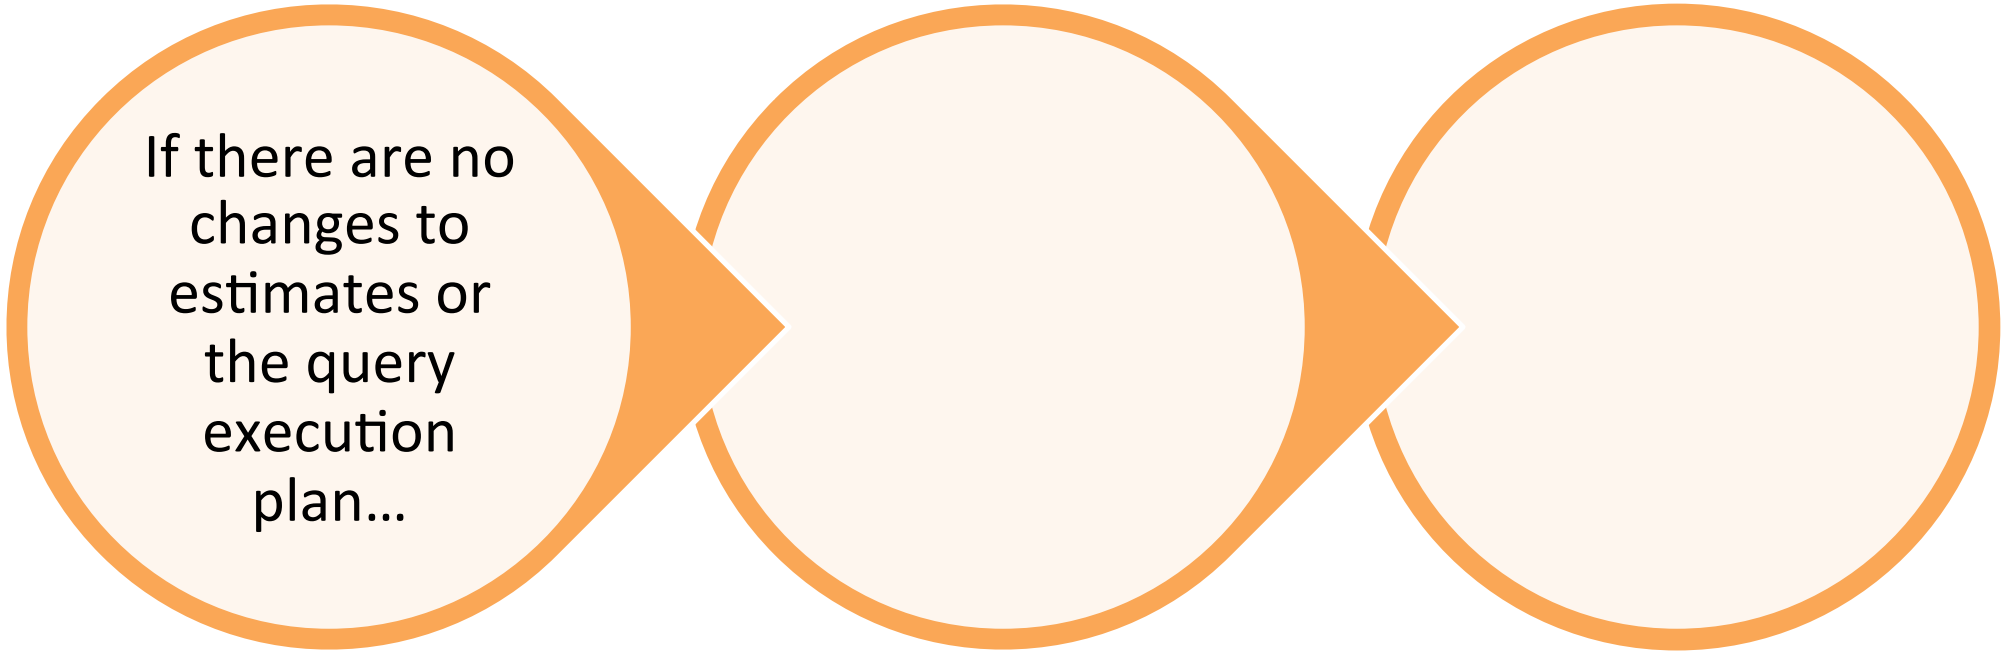

New CE Changes That Require Action

New CE Changes That Require Action

New CE Changes That Require Action

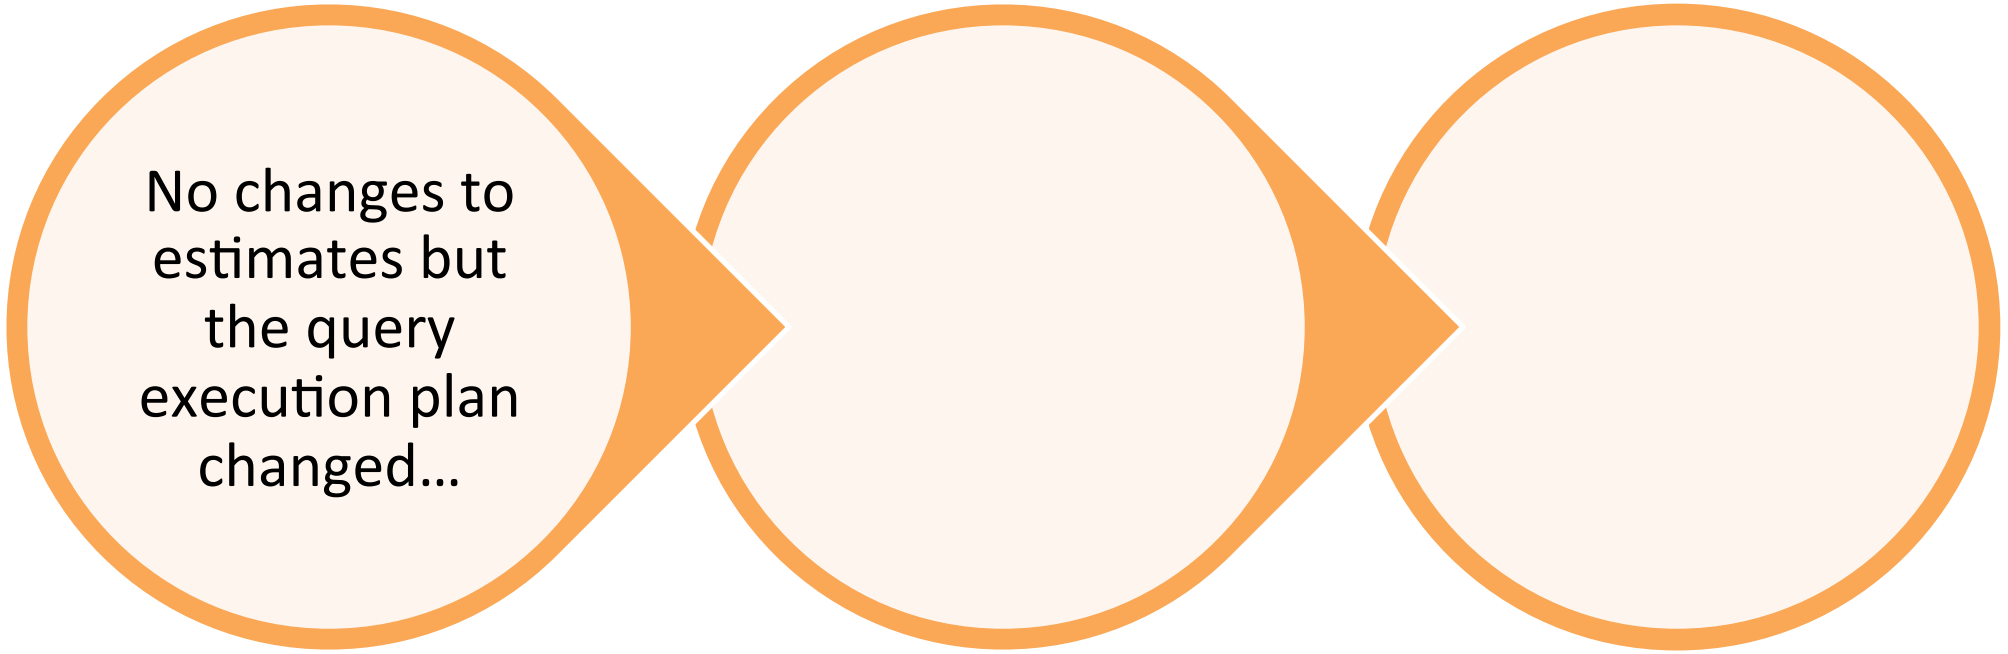

New CE Changes That Require Action

New CE Changes That Require Action

New CE Changes That Require Action

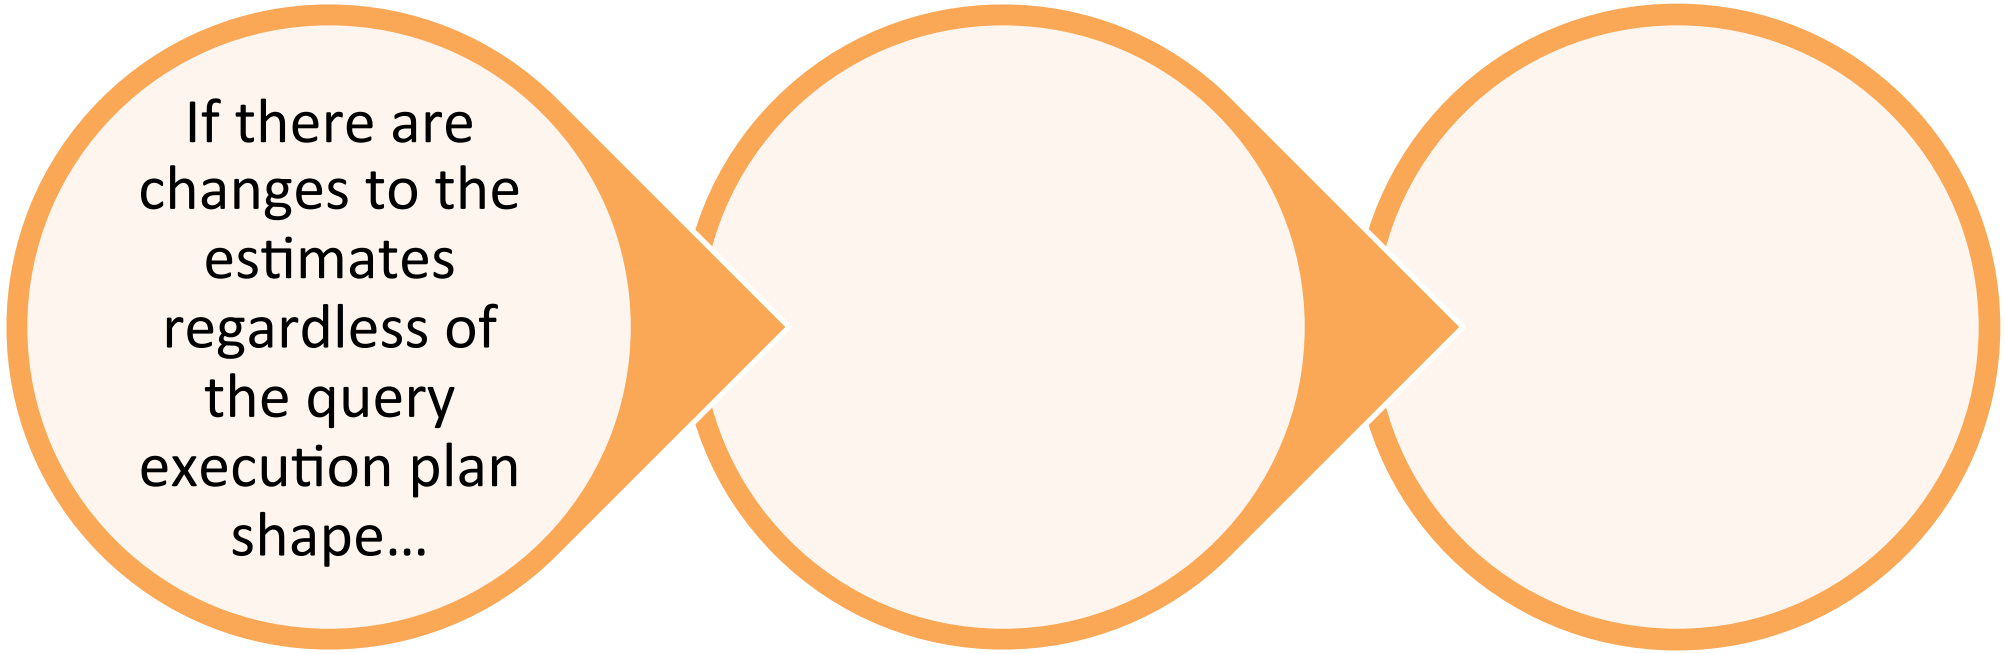

New CE Changes That Require Action

New CE Changes That Require Action

Tips & Tricks

solarwinds

Tips and Tricks

- Auto Create/Update stats
-
-
-

solarwinds

Tips and Tricks

- Auto Create/Update stats
- Database compatibility mode
-
-

solarwinds

Tips and Tricks

- Auto Create/Update stats
- Database compatibility mode
- Query trace flags
-

solarwinds

Tips and Tricks

- Auto Create/Update stats
- Database compatibility mode
- Query trace flags
- XML showplan, xEvents

solarwinds

Agenda

- Why cardinality matters
- What's new in SQL 2014
- See the new CE in action
- Questions

solarwinds

Questions?

solarwinds

Additional info

- http://tinyurl.com/joe-sack-cardinality
- http://tinyurl.com/sql-2014-new-CE
- http://tinyurl.com/sql-2014-new-CE-p2
- http://tinyurl.com/CE-deep-dive-PWhite
- http://tinyurl.com/MSDN-CE-2014