The Neo4j Drivers Manual v4.1 for C#

Table of Contents

Copyright © 2021

This manual covers the following areas:

- Get started — An overview of the official Neo4j drivers and how to connect to a Neo4j database.

- Client applications — How to manage database connections within an application.

- Cypher workflow — How to create units of work and provide a logical context for that work.

- The session API — How the types and values used by Cypher map to native language types.

- Driver terminology — Terminology for drivers.

Who should read this?

This manual is written for the developer of a Neo4j client application.

Get started

About the official drivers

Neo4j provides official drivers for a number of popular programming languages. These drivers are supported by Neo4j.

Community drivers also exist for many languages, but vary greatly in terms of feature sets, maturity, and support. To find more about community drivers, visit https://neo4j.com/developer/language-guides/.

The following languages and frameworks are officially supported by Neo4j:

Table 1. Supported languages and frameworks for the 4.x driver series

Language/framework	Versions supported
.NET	.NET Standard 2.0
Go	Go 1.10
Java	Java 8+ (latest patch releases).
JavaScript	All LTS versions of Node.JS, specifically the 4.x and 6.x series runtimes (https://github.com/nodejs/LTS).
Python	CPython 3.5 and above.

The driver API is intended to be topologically agnostic. This means that the underlying database topology — single instance, Causal Cluster, etc. — can be altered without requiring a corresponding alteration to application code. In the general case, only the connection URI needs to be modified when changes are made to the topology.

> **i** The official drivers do not support HTTP communication. If you need an HTTP driver, choose one of the community drivers. See also the HTTP API documentation.

Driver versions and installation

Starting with **Neo4j 4.0**, *the versioning scheme for the database, driver and protocol have all been aligned*. This simplifies general compatibility concerns.

Cross-version compatibility is still available, and minimum support for current and previous versions between both server and driver is guaranteed. More specifically, this means that **Neo4j 4.0** is guaranteed to be compatible with both **4.0 Drivers** and **1.7 Drivers**, and the **4.0 Drivers** are guaranteed to be compatible with both **Neo4j 4.0** and **Neo4j 3.5**. In cases where at least one peer is below version 4.0, communication will occur in *fallback mode, limiting functionality to that available in the lowest-versioned component*.

Drivers 1.7 do not support multiple databases and Neo4j Fabric, features introduced in Neo4j 4.0. To be able to run multiple databases online concurrently and to do distributed queries over them, you must upgrade Drivers from 1.7 to 4.0. For information, see Upgrade and Migration Guide → Drivers changes.

Wherever possible, it is recommended to use the latest stable driver release available. This will provide the greatest degree of stability and will ensure that the full set of server functionality is available. The drivers, when used with **Neo4j Enterprise Edition**, come with full cluster routing support. The drivers make no explicit distinction between **Enterprise Edition** and **Community Edition** however, and simply operate with the functionality made available by **Neo4j** itself.

To install a driver or to find out more about which driver versions are available, use the relevant language distribution system, described below.

Example 7. Acquire the driver

The .NET driver is distributed via the NuGet Gallery. To find the latest version of the driver, visit
https://www.nuget.org/packages/Neo4j.Driver/.

Dependencies

- .NETStandard (^2.0)
- System.Net.NameResolution (^4.3.0)
- System.Net.Security (^4.3.2)
- System.Net.Sockets (^4.3.0)
- System.Runtime.InteropServices.RuntimeInformation (^4.3.0)
- System.Runtime.Serialization.Primitives (^4.3.0)
- System.Threading.Thread (^4.3.0)
- System.ValueTuple (^4.5.0)

The .NET Reactive API is under package `Neo4j.Driver.Reactive`. It is built upon
`System.Reactive`. There is no extra dependency required use .NET reactive API.

Example 1. Installation via Install-package

To install the latest version of the driver using NuGet in Visual Studio:

```
PM> Install-Package Neo4j.Driver
```

It is also an option to install a certain version of the driver.

```
PM> Install-Package Neo4j.Driver -Version $DOTNET_DRIVER_VERSION
```

The .Net driver uses asynchronous methods by default. If you want to use blocking calls,
the `Neo4j.Driver.Simple` NuGet package contains a set of extensions. This can be found at
https://www.nuget.org/packages/Neo4j.Driver.Simple/. The initial examples make use of
these extensions.

In the following example we are installing driver version 4.1.1.

```
PM> Install-Package Neo4j.Driver -Version 4.1.1
```

The release notes for this driver are available here.

A "Hello World" example

The example below shows the minimal configuration necessary to interact with **Neo4j** through a driver.

Example 8. Hello World

```csharp
public class HelloWorldExample : IDisposable
{
    private readonly IDriver _driver;

    public HelloWorldExample(string uri, string user, string password)
    {
        _driver = GraphDatabase.Driver(uri, AuthTokens.Basic(user, password));
    }

    public void PrintGreeting(string message)
    {
        using (var session = _driver.Session())
        {
            var greeting = session.WriteTransaction(tx =>
            {
                var result = tx.Run("CREATE (a:Greeting) " +
                                    "SET a.message = $message " +
                                    "RETURN a.message + ', from node ' + id(a)",
                    new {message});
                return result.Single()[0].As<string>();
            });
            Console.WriteLine(greeting);
        }
    }

    public void Dispose()
    {
        _driver?.Dispose();
    }

    public static void Main()
    {
        using (var greeter = new HelloWorldExample("bolt://localhost:7687", "neo4j",
"password"))
        {
            greeter.PrintGreeting("hello, world");
        }
    }
}
```

Driver API docs

For a comprehensive listing of all driver functionality, refer to the **API documentation** for the specific driver language.

Example 9. API docs

https://neo4j.com/docs/api/dotnet-driver/4.1/

Client applications

This section contains the following:

The Driver Object

Neo4j client applications require a *Driver Object* which, from a data access perspective, forms the backbone of the application. It is through this object that all Neo4j interaction is carried out, and it should therefore be made available to all parts of the application that require data access.

In languages where thread safety is an issue, the Driver Object can be considered *thread-safe*.

💡	*A note on lifecycle* Applications will typically construct a Driver Object on startup and destroy it on exit. Destroying a Driver Object will immediately shut down any connections previously opened via that Driver Object, by closing the associated connection pool. This will have the consequence of rolling back any open transactions, and closing any unconsumed results.

To construct a driver instance, a connection URI and authentication information must be supplied.

Additional configuration details can be supplied if required. The configuration details are immutable for the lifetime of the Driver Object. Therefore, if multiple configurations are required (such as when working with multiple database users) then multiple Driver Objects must be used.

Example 10. The driver lifecycle

```
public class DriverLifecycleExample : IDisposable
{
    public IDriver Driver { get; }

    public DriverLifecycleExample(string uri, string user, string password)
    {
        Driver = GraphDatabase.Driver(uri, AuthTokens.Basic(user, password));
    }

    public void Dispose()
    {
        Driver?.Dispose();
    }
}
```

Connection URIs

A connection URI identifies a graph database and how to connect to it.

The **encryption** and **trust** settings provide detail to how that connection should be secured.

	There are significant changes to security settings between Neo4j 3.x and Neo4j 4.x
🔥	Please consider the information in this section before upgrading from a previous version.
	The Migration Guide is also a good source of information about changes in functionality.

Starting with **Neo4j 4.0**, client-server communication uses only *unencrypted local connections by default*.

This is a change from previous versions, which switched on encryption by default, but generated a self-signed certificate out of the box.

When a full certificate is installed, and encryption is enabled on the driver, full certificate checks are carried out (refer to Operations Manual → SSL framework). Full certificates provide better overall security than self-signed certificates as they include a complete chain of trust back to a root certificate authority.

	Neo4j Aura is a *secure hosted service* backed by full certificates signed by a root certificate authority.
ℹ️	To connect to **Neo4j Aura**, driver users must *enable encryption* and the complete set of certificate checks (the latter of which are enabled by default).
	For more information, see Examples below.

Table 2. Changes in default security settings between 3.x and 4.x

Setting	Neo4j 4.x	Neo4j 3.x (Drivers 1.x)
Bundled certificate	none	auto-generated, self-signed

Setting	Neo4j 4.x	Neo4j 3.x (Drivers 1.x)
Driver encryption	off	on
Bolt interface	localhost	localhost
Certificate expiry check	on	on
Certificate CA check	on	off
Certificate hostname check	on	off

Initial address resolution

The address provided in a `neo4j://` URI is used for initial and fallback communication only.

This communication occurs **to bootstrap the routing table**, through which all subsequent communication is carried out. Fallback occurs when the driver is unable to contact any of the addresses held in the routing table. The initial address is once again reused to bootstrap the system.

Several options are available for providing this initial logical-to-physical host resolution. These include *regular DNS*, *custom middleware* such as a load balancer, and the Driver Object *resolver function*, all of which are described in the following sections.

DNS resolution

DNS resolution is the default, and always-available option. As it is possible to configure DNS to resolve a single host name down to multiple IP addresses, this can be used to expose all core server IP addresses under a single host name.
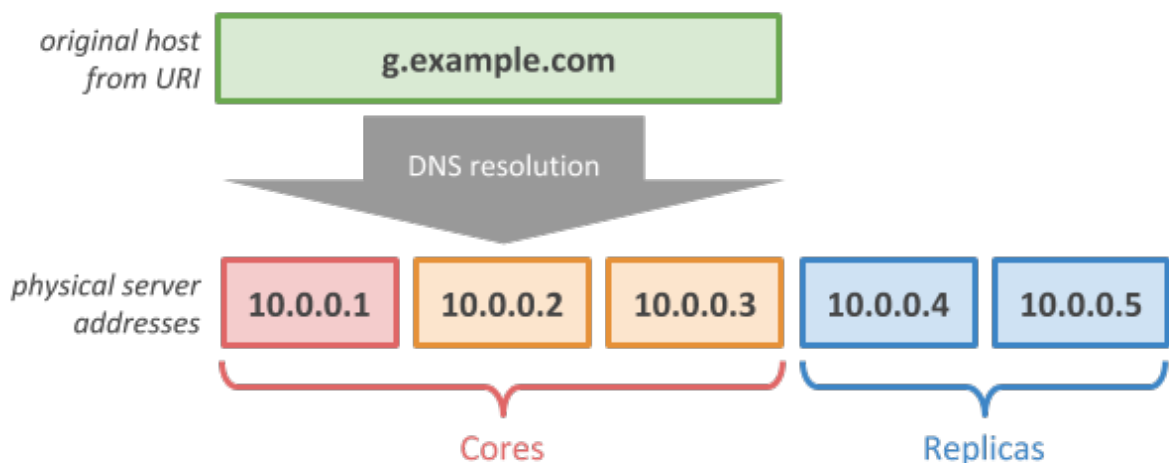

Figure 1. Initial address resolution over DNS

Custom middleware

Middleware, such as a load balancer, can be used to group the core servers under a single public address.
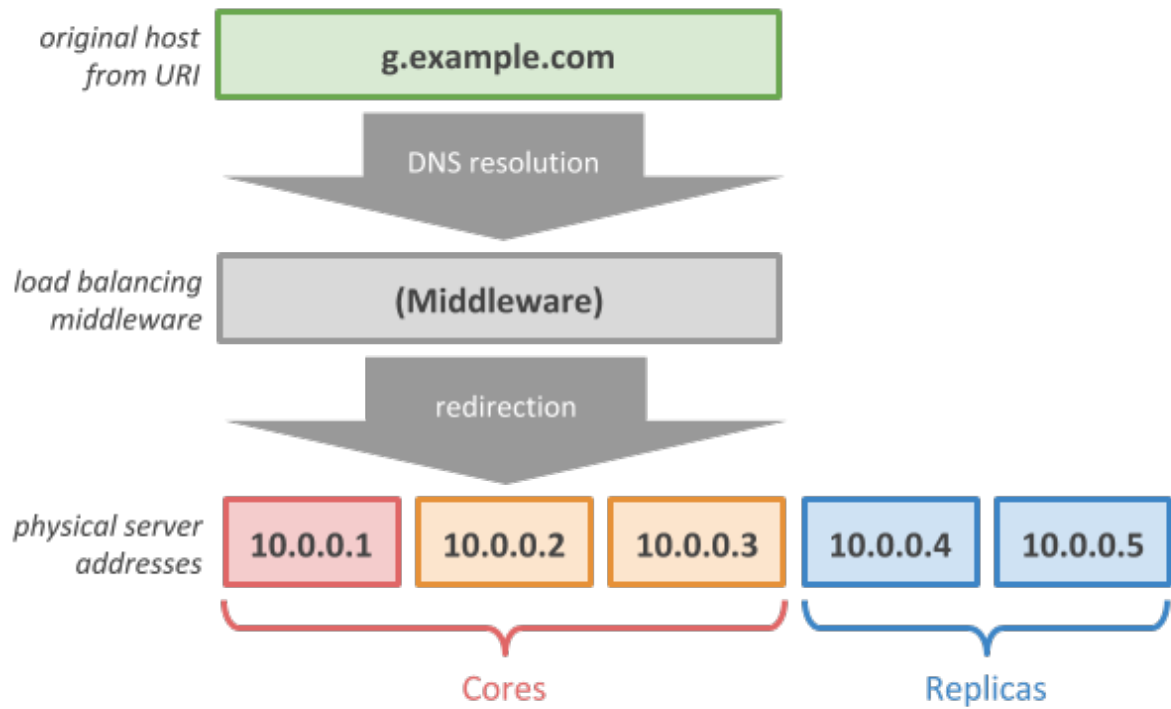
Figure 2. Initial address resolution using custom middleware

Resolver function

Neo4j Drivers also present an address resolution intercept hook called the **resolver function**.

This takes the form of a callback function that accepts a single input address and returns multiple output addresses. The function may hard code the output addresses or may draw them from another configuration source, as required.

Figure 3. Initial address resolution using resolver function

The example below shows how to expand a single address into multiple (hard-coded) output addresses:

Example 11. Custom Address Resolver

```csharp
private IDriver CreateDriverWithCustomResolver(string virtualUri, IAuthToken token,
    params ServerAddress[] addresses)
{
    return GraphDatabase.Driver(virtualUri, token,
        o => o.WithResolver(new ListAddressResolver(addresses)).WithEncryptionLevel
(EncryptionLevel.None));
}

public void AddPerson(string name)
{
    using (var driver = CreateDriverWithCustomResolver("neo4j://x.example.com",
        AuthTokens.Basic(Username, Password),
        ServerAddress.From("a.example.com", 7687), ServerAddress.From("b.example.com", 7877),
        ServerAddress.From("c.example.com", 9092)))
    {
        using (var session = driver.Session())
        {
            session.Run("CREATE (a:Person {name: $name})", new {name});
        }
    }
}

private class ListAddressResolver : IServerAddressResolver
{
    private readonly ServerAddress[] servers;

    public ListAddressResolver(params ServerAddress[] servers)
    {
        this.servers = servers;
    }

    public ISet<ServerAddress> Resolve(ServerAddress address)
    {
        return new HashSet<ServerAddress>(servers);
    }
}
```

Routing table

The routing table acts like the glue between the driver connectivity layer and the database surface. This table contains a list of server addresses, grouped as *readers* and *writers*, and is refreshed automatically by the driver as required.

The driver does not expose any API to work directly with the routing table, but it can sometimes be useful to explore when troubleshooting a system.

Routing context

A routing context can be included as the query part of a `neo4j://` URI.

Routing contexts are defined by means of *server policies* and allow customization of the contents of the routing table.

Example 12. Configure a routing driver with routing context

This example will assume that **Neo4j** has been configured for server policies as described in Neo4j Operations Manual → Load balancing for multi-data center systems. In particular, a server policy called `europe` has been defined. Additionally, we have a server `neo01.graph.example.com` to which we wish to direct the driver.

This URI will use the server policy `europe`:

`neo4j://neo01.graph.example.com?policy=europe`

Server-side configuration to enable routing drivers with routing context

A prerequisite for using a routing driver with routing context is that the **Neo4j** database is operated on a Causal Cluster with the Multi-data center licensing option enabled. Additionally, the routing contexts must be defined within the cluster as *routing policies*.

For details on how to configure *multi-data center routing policies* for a Causal Cluster, please refer to Operations Manual → Causal Clustering.

Exposing a single instance deployment on a remote host using the `neo4j` *URI scheme*

If you are using a *single instance* of **Neo4j**, deployed on a remote machine whilst using the `neo4j` URI scheme, *you will need to complete additional configuration of the server.*

To make the server aware of its deployment environment, you need to configure `default_advertised_address` with your deployment machine's host name, as visible from the client machine.

Examples

Connection URIs are typically formed according to the following pattern:

```
neo4j://<HOST>:<PORT>[?<ROUTING_CONTEXT>]
```

This targets a routed **Neo4j** service that may be fulfilled by either a cluster or a single instance. The `HOST` and `PORT` values contain a logical hostname and port number targeting the entry point to the **Neo4j** service (e.g. `neo4j://graph.example.com:7687`).

In a clustered environment, the URI address will resolve to one of more of the core members; for standalone installations, this will simply point to that server address. The `ROUTING_CONTEXT` option allows for customization of the routing table and is discussed in more detail in Routing context.

An alternative URI form, using the `bolt` URI scheme (e.g. `bolt://graph.example.com:7687`), can be used when a *single point-to-point connection is required*. This variant is useful for the subset client applications (such as admin tooling) that need to be aware of individual servers, as opposed to those which require a highly available database service.

```
bolt://<HOST>:<PORT>
```

Each of the `neo4j` and `bolt` URI schemes permit variants that contain extra encryption and trust information. The `+s` variants enable encryption with a full certificate check, and the `+ssc` variants enable encryption, but with no certificate check. This latter variant is designed specifically for use with self-signed certificates.

Table 3. Available URI schemes

URI scheme	Routing	Description
`neo4j`	Yes	Unsecured
`neo4j+s`	Yes	Secured with full certificate
`neo4j+ssc`	Yes	Secured with self-signed certificate
`bolt`	No	Unsecured
`bolt+s`	No	Secured with full certificate
`bolt+ssc`	No	Secured with self-signed certificate

> ℹ️ **Neo4j 3.x** did not provide a routing table in single instance mode and therefore you should use a `bolt://` URI if targeting an older, non-clustered server.

The table below provides example code snippets for different deployment configurations. Each snippet expects an `auth` variable to have been previously defined, containing the authentication details for that connection.

Example 13. Connecting to a service

The table below illustrates examples of how to connect to a service with routing:

Product	Security	Code snippet	Comments

Neo4j Aura	Secured with full certificate	```GraphDatabase.Driver("neo4j+s://graph.example.com:7687", auth)```
		
		 If you do not have at least the .NET Driver 4.0.1 patch installed, you will need this snippet instead:
		
		 ```String uri = "neo4j://graph.example.com:7687"; IDriver driver = GraphDatabase.Driver(uri, auth, o => o.WithEncryptionLevel(EncryptionLevel.Encrypted));```
			This is the default (and only option) for Neo4j Aura

Neo4j 4.x	Unsecured	```GraphDatabase.Driver("neo4j://graph.example.com:7687", auth);```	This is the default for Neo4j 4.x series

Neo4j 4.x	Secured with full certificate	```GraphDatabase.Driver("neo4j+s://graph.example.com:7687", auth)```
		
		 If you do not have at least the .NET Driver 4.0.1 patch installed, you will need this snippet instead:
		
		 ```String uri = "neo4j://graph.example.com:7687"; IDriver driver = GraphDatabase.Driver(uri, auth, o => o.WithEncryptionLevel(EncryptionLevel.Encrypted));```	

Product	Security	Code snippet	Comments

Neo4j 4.x	Secured with self-signed certificate	```GraphDatabase.Driver("neo4j+ssc://graph.example.com:7687", auth)```
		
		If you do not have at least the .NET Driver 4.0.1 patch installed, you will need this snippet instead:
		
		```String uri = "neo4j://graph.example.com:7687";\nIDriver driver = GraphDatabase.Driver(uri, auth,\n o => o.WithEncryptionLevel(EncryptionLevel.Encrypted)\n\n.WithTrustManager(TrustManager.CreateInsecure()));```	

Neo4j 3.x	Secured with full certificate	```GraphDatabase.Driver("neo4j+s://graph.example.com:7687", auth)```
		
		If you do not have at least the .NET Driver 4.0.1 patch installed, you will need this snippet instead:
		
		```String uri = "neo4j://graph.example.com:7687";\nIDriver driver = GraphDatabase.Driver(uri, auth,\n o => o.WithEncryptionLevel(EncryptionLevel.Encrypted));```	

Product	Security	Code snippet	Comments

Neo4j 3.x	Secured with self-signed certificate	```GraphDatabase.Driver("neo4j+ssc://graph.example.com:7687", auth)```
		
		 If you do not have at least the .NET Driver 4.0.1 patch installed, you will need this snippet instead:
		
		 ```String uri = "neo4j://graph.example.com:7687"; IDriver driver = GraphDatabase.Driver(uri, auth, o => o.WithEncryptionLevel(EncryptionLevel.Encrypted) .WithTrustManager(TrustManager.CreateInsecure()));```
			This is the default for Neo4j 3.x series

Neo4j 3.x	Unsecured	```GraphDatabase.Driver("neo4j://graph.example.com:7687", auth);```	

> **ℹ** To connect to a service without routing, you can replace `neo4j` with `bolt`.

Authentication

Authentication details are provided as an auth token which contains the user names, passwords or other credentials required to access the database. **Neo4j** supports multiple authentication standards but uses *basic authentication* by default.

Basic authentication

The basic authentication scheme is backed by a password file stored within the server and requires applications to provide a user name and password. For this, use the basic auth helper:

Example 14. Basic authentication

```
public IDriver CreateDriverWithBasicAuth(string uri, string user, string password)
{
    return GraphDatabase.Driver(uri, AuthTokens.Basic(user, password));
}
```

> ℹ️ The basic authentication scheme can also be used to authenticate against an LDAP server.

Kerberos authentication

The Kerberos authentication scheme provides a simple way to create a Kerberos authentication token with a base64 encoded server authentication ticket. The best way to create a Kerberos authentication token is shown below:

Example 15. Kerberos authentication

```
public IDriver CreateDriverWithKerberosAuth(string uri, string ticket)
{
    return GraphDatabase.Driver(uri, AuthTokens.Kerberos(ticket),
        o => o.WithEncryptionLevel(EncryptionLevel.None));
}
```

> ℹ️ The Kerberos authentication token can only be understood by the server if the server has the Kerberos Add-on installed.

Custom authentication

For advanced deployments, where a custom security provider has been built, the custom authentication helper can be used.

Example 16. Custom authentication

```
public IDriver CreateDriverWithCustomizedAuth(string uri,
    string principal, string credentials, string realm, string scheme,
    Dictionary<string, object> parameters)
{
    return GraphDatabase.Driver(uri, AuthTokens.Custom(principal, credentials, realm, scheme,
parameters),
        o=>o.WithEncryptionLevel(EncryptionLevel.None));
}
```

Configuration

`ConnectionAcquisitionTimeout`

The maximum amount of time a session will wait when requesting a connection from the connection pool. For connection pools where all connections are currently being used and the `MaxConnectionPoolSize` limit has been reached, a session will wait this duration for a connection to be made available. Since the process of acquiring a connection may involve creating a new connection, *ensure that the value of this configuration is higher than the configured* `ConnectionTimeout`.

Setting a low value will allow for transactions to *fail fast* when all connections in the pool have been acquired by other transactions. Setting a higher value will result in these transactions being queued, increasing the chances of eventually acquiring a connection at the cost of longer time to receive feedback on failure. Finding an optimal value may require an element of experimentation, taking into consideration the expected levels of parallelism within your application as well as the `MaxConnectionPoolSize`.

Default: 60 seconds

Example 17. Configure connection pool

```
public IDriver CreateDriverWithCustomizedConnectionPool(string uri, string user, string
password)
{
    return GraphDatabase.Driver(uri, AuthTokens.Basic(user, password),
        o => o.WithMaxConnectionLifetime(TimeSpan.FromMinutes(30))
            .WithMaxConnectionPoolSize(50)
            .WithConnectionAcquisitionTimeout(TimeSpan.FromMinutes(2)));
}
```

`ConnectionTimeout`

The maximum amount of time to wait for a TCP connection to be established. Connections are only created when a session requires one unless there is an available connection in the connection pool. The driver maintains a pool of open connections which can be loaned to a session when one is available. If a connection is not available, then an attempt to create a new connection (provided the `MaxConnectionPoolSize` limit has not been reached) is made with this configuration option, providing the maximum amount of time to wait for the connection to be established.

In environments with high latency and high occurrences of connection timeouts it is recommended to configure a higher value. For lower latency environments and quicker feedback on potential network issues configure with a lower value.

Default: 30 seconds

Example 18. Configure connection timeout

```
public IDriver CreateDriverWithCustomizedConnectionTimeout(string uri, string user, string
password)
{
    return GraphDatabase.Driver(uri, AuthTokens.Basic(user, password),
        o => o.WithConnectionTimeout(TimeSpan.FromSeconds(15)));
}
```

CustomResolver

Specify a custom server address resolver used by the routing driver to resolve the initial address used to create the driver. See Resolver function for more details.

Encryption

Specify whether to use an encrypted connection between the driver and server.

Default: None

Example 19. Unencrypted configuration

```
public IDriver CreateDriverWithCustomizedSecurityStrategy(string uri, string user, string
password)
{
    return GraphDatabase.Driver(uri, AuthTokens.Basic(user, password),
        o=>o.WithEncryptionLevel(EncryptionLevel.None));
}
```

MaxConnectionLifetime

The maximum duration the driver will keep a connection for before being removed from the pool. Note that while the driver will respect this value, it is possible that the network environment will close connections inside this lifetime. This is beyond the control of the driver. The check on the connection's lifetime happens when a session requires a connection. If the available connection's lifetime is over this limit it is closed and a new connection is created, added to the pool and returned to the requesting session. Changing this configuration value would be useful in environments where users don't have full control over the network environment and wish to proactively ensure all connections are ready.

Setting this option to a low value will cause a high connection churn rate, and can result in a performance drop. It is recommended to pick a value smaller than the maximum lifetime exposed by the surrounding system infrastructure (such as operating system, router, load balancer, proxy and firewall). Negative values result in lifetime not being checked.

Default: 1 hour (3600 seconds)

MaxConnectionPoolSize

The maximum total number of connections allowed, per host (i.e. cluster nodes), to be managed by the connection pool. In other words, for a direct driver using the bolt:// scheme, this sets the maximum

number of connections towards a single database server. For a driver connected to a cluster using the `neo4j://` scheme, this sets the maximum amount of connections per cluster member. If a session or transaction tries to acquire a connection at a time when the pool size is at its full capacity, it must wait until a free connection is available in the pool or the request to acquire a new connection times out. The connection acquiring timeout is configured via `ConnectionAcquisitionTimeout`.

This configuration option allows you to manage the memory and I/O resources being used by the driver and tuning this option is dependent on these factors, in addition to number of cluster members.

Default: 100 connections (except .NET Driver with 500 connections)

MaxTransactionRetryTime

The maximum amount of time that a managed transaction will retry for before failing. Queries that are executed within a managed transaction gain the benefit of being retried when a transient error occurs. When this happens the transaction is retired multiple times up to the `MaxTransactionRetryTime`.

Configure this option higher in high latency environments or if you are executing many large transactions which could limit the number of times that they are retired and therefore their chance to succeed. Configure lower in low latency environments and where your workload mainly consists of many smaller transactions. Failing transactions faster may highlight the reasons behind the transient errors making it easier to fix underlying issues.

Default: 30 seconds

Example 20. Configure maximum retry time

```
public IDriver CreateDriverWithCustomizedMaxRetryTime(string uri, string user, string password)
{
    return GraphDatabase.Driver(uri, AuthTokens.Basic(user, password),
        o => o.WithMaxTransactionRetryTime(TimeSpan.FromSeconds(15)));
}
```

TrustStrategy

Specify how to determine the authenticity of encryption certificates provided by the Neo4j instance that you are connecting to. There are three choices as to which strategy to use:

- `TRUST_SYSTEM_CA_SIGNED_CERTIFICATES` - Accept any certificate that can be verified against system store.

- `TRUST_CUSTOM_CA_SIGNED_CERTIFICATES` - Accept any certificate that can be verified against the a custom CA.

- `TRUST_ALL_CERTIFICATES` - Accept any certificate provided by the server, regardless of CA chain. **We do not recommend using this setting for production environments.**

Default: `TRUST_SYSTEM_CA_SIGNED_CERTIFICATES` (Note - only when encryption is enabled)

Example 21. Configure trusted certificates

```
public IDriver CreateDriverWithCustomizedTrustStrategy(string uri, string user, string password)
{
    return GraphDatabase.Driver(uri, AuthTokens.Basic(user, password),
        o => o.WithTrustManager(TrustManager.CreateInsecure()));
}
```

KeepAlive

Specify whether TCP keep-alive should be enabled. To ensure that the connection between the driver and server is still operational, the TCP layer can periodically send messages to check the connection.

Default: True

> ℹ️ This is currently only configurable in the .NET Driver

Logging

All official **Neo4j Drivers** log information to standard logging channels. This can typically be accessed in an ecosystem-specific way.

The code snippets below show how to redirect log messages to standard output.

| .NET | ```
#Please note that you will have to provide your own console logger implementing the ILogger interface.
IDriver driver = GraphDatabase.Driver(..., o => o.WithLogger(logger));
``` |
|---|---|
| Go | ```
driver, err := neo4j.NewDriver(..., func(c *neo4j.Config) { c.Log = neo4j.ConsoleLogger(neo4j.DEBUG) })
``` |
| Java | ```
ConfigBuilder.withLogging(Logging.console(Level.DEBUG))
``` |
| JavaScript | ```
const loggingConfig = {logging: neo4j.logging.console('debug')};
const driver = neo4j.driver(..., loggingConfig);
``` |
| Python | ```
from logging import getLogger, StreamHandler, DEBUG
handler = StreamHandler()
handler.setLevel(DEBUG)
getLogger("neo4j").addHandler(handler)
``` |

Cypher workflow

Overview

The Neo4j Drivers expose a Cypher Channel over which database work can be carried out (see the Cypher Manual for more information on the Cypher Query Language).

Work itself is organized into **sessions**, **transactions** and **queries**, as follows:

Figure 4. Sessions, queries and transactions

Sessions are always bound to **a single transaction context**, which is typically an individual database.

Using the bookmark mechanism, sessions also provide a guarantee of correct transaction sequencing, even when transactions occur over multiple cluster members. This effect is known as causal chaining.

Sessions

Sessions are **lightweight containers for causally chained sequences of transactions** (see Operations Manual → Causal consistency). They essentially provide context for storing transaction sequencing information in the form of bookmarks.

When a transaction begins, the session in which it is contained acquires a connection from the driver connection pool. On commit (or rollback) of the transaction, the session releases that connection again. This means that it is only when a session is carrying out work that it occupies a connection resource. When idle, no such resource is in use.

Due to the sequencing guaranteed by a session, sessions may only host **one transaction at a time**. For parallel execution, multiple sessions should be used. In languages where thread safety is an issue, **sessions should not be considered thread-safe**.

Closing a session forces any open transaction to be rolled back, and its associated connection to consequently be released back into the pool.

Sessions are bound to a single transactional context, specified on construction. **Neo4j** exposes each database inside its own context, thereby prohibiting cross-database transactions (or sessions) by design.

Similarly, *sessions bound to different databases may not be causally chained by propagating bookmarks between them*.

Individual language drivers provide several session classes, each oriented around a particular programming style. Each session class provides a similar set of functionality but offers client applications a choice based on how the application is structured and what frameworks are in use, if any.

The session classes are described in The session API. For more details, please see the API documentation for your language.

Transactions

Transactions are **atomic units of work** containing one or more **Cypher Queries**. Transactions may contain read or write work, and will generally be routed to an appropriate server for execution, where they will be carried out in their entirety. In case of a transaction failure, the transaction needs to be retried from the beginning. This is the responsibility of the transaction manager.

The Neo4j Drivers provide **transaction management** via the transaction function mechanism. This mechanism is exposed through methods on the Session object which accept a function object that can be played multiple times against different servers until it either succeeds or a timeout is reached. This approach is recommended for most client applications.

A convenient **short-form alternative** is the auto-commit transaction mechanism. This provides a limited form of transaction management for single-query transactions, as a trade-off for a slightly smaller code overhead. This form of transaction is useful for **quick scripts and environments where high availability guarantees are not required**. It is also the required form of transaction for running `PERIODIC COMMIT` queries, which are the only type of Cypher Query to manage their own transactions.

A lower-level **unmanaged transaction API** is also available for advanced use cases. This is useful when an alternative transaction management layer is applied by the client, in which error handling and retries need to be managed in a custom way. To learn more about how to use unmanaged transactions, see API documentation for the relevant language.

Queries and results

Queries consist of a request to the server to execute a Cypher statement, followed by a response back to the client with the result. Results are transmitted as a **stream of records**, along with header and footer metadata, and can be incrementally consumed by a client application. With reactive capabilities, the semantics of the record stream can be enhanced by allowing a Cypher result to be paused or cancelled part-way through.

To execute a Cypher Query, the **query text** is required along with an optional set of **named parameters**. The text can contain **parameter placeholders** that are substituted with the corresponding values at runtime. While it is possible to run non-parameterized Cypher Queries, **good programming practice is to use parameters in Cypher Queries whenever possible**. This allows for the caching of queries within the Cypher Engine, which is beneficial for performance. Parameter values should adhere to /docs/cypher-manual/4.1/syntax/values.

A result summary is also generally available. This contains additional information relating to the query

execution and the result content. For an `EXPLAIN` or `PROFILE` query, this is where the query plan is returned. See Cypher Manual → Profiling a query for more information on these queries.

Causal chaining and bookmarks

When working with a *Causal Cluster*, transactions can be chained, via a session, to ensure *causal consistency*. This means that for any two transactions, it is guaranteed that the second transaction will begin only after the first has been successfully committed. This is true even if the transactions are carried out on different physical cluster members. For more information on Causal Clusters, please refer to Operations Manual → Clustering.

Internally, causal chaining is carried out by passing bookmarks between transactions. Each bookmark records one or more points in transactional history for a particular database, and can be used to inform cluster members to carry out units of work in a particular sequence. On receipt of a bookmark, the server will block until it has caught up with the relevant transactional point in time.

An initial bookmark is sent from client to server on beginning a new transaction, and a final bookmark is returned on successful completion. Note that this applies to both read and write transactions.

Bookmark propagation is carried out automatically within sessions and does not require any explicit signal or setting from the application. To opt out of this mechanism, for unrelated units of work, applications can use multiple sessions. This avoids the small latency overhead of the causal chain.

Bookmarks can be passed between sessions by extracting the last bookmark from a session and passing this into the construction of another. Multiple bookmarks can also be combined if a transaction has more than one logical predecessor. Note that it is only when chaining across sessions that an application will need to work with bookmarks directly.

Figure 5. Passing bookmarks

Example 22. Pass bookmarks

Pass bookmarks

```csharp
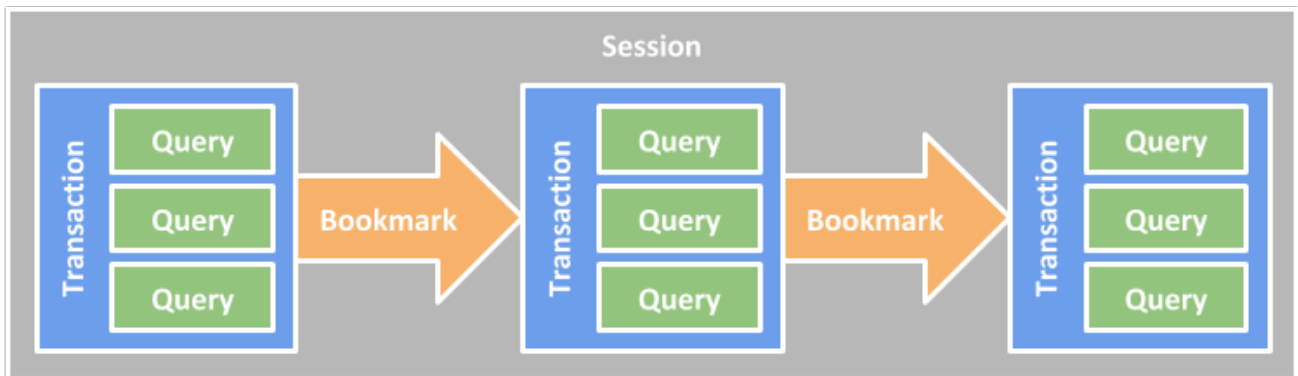// Create a company node
private IResult AddCompany(ITransaction tx, string name)
{
    return tx.Run("CREATE (a:Company {name: $name})", new {name});
}

// Create a person node
private IResult AddPerson(ITransaction tx, string name)
{
    return tx.Run("CREATE (a:Person {name: $name})", new {name});
}

// Create an employment relationship to a pre-existing company node.
// This relies on the person first having been created.
private IResult Employ(ITransaction tx, string personName, string companyName)
{
    return tx.Run(@"MATCH (person:Person {name: $personName})
            MATCH (company:Company {name: $companyName})
            CREATE (person)-[:WORKS_FOR]->(company)", new {personName, companyName});
}

// Create a friendship between two people.
private IResult MakeFriends(ITransaction tx, string name1, string name2)
{
    return tx.Run(@"MATCH (a:Person {name: $name1})
            MATCH (b:Person {name: $name2})
            MERGE (a)-[:KNOWS]->(b)", new {name1, name2});
}

// Match and display all friendships.
private int PrintFriendships(ITransaction tx)
{
    var result = tx.Run("MATCH (a)-[:KNOWS]->(b) RETURN a.name, b.name");

    var count = 0;
    foreach (var record in result)
    {
        count++;
        Console.WriteLine($"{record["a.name"]} knows {record["b.name"]}");
    }

    return count;
}

public void AddEmployAndMakeFriends()
{
    // To collect the session bookmarks
    var savedBookmarks = new List<Bookmark>();

    // Create the first person and employment relationship.
    using (var session1 = Driver.Session(o => o.WithDefaultAccessMode(AccessMode.Write)))
    {
        session1.WriteTransaction(tx => AddCompany(tx, "Wayne Enterprises"));
        session1.WriteTransaction(tx => AddPerson(tx, "Alice"));
        session1.WriteTransaction(tx => Employ(tx, "Alice", "Wayne Enterprises"));

        savedBookmarks.Add(session1.LastBookmark);
    }

    // Create the second person and employment relationship.
    using (var session2 = Driver.Session(o => o.WithDefaultAccessMode(AccessMode.Write)))
    {
        session2.WriteTransaction(tx => AddCompany(tx, "LexCorp"));
        session2.WriteTransaction(tx => AddPerson(tx, "Bob"));
        session2.WriteTransaction(tx => Employ(tx, "Bob", "LexCorp"));

        savedBookmarks.Add(session2.LastBookmark);
    }

    // Create a friendship between the two people created above.
    using (var session3 = Driver.Session(o =>
        o.WithDefaultAccessMode(AccessMode.Write).WithBookmarks(savedBookmarks.ToArray())))
    {
```

```
            session3.WriteTransaction(tx => MakeFriends(tx, "Alice", "Bob"));

            session3.ReadTransaction(PrintFriendships);
        }
    }
```

Routing transactions using access modes

Transactions can be executed in either **read** or **write** mode; this is known as the access mode. In a Causal Cluster, each transaction will be routed to an appropriate server based on the mode. When using a single instance, all transactions will be passed to that one server.

Routing Cypher by identifying reads and writes can improve the utilization of available cluster resources. Since read servers are typically more plentiful than write servers, it is beneficial to direct read traffic to read servers instead of the write server. Doing so helps in keeping write servers available for write transactions.

Access mode is generally specified by the method used to call the transaction function. Session classes provide a method for calling reads and another for writes.

As a fallback for **auto-commit** and **unmanaged transactions**, a *default access mode* can also be provided *at session level*. This is only used in cases when the access mode cannot otherwise be specified. *In case a transaction function is used within that session, the default access mode will be overridden.*

> ℹ The driver does not parse **Cypher** and therefore cannot automatically determine whether a transaction is intended to carry out read or write operations. As a result, a write transaction tagged as a read will still be sent to a read server, but will fail on execution.

Example 23. Read-write transaction

Read write transaction

```
public long AddPerson(string name)
{
    using (var session = Driver.Session())
    {
        session.WriteTransaction(tx => CreatePersonNode(tx, name));
        return session.ReadTransaction(tx => MatchPersonNode(tx, name));
    }
}

private static IResult CreatePersonNode(ITransaction tx, string name)
{
    return tx.Run("CREATE (a:Person {name: $name})", new {name});
}

private static long MatchPersonNode(ITransaction tx, string name)
{
    var result = tx.Run("MATCH (a:Person {name: $name}) RETURN id(a)", new {name});
    return result.Single()[0].As<long>();
}
```

Databases and execution context

Neo4j offers the ability to work with *multiple databases* within the same DBMS.

For *Community Edition*, this is limited to **one user database**, plus the `system` database.

From a driver API perspective, sessions have a DBMS scope, and the default database for a session can be selected on session construction. The default database is used as a target for queries that don't explicitly specify a database with a `USE` clause. See Cypher Manual → USE for details about the `USE` clause.

In a multi-database environment, **the server tags one database as default**. This is selected whenever a session is created without naming a particular database as default. In an environment with a single database, that database is always the default.

For more information about managing multiple databases within the same DBMS, refer to Cypher Manual → Neo4j databases and graphs, which has a full breakdown of the **Neo4j** data storage hierarchy.

The following example illustrates how to work with databases:

```
var session = driver.session(SessionConfig.forDatabase( "foo" ))
// lists nodes from database foo
session.run("MATCH (n) RETURN n").list()

// lists nodes from database bar
session.run("USE bar MATCH (n) RETURN n").list()

// creates an index in foo
session.run("CREATE INDEX foo_idx FOR (n:Person) ON n.id").consume()

// creates an index in bar
session.run("USE bar CREATE INDEX bar_idx FOR (n:Person) ON n.id").consume()

// targets System database
session.run("SHOW USERS")
```

Database selection

The database selection on the client side happens on the session API.

You pass the name of the database to the driver during session creation. If you don't specify a name, the default database is used. The database name must not be `null`, nor an empty string.

> ⓘ *The selection of database is only possible when the driver is connected against Neo4j Enterprise Edition.* Changing to any other database than the default database in Neo4j Community Edition will lead to a runtime error.

The following example illustrates the concept of DBMS Cypher Manual → Transactions and shows how queries to multiple databases can be issued in one driver transaction. It is annotated with comments describing how each operation impacts database transactions.

```
var session = driver.session(SessionConfig.forDatabase( "foo" ))
// a DBMS-level transaction is started
var transaction = session.begin()

// a transaction on database "foo" is started automatically with the first query targeting foo
transaction.run("MATCH (n) RETURN n").list()

// a transaction on database "bar" is started
transaction.run("USE bar MATCH (n) RETURN n").list()

// executes in the transaction on database "foo"
transaction.run("RETURN 1").consume()

// executes in the transaction on database "bar"
transaction.run("USE bar RETURN 1").consume()

// commits the DBMS-level transaction which commits the transactions on databases "foo" and
// "bar"
transaction.commit()
```

Please note that the database that is requested must exist.

Example 24. Database selection on session creation

Database selection on session creation

```
using (var session = _driver.Session(SessionConfigBuilder.ForDatabase("examples")))
{
    session.Run("CREATE (a:Greeting {message: 'Hello, Example-Database'}) RETURN a").Consume();
}

void SessionConfig(SessionConfigBuilder configBuilder) =>
    configBuilder.WithDatabase("examples")
        .WithDefaultAccessMode(AccessMode.Read)
        .Build();

using (var session = _driver.Session(SessionConfig))
{
    var result = session.Run("MATCH (a:Greeting) RETURN a.message as msg");
    var msg = result.Single()[0].As<string>();
    Console.WriteLine(msg);
}
```

Type mapping

Drivers translate between *application language types* and the *Cypher Types.*

To pass parameters and process results, it is important to know the basics of how Cypher works with types and to understand how the Cypher Types are mapped in the driver.

The table below shows the available data types. All types can be potentially found in the result, although not all types can be used as parameters.

Cypher Type	Parameter	Result
null*		
List		
Map		
Boolean		
Integer		
Float		
String		
ByteArray		
Date		
Time		
LocalTime		
DateTime		
LocalDateTime		
Duration		
Point		
Node**		
Relationship**		
Path**		

* The null marker is not a type but a placeholder for absence of value. For information on how to work with null in Cypher, please refer to Cypher Manual → Working with null.

** Nodes, relationships and paths are passed in results as snapshots of the original graph entities. While the original entity IDs are included in these snapshots, no permanent link is retained back to the underlying server-side entities, which may be deleted or otherwise altered independently of the client copies. Graph structures may not be used as parameters because it depends on application context whether such a parameter would be passed by reference or by value, and Cypher provides no mechanism to denote this. Equivalent functionality is available by simply passing either the ID for pass-by-reference, or an extracted map of properties for pass-by-value.

The Neo4j Drivers map Cypher Types to and from native language types as depicted in the table below. Custom types (those not available in the language or standard library) are highlighted in **bold**.

Example 25. Map Neo4j types to native language types

Neo4j Cypher Type	.NET Type
`null`	`null`
`List`	`IList<object>`
`Map`	`IDictionary<string, object>`
`Boolean`	`bool`
`Integer`	`long`
`Float`	`double`
`String`	`string`
`ByteArray`	`byte[]`
`Date`	**`LocalDate`**
`Time`	**`OffsetTime`**
`LocalTime`	**`LocalTime`**
`DateTime`*	**`ZonedDateTime`**
`LocalDateTime`	**`LocalDateTime`**
`Duration`	**`Duration`**
`Point`	**`Point`**
`Node`	**`INode`**
`Relationship`	**`IRelationship`**
`Path`	**`IPath`**

* Time zone names adhere to the IANA system, rather than the Windows system. Inbound conversion is carried out using Extended Windows-Olson zid mapping as defined by Unicode CLDR.

Exceptions and error handling

When executing **Cypher** or carrying out other operations with the driver, certain exceptions and error cases may arise. **Server-generated exceptions** are each associated with a status code that describes the nature of the problem and a message that provides more detail.

The classifications are listed in the table below.

Table 4. Server status code classifications

Classification	Description
ClientError	The client application has caused an error. The application should amend and retry the operation.

Classification	Description
DatabaseError	The server has caused an error. Retrying the operation will generally be unsuccessful.
TransientError	A temporary error has occurred. The application should retry the operation.

Service unavailable

A *Service Unavailable* exception will be signalled when the driver is no longer able to establish communication with the server, even after retries.

Encountering this condition usually indicates a fundamental networking or database problem.

While certain mitigations can be made by the driver to avoid this issue, there will always be cases when this is impossible. As such, it is highly recommended to ensure client applications contain a code path that can be followed when the client is no longer able to communicate with the server.

Transient errors

Transient errors are those which are generated by the server and marked as safe to retry without alteration to the original request. Examples of such errors are deadlocks and memory issues.

When using transaction functions, the driver will usually be able to automatically retry when a transient failure occurs.

The session API

This chapter describes the following:

Simple sessions

i	Available in: - .NET Driver - Go Driver - Java Driver - Python Driver

Simple sessions provide a "classic" *blocking style* API for Cypher execution. In general, simple sessions provide the easiest programming style to work with since API calls are executed in a strictly sequential fashion.

Lifecycle

The session lifetime extends from session construction to session closure. In languages that support them, simple sessions are usually scoped within a context block; this ensures that they are properly closed and

that any underlying connections are released and not leaked.

Example 26. Session construction and closure

```
using (var session = driver.Session(...)) {
    // transactions go here
}
```

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Transaction functions

Transaction functions are used for containing transactional units of work. This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic.

Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Transaction functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are operating in a single instance environment, this routing has no impact. It does give you flexibility if you choose to adopt a clustered environment later on.

Before writing a transaction function *it is important to ensure that it is designed to be idempotent*. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

	Transaction functions are the recommended form for containing transactional units of work. *When a transaction fails, the driver retry logic is invoked*. For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors.

Example 27. Transaction function

```csharp
public void AddPerson(string name)
{
    using (var session = Driver.Session())
    {
        session.WriteTransaction(tx => tx.Run("CREATE (a:Person {name: $name})", new {name}));
    }
}
```

Auto-commit transactions

An **auto-commit transaction** is a **basic but limited form of transaction**. Such a transaction consists of only one Cypher query and is not automatically replayed on failure. Therefore any error scenarios must be handled by the client application itself.

Auto-commit transactions are **intended to be used for simple use cases** such as when learning **Cypher** or writing one-off scripts.

⚠️	It is not recommended to use auto-commit transactions in production environments.

ℹ️	Unlike other kinds of Cypher Query, `PERIODIC COMMIT` queries do not participate in the causal chain. Therefore, the only way to execute `PERIODIC COMMIT` from a driver is to use auto-commit transactions. Please refer to the Cypher Manual → `PERIODIC COMMIT` query hint.

Example 28. Simple auto-commit transactions

```csharp
public void AddPerson(string name)
{
    using (var session = Driver.Session())
    {
        session.Run("CREATE (a:Person {name: $name})", new {name});
    }
}
```

Consuming results

Query results are typically consumed as a stream of records. The drivers provide a language-idiomatic way to iterate through that stream.

Example 29. Consuming results

```csharp
public List<string> GetPeople()
{
    using (var session = Driver.Session())
    {
        return session.ReadTransaction(tx =>
        {
            var result = tx.Run("MATCH (a:Person) RETURN a.name ORDER BY a.name");
            return result.Select(record => record[0].As<string>()).ToList();
        });
    }
}
```

Retaining results

Within a session, **only one result stream can be active at any one time**. Therefore, if the result of one query is not fully consumed before another query is executed, the remainder of the first result will be automatically buffered within the result object.
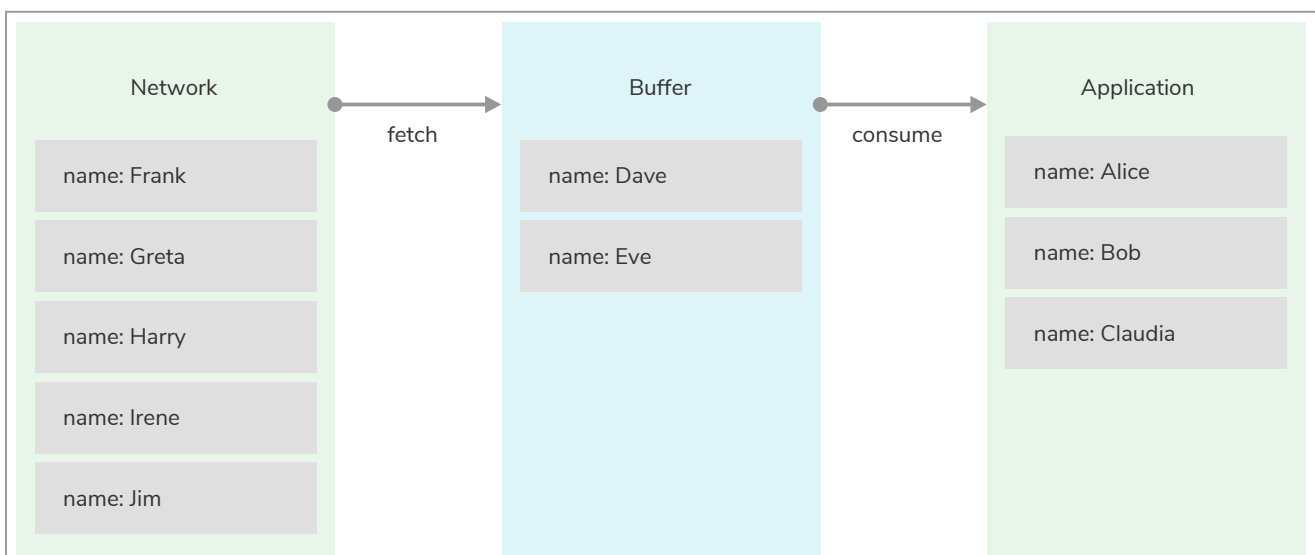
Figure 6. Result buffer

This buffer provides a staging point for results, and divides result handling into *fetching* (moving from the network to the buffer) and *consuming* (moving from the buffer to the application).

> For large results, the result buffer may require a significant amount of memory.
>
> For this reason, **it is recommended to consume results in order wherever possible**.

Client applications can choose to take control of more advanced query patterns by explicitly retaining results. Such explicit retention may also be useful when a result needs to be saved for future processing. The drivers offer support for this process, as per the example below:

Example 30. Retain results for further processing

```csharp
public int AddEmployees(string companyName)
{
    using (var session = Driver.Session())
    {
        var persons =
            session.ReadTransaction(tx => tx.Run("MATCH (a:Person) RETURN a.name AS name"
).ToList());
        return persons.Sum(person => session.WriteTransaction(tx =>
        {
            tx.Run("MATCH (emp:Person {name: $person_name}) " +
                "MERGE (com:Company {name: $company_name}) " +
                "MERGE (emp)-[:WORKS_FOR]->(com)",
                new {person_name = person["name"].As<string>(), company_name = companyName});
            return 1;
        }));
    }
}
```

Asynchronous sessions

Available in:

- .NET Driver
- Java Driver
- JavaScript Driver

Asynchronous sessions provide an API wherein function calls typically return available objects such as futures. This allows client applications to work within asynchronous frameworks and take advantage of cooperative multitasking.

Lifecycle

Session lifetime begins with session construction. A session then exists until it is closed, which is typically set to occur after its contained query results have been consumed.

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Transaction functions

Transaction functions are **the recommended form for containing transactional units of work**. This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic. Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are in a single instance environment, this routing has no impact but it does give you the flexibility should you choose to later adopt a clustered environment.

Before writing a transaction function it is important to ensure that any side-effects carried out by a transaction function should be designed to be idempotent. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

> When a transaction fails, the driver retry logic is invoked. For several failure cases, the transaction can be immediately retried against a different server.
>
> These cases include connection issues, server role changes (e.g. leadership elections) and transient errors. Retry logic can be configured when creating a session.

Example 31. Asynchronous transaction functions

Transaction function

```csharp
public async Task<List<string>> PrintAllProducts()
{
    List<string> result = null;
    var session = Driver.AsyncSession();

    try
    {
        // Wrap whole operation into an managed transaction and
        // get the results back.
        result = await session.ReadTransactionAsync(async tx =>
        {
            var products = new List<string>();

            // Send cypher query to the database
            var reader = await tx.RunAsync(
                "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
                new {id = 0} // Parameters in the query, if any
            );

            // Loop through the records asynchronously
            while (await reader.FetchAsync())
            {
                // Each current read in buffer can be reached via Current
                products.Add(reader.Current[0].ToString());
            }

            return products;
        });
    }
    finally
    {
        // asynchronously close session
        await session.CloseAsync();
    }

    return result;
}
```

Auto-commit transactions

An auto-commit transaction is a basic but limited form of transaction. Such a transaction consists of only

one Cypher query and is not automatically replayed on failure. Therefore any error scenarios must be handled by the client application itself.

Auto-commit transactions are intended to be used for simple use cases such as when learning Cypher or writing one-off scripts.

> ⚠️ It is not recommended to use auto-commit transactions in production environments.

> ℹ️ Unlike other kinds of Cypher Query, `PERIODIC COMMIT` queries do not participate in the causal chain.
>
> Therefore, the only way to execute them from a driver is to use auto-commit transactions.
>
> Please refer to the Cypher Manual → `PERIODIC COMMIT` query hint.

Example 32. Asynchronous auto-commit transactions

Auto-commit transaction

```
public async Task<List<string>> ReadProductTitles()
{
    var records = new List<string>();
    var session = Driver.AsyncSession();

    try
    {
        // Send cypher query to the database.
        // The existing IResult interface implements IEnumerable
        // and does not play well with asynchronous use cases. The replacement
        // IResultCursor interface is returned from the RunAsync
        // family of methods instead and provides async capable methods.
        var reader = await session.RunAsync(
            "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
            new {id = 0} // Parameters in the query, if any
        );

        // Loop through the records asynchronously
        while (await reader.FetchAsync())
        {
            // Each current read in buffer can be reached via Current
            records.Add(reader.Current[0].ToString());
        }
    }
    finally
    {
        // asynchronously close session
        await session.CloseAsync();
    }

    return records;
}
```

Consuming results

The asynchronous session API provides language-idiomatic methods to aid integration with asynchronous applications and frameworks.

Example 33. Asynchronous consuming results

Consuming results

```csharp
public async Task<List<string>> GetPeopleAsync()
{
    var session = Driver.AsyncSession();
    try
    {
        return await session.ReadTransactionAsync(async tx =>
        {
            var result = await tx.RunAsync("MATCH (a:Person) RETURN a.name ORDER BY a.name");

            return await result.ToListAsync(r => r[0].As<string>());
        });
    }
    finally
    {
        await session.CloseAsync();
    }
}
```

Reactive sessions

Available in:

- .NET Driver

- Java Driver

- JavaScript Driver

Starting with **Neo4j 4.0**, the reactive processing of queries is supported. This can be achieved through reactive sessions. Reactive sessions allow for dynamic management of the data that is being exchanged between the driver and the server.

Typical of reactive programming, consumers control the rate at which they consume records from queries and the driver in turn manages the rate at which records are requested from the server. Flow control is supported throughout the entire Neo4j stack, meaning that the query engine responds correctly to the flow control signals. This results in far more efficient resource handling and ensures that the receiving side is not forced to buffer arbitrary amounts of data.

For more information about reactive stream, please see the following:

- The Reactive Manifesto

- Reactive Streams for JVM and JavaScript

- Project Reactor reference documentation

- Introduction to Reactive Extensions to .NET

Lifecycle

Session lifetime begins with session construction. A session then exists until it is closed, which is typically set to occur after its contained query results have been consumed.

Transaction functions

This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic. Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are in a single instance environment, this routing has no impact but it does give you the flexibility should you choose to later adopt a clustered environment.

Before writing a transaction function it is important to ensure that any side-effects carried out by a transaction function should be designed to be idempotent. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

When a transaction fails, the driver retry logic is invoked. For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors. Retry logic can be configured when creating a session.

Example 34. Reactive transaction functions

Reactive transaction function

```
public IObservable<string> PrintAllProducts()
{
    var session = Driver.RxSession();

    return session.ReadTransaction(tx =>
    {
        return tx.Run(
                "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
                new {id = 0} // Parameters in the query, if any
            )
            .Records()
            .Select(record => record[0].ToString());
    }).OnErrorResumeNext(session.Close<string>());
}
```

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Auto-commit transactions

An auto-commit transaction is a basic but limited form of transaction. Such a transaction consists of only one Cypher Query and is not automatically replayed on failure. Therefore any error scenarios will need to be handled by the client application itself.

Auto-commit transactions are intended to be used for simple use cases such as when learning **Cypher** or writing one-off scripts.

⚠	It is not recommended to use auto-commit transactions in production environments.

ℹ	The only way to execute PERIODIC COMMIT Cypher Queries is to auto-commit the transaction.
	Unlike other kinds of Cypher query, PERIODIC COMMIT queries do not participate in the causal chain.
	Please refer to the Cypher Manual → PERIODIC COMMIT query hint.

Example 35. Reactive auto-commit transactions

> **Auto-commit transactions**
>
> ```csharp
> public IObservable<string> ReadProductTitles()
> {
>     var session = Driver.RxSession();
>
>     return session.Run(
>             "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
>             new {id = 0} // Parameters in the query, if any
>         )
>         .Records()
>         .Select(record => record[0].ToString())
>         .OnErrorResumeNext(session.Close<string>());
> }
> ```

Consuming results

To consume data from a query in a reactive session, a subscriber is required to handle the results that are being returned by the publisher.

Each transaction corresponds to a data flow which supplies the data from the server. Result processing begins when records are pulled from this flow. Only one subscriber may pull data from a given flow.

Example 36. Reactive consuming results

> **Consuming results**
>
> ```csharp
> public IObservable<string> GetPeople()
> {
>     var session = Driver.RxSession();
>     return session.ReadTransaction(tx =>
>     {
>         return tx.Run("MATCH (a:Person) RETURN a.name ORDER BY a.name")
>             .Records()
>             .Select(record => record[0].As<string>());
>     }).OnErrorResumeNext(session.Close<string>());
> }
> ```

Cancellation

As per the reactive stream specification, a reactive data flow can be cancelled part way through. This prematurely commits or rolls back the transaction and stops the query engine from producing any more records.

Session configuration

Bookmarks

    The mechanism which ensures causal consistency between transactions within a session. Bookmarks

are implicitly passed between transactions within a single session to meet the causal consistency requirements. There may be scenarios where you might want to use the bookmark from one session in a different new session.

Default: None (Sessions will initially be created without a bookmark)

Default Access Mode

A fallback for the access mode setting when transaction functions are not used. Typically, access mode is set per transaction by calling the appropriate transaction function method. In other cases, this setting is inherited. Note that transaction functions will ignore/override this setting.

Default: Write

Database

The database with which the session will interact. When you are working with a database which is not the default (i.e. the `system` database or another database in Neo4j 4.0 Enterprise Edition), you can explicitly configure the database which the driver is executing transactions against. See Operations Manual → The default database for more information on databases.

Default: the default database as configured on the server.

`Fetch Size`

The number of records to fetch in each batch from the server. Neo4j 4.0 introduces the ability to pull records in batches, allowing the client application to take control of data population and apply back pressure to the server. This `FetchSize` applies to simple sessions and async-sessions whereas reactive sessions can be controlled directly using the request method of the subscription.

Default: 1000 records

The session API

This chapter describes the following:

- Simple sessions
  - Lifecycle
  - Transaction functions
  - Auto-commit transactions
  - Consuming results
  - Retaining results
- Asynchronous sessions
  - Transaction functions
  - Auto-commit transaction
  - Consuming results
- Reactive sessions

Simple sessions

ℹ	Available in: • .NET Driver • Go Driver • Java Driver • Python Driver

Simple sessions provide a "classic" *blocking style* API for Cypher execution. In general, simple sessions provide the easiest programming style to work with since API calls are executed in a strictly sequential fashion.

Lifecycle

The session lifetime extends from session construction to session closure. In languages that support them, simple sessions are usually scoped within a context block; this ensures that they are properly closed and that any underlying connections are released and not leaked.

Example 37. Session construction and closure

```
using (var session = driver.Session(...)) {
    // transactions go here
}
```

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See [Session configuration](#) for more details.

Transaction functions

Transaction functions are used for containing transactional units of work. This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic.

Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Transaction functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are operating in a single instance environment, this routing has no impact. It does give you flexibility if you choose to adopt a clustered environment later on.

Before writing a transaction function *it is important to ensure that it is designed to be idempotent*. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

> *Transaction functions are the recommended form for containing transactional units of work.*
>
> *When a transaction fails, the driver retry logic is invoked*. For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors.

Example 38. Transaction function

```
public void AddPerson(string name)
{
    using (var session = Driver.Session())
    {
        session.WriteTransaction(tx => tx.Run("CREATE (a:Person {name: $name})", new {name}));
    }
}
```

Auto-commit transactions

An *auto-commit transaction* is a *basic but limited form of transaction*. Such a transaction consists of only one Cypher query and is not automatically replayed on failure. Therefore any error scenarios must be handled by the client application itself.

Auto-commit transactions are *intended to be used for simple use cases* such as when learning **Cypher** or writing one-off scripts.

> ⚠ It is not recommended to use auto-commit transactions in production environments.

> Unlike other kinds of Cypher Query, `PERIODIC COMMIT` queries do not participate in the causal chain.
>
> Therefore, the only way to execute `PERIODIC COMMIT` from a driver is to use auto-commit transactions.
>
> Please refer to the Cypher Manual → `PERIODIC COMMIT` query hint.

Example 39. Simple auto-commit transactions

```csharp
public void AddPerson(string name)
{
    using (var session = Driver.Session())
    {
        session.Run("CREATE (a:Person {name: $name})", new {name});
    }
}
```

Consuming results

Query results are typically consumed as a stream of records. The drivers provide a language-idiomatic way to iterate through that stream.

Example 40. Consuming results

```csharp
public List<string> GetPeople()
{
    using (var session = Driver.Session())
    {
        return session.ReadTransaction(tx =>
        {
            var result = tx.Run("MATCH (a:Person) RETURN a.name ORDER BY a.name");
            return result.Select(record => record[0].As<string>()).ToList();
        });
    }
}
```

Retaining results

Within a session, *only one result stream can be active at any one time*. Therefore, if the result of one query is not fully consumed before another query is executed, the remainder of the first result will be automatically buffered within the result object.

Figure 7. Result buffer

This buffer provides a staging point for results, and divides result handling into *fetching* (moving from the network to the buffer) and *consuming* (moving from the buffer to the application).

> ℹ️ For large results, the result buffer may require a significant amount of memory.
>
> For this reason, *it is recommended to consume results in order wherever possible*.

Client applications can choose to take control of more advanced query patterns by explicitly retaining results. Such explicit retention may also be useful when a result needs to be saved for future processing. The drivers offer support for this process, as per the example below:

Example 41. Retain results for further processing

```csharp
public int AddEmployees(string companyName)
{
    using (var session = Driver.Session())
    {
        var persons =
            session.ReadTransaction(tx => tx.Run("MATCH (a:Person) RETURN a.name AS name"
).ToList());
        return persons.Sum(person => session.WriteTransaction(tx =>
        {
            tx.Run("MATCH (emp:Person {name: $person_name}) " +
                    "MERGE (com:Company {name: $company_name}) " +
                    "MERGE (emp)-[:WORKS_FOR]->(com)",
                new {person_name = person["name"].As<string>(), company_name = companyName});
            return 1;
        }));
    }
}
```

Asynchronous sessions

	Available in:
i	• .NET Driver • Java Driver • JavaScript Driver

Asynchronous sessions provide an API wherein function calls typically return available objects such as futures. This allows client applications to work within asynchronous frameworks and take advantage of cooperative multitasking.

Lifecycle

Session lifetime begins with session construction. A session then exists until it is closed, which is typically set to occur after its contained query results have been consumed.

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Transaction functions

Transaction functions are *the recommended form for containing transactional units of work*. This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic. Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are in a single instance environment, this routing has no impact but it does give you the flexibility should you choose to later adopt a clustered environment.

Before writing a transaction function it is important to ensure that any side-effects carried out by a transaction function should be designed to be idempotent. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

	When a transaction fails, the driver retry logic is invoked. For several failure cases, the transaction can be immediately retried against a different server.
i	These cases include connection issues, server role changes (e.g. leadership elections) and transient errors. Retry logic can be configured when creating a session.

Example 42. Asynchronous transaction functions

Transaction function

```
public async Task<List<string>> PrintAllProducts()
{
    List<string> result = null;
    var session = Driver.AsyncSession();

    try
    {
        // Wrap whole operation into an managed transaction and
        // get the results back.
        result = await session.ReadTransactionAsync(async tx =>
        {
            var products = new List<string>();

            // Send cypher query to the database
            var reader = await tx.RunAsync(
                "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
                new {id = 0} // Parameters in the query, if any
            );

            // Loop through the records asynchronously
            while (await reader.FetchAsync())
            {
                // Each current read in buffer can be reached via Current
                products.Add(reader.Current[0].ToString());
            }

            return products;
        });
    }
    finally
    {
        // asynchronously close session
        await session.CloseAsync();
    }

    return result;
}
```

Auto-commit transactions

An auto-commit transaction is a basic but limited form of transaction. Such a transaction consists of only one Cypher query and is not automatically replayed on failure. Therefore any error scenarios must be handled by the client application itself.

Auto-commit transactions are intended to be used for simple use cases such as when learning Cypher or writing one-off scripts.

> ⚠️ It is not recommended to use auto-commit transactions in production environments.

> **ⓘ** Unlike other kinds of Cypher Query, `PERIODIC COMMIT` queries do not participate in the causal chain.
>
> Therefore, the only way to execute them from a driver is to use auto-commit transactions.
>
> Please refer to the Cypher Manual → `PERIODIC COMMIT` query hint.

Example 43. Asynchronous auto-commit transactions

Auto-commit transaction

```csharp
public async Task<List<string>> ReadProductTitles()
{
    var records = new List<string>();
    var session = Driver.AsyncSession();

    try
    {
        // Send cypher query to the database.
        // The existing IResult interface implements IEnumerable
        // and does not play well with asynchronous use cases. The replacement
        // IResultCursor interface is returned from the RunAsync
        // family of methods instead and provides async capable methods.
        var reader = await session.RunAsync(
            "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
            new {id = 0} // Parameters in the query, if any
        );

        // Loop through the records asynchronously
        while (await reader.FetchAsync())
        {
            // Each current read in buffer can be reached via Current
            records.Add(reader.Current[0].ToString());
        }
    }
    finally
    {
        // asynchronously close session
        await session.CloseAsync();
    }

    return records;
}
```

Consuming results

The asynchronous session API provides language-idiomatic methods to aid integration with asynchronous applications and frameworks.

Example 44. Asynchronous consuming results

Consuming results

```csharp
public async Task<List<string>> GetPeopleAsync()
{
    var session = Driver.AsyncSession();
    try
    {
        return await session.ReadTransactionAsync(async tx =>
        {
            var result = await tx.RunAsync("MATCH (a:Person) RETURN a.name ORDER BY a.name");

            return await result.ToListAsync(r => r[0].As<string>());
        });
    }
    finally
    {
        await session.CloseAsync();
    }
}
```

Reactive sessions

Available in:

- .NET Driver
- Java Driver
- JavaScript Driver

Starting with **Neo4j 4.0**, the reactive processing of queries is supported. This can be achieved through reactive sessions. Reactive sessions allow for dynamic management of the data that is being exchanged between the driver and the server.

Typical of reactive programming, consumers control the rate at which they consume records from queries and the driver in turn manages the rate at which records are requested from the server. Flow control is supported throughout the entire Neo4j stack, meaning that the query engine responds correctly to the flow control signals. This results in far more efficient resource handling and ensures that the receiving side is not forced to buffer arbitrary amounts of data.

For more information about reactive stream, please see the following:

- The Reactive Manifesto
- Reactive Streams for JVM and JavaScript
- Project Reactor reference documentation
- Introduction to Reactive Extensions to .NET

> ℹ️ Reactive sessions will typically be used in a client application that is already oriented towards the reactive style; it is expected that a reactive dependency or framework is in place.
>
> Refer to Get started for more information on recommended dependencies.

Lifecycle

Session lifetime begins with session construction. A session then exists until it is closed, which is typically set to occur after its contained query results have been consumed.

Transaction functions

This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic. Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are in a single instance environment, this routing has no impact but it does give you the flexibility should you choose to later adopt a clustered environment.

Before writing a transaction function it is important to ensure that any side-effects carried out by a transaction function should be designed to be idempotent. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

> ℹ️ When a transaction fails, the driver retry logic is invoked. For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors. Retry logic can be configured when creating a session.

Example 45. Reactive transaction functions

Reactive transaction function

```csharp
public IObservable<string> PrintAllProducts()
{
    var session = Driver.RxSession();

    return session.ReadTransaction(tx =>
    {
        return tx.Run(
                "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
                new {id = 0} // Parameters in the query, if any
            )
            .Records()
            .Select(record => record[0].ToString());
    }).OnErrorResumeNext(session.Close<string>());
}
```

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Auto-commit transactions

An auto-commit transaction is a basic but limited form of transaction. Such a transaction consists of only one Cypher Query and is not automatically replayed on failure. Therefore any error scenarios will need to be handled by the client application itself.

Auto-commit transactions are intended to be used for simple use cases such as when learning **Cypher** or writing one-off scripts.

⚠️	It is not recommended to use auto-commit transactions in production environments.

ℹ️	The only way to execute `PERIODIC COMMIT` Cypher Queries is to auto-commit the transaction. Unlike other kinds of Cypher query, `PERIODIC COMMIT` queries do not participate in the causal chain. Please refer to the Cypher Manual → PERIODIC COMMIT query hint.

Example 46. Reactive auto-commit transactions

Auto-commit transactions

```csharp
public IObservable<string> ReadProductTitles()
{
    var session = Driver.RxSession();

    return session.Run(
            "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
            new {id = 0} // Parameters in the query, if any
        )
        .Records()
        .Select(record => record[0].ToString())
        .OnErrorResumeNext(session.Close<string>());
}
```

Consuming results

To consume data from a query in a reactive session, a subscriber is required to handle the results that are being returned by the publisher.

Each transaction corresponds to a data flow which supplies the data from the server. Result processing begins when records are pulled from this flow. Only one subscriber may pull data from a given flow.

Example 47. Reactive consuming results

Consuming results

```csharp
public IObservable<string> GetPeople()
{
    var session = Driver.RxSession();
    return session.ReadTransaction(tx =>
    {
        return tx.Run("MATCH (a:Person) RETURN a.name ORDER BY a.name")
            .Records()
            .Select(record => record[0].As<string>());
    }).OnErrorResumeNext(session.Close<string>());
}
```

Cancellation

As per the reactive stream specification, a reactive data flow can be cancelled part way through. This prematurely commits or rolls back the transaction and stops the query engine from producing any more records.

Session configuration

Bookmarks

The mechanism which ensures causal consistency between transactions within a session. Bookmarks are implicitly passed between transactions within a single session to meet the causal consistency

requirements. There may be scenarios where you might want to use the bookmark from one session in a different new session.

Default: None (Sessions will initially be created without a bookmark)

Default Access Mode

A fallback for the access mode setting when transaction functions are not used. Typically, access mode is set per transaction by calling the appropriate transaction function method. In other cases, this setting is inherited. Note that transaction functions will ignore/override this setting.

Default: Write

Database

The database with which the session will interact. When you are working with a database which is not the default (i.e. the `system` database or another database in Neo4j 4.0 Enterprise Edition), you can explicitly configure the database which the driver is executing transactions against. See Operations Manual → The default database for more information on databases.

Default: the default database as configured on the server.

`Fetch Size`

The number of records to fetch in each batch from the server. Neo4j 4.0 introduces the ability to pull records in batches, allowing the client application to take control of data population and apply back pressure to the server. This `FetchSize` applies to simple sessions and async-sessions whereas reactive sessions can be controlled directly using the request method of the subscription.

Default: 1000 records

The session API

This chapter describes the following:

- Simple sessions
  - Lifecycle
  - Transaction functions
  - Auto-commit transactions
  - Consuming results
  - Retaining results
- Asynchronous sessions
  - Transaction functions
  - Auto-commit transaction
  - Consuming results
- Reactive sessions
  - Lifecycle

Simple sessions

i	Available in: - .NET Driver - Go Driver - Java Driver - Python Driver

Simple sessions provide a "classic" *blocking style* API for Cypher execution. In general, simple sessions provide the easiest programming style to work with since API calls are executed in a strictly sequential fashion.

Lifecycle

The session lifetime extends from session construction to session closure. In languages that support them, simple sessions are usually scoped within a context block; this ensures that they are properly closed and that any underlying connections are released and not leaked.

Example 48. Session construction and closure

```
using (var session = driver.Session(...)) {
    // transactions go here
  }
```

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See [Session configuration](#) for more details.

Transaction functions

Transaction functions are used for containing transactional units of work. This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic.

Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Transaction functions can be called as either read or write operations. This choice will route the transaction

to an appropriate server within a clustered environment. If you are operating in a single instance environment, this routing has no impact. It does give you flexibility if you choose to adopt a clustered environment later on.

Before writing a transaction function *it is important to ensure that it is designed to be idempotent*. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

> ℹ️ *Transaction functions are the recommended form for containing transactional units of work*.
>
> *When a transaction fails, the driver retry logic is invoked*. For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors.

Example 49. Transaction function

```
public void AddPerson(string name)
{
    using (var session = Driver.Session())
    {
        session.WriteTransaction(tx => tx.Run("CREATE (a:Person {name: $name})", new {name}));
    }
}
```

Auto-commit transactions

An *auto-commit transaction* is a *basic but limited form of transaction*. Such a transaction consists of only one Cypher query and is not automatically replayed on failure. Therefore any error scenarios must be handled by the client application itself.

Auto-commit transactions are *intended to be used for simple use cases* such as when learning **Cypher** or writing one-off scripts.

> ⚠️ It is not recommended to use auto-commit transactions in production environments.

> ℹ️ Unlike other kinds of Cypher Query, `PERIODIC COMMIT` queries do not participate in the causal chain.
>
> Therefore, the only way to execute `PERIODIC COMMIT` from a driver is to use auto-commit transactions.
>
> Please refer to the Cypher Manual → `PERIODIC COMMIT` query hint.

Example 50. Simple auto-commit transactions

```csharp
public void AddPerson(string name)
{
    using (var session = Driver.Session())
    {
        session.Run("CREATE (a:Person {name: $name})", new {name});
    }
}
```

Consuming results

Query results are typically consumed as a stream of records. The drivers provide a language-idiomatic way to iterate through that stream.

Example 51. Consuming results

```csharp
public List<string> GetPeople()
{
    using (var session = Driver.Session())
    {
        return session.ReadTransaction(tx =>
        {
            var result = tx.Run("MATCH (a:Person) RETURN a.name ORDER BY a.name");
            return result.Select(record => record[0].As<string>()).ToList();
        });
    }
}
```

Retaining results

Within a session, *only one result stream can be active at any one time*. Therefore, if the result of one query is not fully consumed before another query is executed, the remainder of the first result will be automatically buffered within the result object.

Figure 8. Result buffer

This buffer provides a staging point for results, and divides result handling into *fetching* (moving from the network to the buffer) and *consuming* (moving from the buffer to the application).

> For large results, the result buffer may require a significant amount of memory.
>
> For this reason, *it is recommended to consume results in order wherever possible*.

Client applications can choose to take control of more advanced query patterns by explicitly retaining results. Such explicit retention may also be useful when a result needs to be saved for future processing. The drivers offer support for this process, as per the example below:

Example 52. Retain results for further processing

```csharp
public int AddEmployees(string companyName)
{
    using (var session = Driver.Session())
    {
        var persons =
            session.ReadTransaction(tx => tx.Run("MATCH (a:Person) RETURN a.name AS name"
).ToList());
        return persons.Sum(person => session.WriteTransaction(tx =>
        {
            tx.Run("MATCH (emp:Person {name: $person_name}) " +
                "MERGE (com:Company {name: $company_name}) " +
                "MERGE (emp)-[:WORKS_FOR]->(com)",
                new {person_name = person["name"].As<string>(), company_name = companyName});
            return 1;
        }));
    }
}
```

Asynchronous sessions

	Available in: • .NET Driver • Java Driver • JavaScript Driver

Asynchronous sessions provide an API wherein function calls typically return available objects such as futures. This allows client applications to work within asynchronous frameworks and take advantage of cooperative multitasking.

Lifecycle

Session lifetime begins with session construction. A session then exists until it is closed, which is typically set to occur after its contained query results have been consumed.

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Transaction functions

Transaction functions are *the recommended form for containing transactional units of work*. This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic. Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are in a single instance environment, this routing has no impact but it does give you the flexibility should you choose to later adopt a clustered environment.

Before writing a transaction function it is important to ensure that any side-effects carried out by a transaction function should be designed to be idempotent. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

	When a transaction fails, the driver retry logic is invoked. For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors. Retry logic can be configured when creating a session.

Example 53. Asynchronous transaction functions

Transaction function

```csharp
public async Task<List<string>> PrintAllProducts()
{
    List<string> result = null;
    var session = Driver.AsyncSession();

    try
    {
        // Wrap whole operation into an managed transaction and
        // get the results back.
        result = await session.ReadTransactionAsync(async tx =>
        {
            var products = new List<string>();

            // Send cypher query to the database
            var reader = await tx.RunAsync(
                "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
                new {id = 0} // Parameters in the query, if any
            );

            // Loop through the records asynchronously
            while (await reader.FetchAsync())
            {
                // Each current read in buffer can be reached via Current
                products.Add(reader.Current[0].ToString());
            }

            return products;
        });
    }
    finally
    {
        // asynchronously close session
        await session.CloseAsync();
    }

    return result;
}
```

Auto-commit transactions

An auto-commit transaction is a basic but limited form of transaction. Such a transaction consists of only one Cypher query and is not automatically replayed on failure. Therefore any error scenarios must be handled by the client application itself.

Auto-commit transactions are intended to be used for simple use cases such as when learning Cypher or writing one-off scripts.

⚠	It is not recommended to use auto-commit transactions in production environments.

> ℹ️ Unlike other kinds of Cypher Query, `PERIODIC COMMIT` queries do not participate in the causal chain.
>
> Therefore, the only way to execute them from a driver is to use auto-commit transactions.
>
> Please refer to the Cypher Manual → `PERIODIC COMMIT` query hint.

Example 54. Asynchronous auto-commit transactions

Auto-commit transaction

```csharp
public async Task<List<string>> ReadProductTitles()
{
    var records = new List<string>();
    var session = Driver.AsyncSession();

    try
    {
        // Send cypher query to the database.
        // The existing IResult interface implements IEnumerable
        // and does not play well with asynchronous use cases. The replacement
        // IResultCursor interface is returned from the RunAsync
        // family of methods instead and provides async capable methods.
        var reader = await session.RunAsync(
            "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
            new {id = 0} // Parameters in the query, if any
        );

        // Loop through the records asynchronously
        while (await reader.FetchAsync())
        {
            // Each current read in buffer can be reached via Current
            records.Add(reader.Current[0].ToString());
        }
    }
    finally
    {
        // asynchronously close session
        await session.CloseAsync();
    }

    return records;
}
```

Consuming results

The asynchronous session API provides language-idiomatic methods to aid integration with asynchronous applications and frameworks.

Example 55. Asynchronous consuming results

Consuming results

```csharp
public async Task<List<string>> GetPeopleAsync()
{
    var session = Driver.AsyncSession();
    try
    {
        return await session.ReadTransactionAsync(async tx =>
        {
            var result = await tx.RunAsync("MATCH (a:Person) RETURN a.name ORDER BY a.name");

            return await result.ToListAsync(r => r[0].As<string>());
        });
    }
    finally
    {
        await session.CloseAsync();
    }
}
```

Reactive sessions

Available in:

- .NET Driver
- Java Driver
- JavaScript Driver

Starting with **Neo4j 4.0**, the reactive processing of queries is supported. This can be achieved through reactive sessions. Reactive sessions allow for dynamic management of the data that is being exchanged between the driver and the server.

Typical of reactive programming, consumers control the rate at which they consume records from queries and the driver in turn manages the rate at which records are requested from the server. Flow control is supported throughout the entire Neo4j stack, meaning that the query engine responds correctly to the flow control signals. This results in far more efficient resource handling and ensures that the receiving side is not forced to buffer arbitrary amounts of data.

For more information about reactive stream, please see the following:

- The Reactive Manifesto
- Reactive Streams for JVM and JavaScript
- Project Reactor reference documentation
- Introduction to Reactive Extensions to .NET

> ℹ️ Reactive sessions will typically be used in a client application that is already oriented towards the reactive style; it is expected that a reactive dependency or framework is in place.
>
> Refer to Get started for more information on recommended dependencies.

Lifecycle

Session lifetime begins with session construction. A session then exists until it is closed, which is typically set to occur after its contained query results have been consumed.

Transaction functions

This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic. Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are in a single instance environment, this routing has no impact but it does give you the flexibility should you choose to later adopt a clustered environment.

Before writing a transaction function it is important to ensure that any side-effects carried out by a transaction function should be designed to be idempotent. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

> ℹ️ When a transaction fails, the driver retry logic is invoked. For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors. Retry logic can be configured when creating a session.

Example 56. Reactive transaction functions

Reactive transaction function

```csharp
public IObservable<string> PrintAllProducts()
{
    var session = Driver.RxSession();

    return session.ReadTransaction(tx =>
    {
        return tx.Run(
                "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
                new {id = 0} // Parameters in the query, if any
            )
            .Records()
            .Select(record => record[0].ToString());
    }).OnErrorResumeNext(session.Close<string>());
}
```

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Auto-commit transactions

An auto-commit transaction is a basic but limited form of transaction. Such a transaction consists of only one Cypher Query and is not automatically replayed on failure. Therefore any error scenarios will need to be handled by the client application itself.

Auto-commit transactions are intended to be used for simple use cases such as when learning **Cypher** or writing one-off scripts.

⚠️	It is not recommended to use auto-commit transactions in production environments.

ℹ️	The only way to execute `PERIODIC COMMIT` Cypher Queries is to auto-commit the transaction.
	Unlike other kinds of Cypher query, `PERIODIC COMMIT` queries do not participate in the causal chain.
	Please refer to the Cypher Manual → PERIODIC COMMIT query hint.

Example 57. Reactive auto-commit transactions

> **Auto-commit transactions**
>
> ```csharp
> public IObservable<string> ReadProductTitles()
> {
>     var session = Driver.RxSession();
>
>     return session.Run(
>             "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
>             new {id = 0} // Parameters in the query, if any
>         )
>         .Records()
>         .Select(record => record[0].ToString())
>         .OnErrorResumeNext(session.Close<string>());
> }
> ```

Consuming results

To consume data from a query in a reactive session, a subscriber is required to handle the results that are being returned by the publisher.

Each transaction corresponds to a data flow which supplies the data from the server. Result processing begins when records are pulled from this flow. Only one subscriber may pull data from a given flow.

Example 58. Reactive consuming results

> **Consuming results**
>
> ```csharp
> public IObservable<string> GetPeople()
> {
>     var session = Driver.RxSession();
>     return session.ReadTransaction(tx =>
>     {
>         return tx.Run("MATCH (a:Person) RETURN a.name ORDER BY a.name")
>             .Records()
>             .Select(record => record[0].As<string>());
>     }).OnErrorResumeNext(session.Close<string>());
> }
> ```

Cancellation

As per the reactive stream specification, a reactive data flow can be cancelled part way through. This prematurely commits or rolls back the transaction and stops the query engine from producing any more records.

Session configuration

Bookmarks

> The mechanism which ensures causal consistency between transactions within a session. Bookmarks are implicitly passed between transactions within a single session to meet the causal consistency

requirements. There may be scenarios where you might want to use the bookmark from one session in a different new session.

Default: None (Sessions will initially be created without a bookmark)

Default Access Mode

A fallback for the access mode setting when transaction functions are not used. Typically, access mode is set per transaction by calling the appropriate transaction function method. In other cases, this setting is inherited. Note that transaction functions will ignore/override this setting.

Default: Write

Database

The database with which the session will interact. When you are working with a database which is not the default (i.e. the `system` database or another database in Neo4j 4.0 Enterprise Edition), you can explicitly configure the database which the driver is executing transactions against. See Operations Manual → The default database for more information on databases.

Default: the default database as configured on the server.

end::Database[]

`Fetch Size`

The number of records to fetch in each batch from the server. Neo4j 4.0 introduces the ability to pull records in batches, allowing the client application to take control of data population and apply back pressure to the server. This `FetchSize` applies to simple sessions and async-sessions whereas reactive sessions can be controlled directly using the request method of the subscription.

Default: 1000 records

The session API

This chapter describes the following:

- Simple sessions
    - Lifecycle
    - Transaction functions
    - Auto-commit transactions
    - Consuming results
    - Retaining results
- Asynchronous sessions
    - Transaction functions
    - Auto-commit transaction
    - Consuming results
- Reactive sessions

Simple sessions

<table>
<tr>
<td>

ℹ️

</td>
<td>

Available in:

- .NET Driver

- Go Driver

- Java Driver

- Python Driver

</td>
</tr>
</table>

Simple sessions provide a "classic" *blocking style* API for Cypher execution. In general, simple sessions provide the easiest programming style to work with since API calls are executed in a strictly sequential fashion.

Lifecycle

The session lifetime extends from session construction to session closure. In languages that support them, simple sessions are usually scoped within a context block; this ensures that they are properly closed and that any underlying connections are released and not leaked.

Example 59. Session construction and closure

```
using (var session = driver.Session(...)) {
    // transactions go here
}
```

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See [Session configuration](#) for more details.

Transaction functions

Transaction functions are used for containing transactional units of work. This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic.

Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Transaction functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are operating in a single instance environment, this routing has no impact. It does give you flexibility if you choose to adopt a clustered environment later on.

Before writing a transaction function *it is important to ensure that it is designed to be idempotent*. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

> **i** *Transaction functions are the recommended form for containing transactional units of work.*
>
> *When a transaction fails, the driver retry logic is invoked*. For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors.

Example 60. Transaction function

```
public void AddPerson(string name)
{
    using (var session = Driver.Session())
    {
        session.WriteTransaction(tx => tx.Run("CREATE (a:Person {name: $name})", new {name}));
    }
}
```

Auto-commit transactions

An **auto-commit transaction** is a **basic but limited form of transaction**. Such a transaction consists of only one Cypher query and is not automatically replayed on failure. Therefore any error scenarios must be handled by the client application itself.

Auto-commit transactions are *intended to be used for simple use cases* such as when learning **Cypher** or writing one-off scripts.

> **!** It is not recommended to use auto-commit transactions in production environments.

> **i** Unlike other kinds of Cypher Query, `PERIODIC COMMIT` queries do not participate in the causal chain.
>
> Therefore, the only way to execute `PERIODIC COMMIT` from a driver is to use auto-commit transactions.
>
> Please refer to the Cypher Manual → `PERIODIC COMMIT` query hint.

Example 61. Simple auto-commit transactions

```
public void AddPerson(string name)
{
    using (var session = Driver.Session())
    {
        session.Run("CREATE (a:Person {name: $name})", new {name});
    }
}
```

Consuming results

Query results are typically consumed as a stream of records. The drivers provide a language-idiomatic way to iterate through that stream.

Example 62. Consuming results

```
public List<string> GetPeople()
{
    using (var session = Driver.Session())
    {
        return session.ReadTransaction(tx =>
        {
            var result = tx.Run("MATCH (a:Person) RETURN a.name ORDER BY a.name");
            return result.Select(record => record[0].As<string>()).ToList();
        });
    }
}
```

Retaining results

Within a session, *only one result stream can be active at any one time*. Therefore, if the result of one query is not fully consumed before another query is executed, the remainder of the first result will be automatically buffered within the result object.

Figure 9. Result buffer

This buffer provides a staging point for results, and divides result handling into *fetching* (moving from the network to the buffer) and *consuming* (moving from the buffer to the application).

> For large results, the result buffer may require a significant amount of memory.
>
> For this reason, *it is recommended to consume results in order wherever possible*.

Client applications can choose to take control of more advanced query patterns by explicitly retaining results. Such explicit retention may also be useful when a result needs to be saved for future processing. The drivers offer support for this process, as per the example below:

Example 63. Retain results for further processing

```csharp
public int AddEmployees(string companyName)
{
    using (var session = Driver.Session())
    {
        var persons =
            session.ReadTransaction(tx => tx.Run("MATCH (a:Person) RETURN a.name AS name"
).ToList());
        return persons.Sum(person => session.WriteTransaction(tx =>
        {
            tx.Run("MATCH (emp:Person {name: $person_name}) " +
                    "MERGE (com:Company {name: $company_name}) " +
                    "MERGE (emp)-[:WORKS_FOR]->(com)",
                new {person_name = person["name"].As<string>(), company_name = companyName});
            return 1;
        }));
    }
}
```

Asynchronous sessions

Asynchronous sessions provide an API wherein function calls typically return available objects such as futures. This allows client applications to work within asynchronous frameworks and take advantage of cooperative multitasking.

Lifecycle

Session lifetime begins with session construction. A session then exists until it is closed, which is typically set to occur after its contained query results have been consumed.

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Transaction functions

Transaction functions are *the recommended form for containing transactional units of work*. This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic. Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are in a single instance environment, this routing has no impact but it does give you the flexibility should you choose to later adopt a clustered environment.

Before writing a transaction function it is important to ensure that any side-effects carried out by a transaction function should be designed to be idempotent. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

i

When a transaction fails, the driver retry logic is invoked. For several failure cases, the transaction can be immediately retried against a different server.

These cases include connection issues, server role changes (e.g. leadership elections) and transient errors. Retry logic can be configured when creating a session.

Example 64. Asynchronous transaction functions

Transaction function

```csharp
public async Task<List<string>> PrintAllProducts()
{
    List<string> result = null;
    var session = Driver.AsyncSession();

    try
    {
        // Wrap whole operation into an managed transaction and
        // get the results back.
        result = await session.ReadTransactionAsync(async tx =>
        {
            var products = new List<string>();

            // Send cypher query to the database
            var reader = await tx.RunAsync(
                "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
                new {id = 0} // Parameters in the query, if any
            );

            // Loop through the records asynchronously
            while (await reader.FetchAsync())
            {
                // Each current read in buffer can be reached via Current
                products.Add(reader.Current[0].ToString());
            }

            return products;
        });
    }
    finally
    {
        // asynchronously close session
        await session.CloseAsync();
    }

    return result;
}
```

Auto-commit transactions

An auto-commit transaction is a basic but limited form of transaction. Such a transaction consists of only one Cypher query and is not automatically replayed on failure. Therefore any error scenarios must be handled by the client application itself.

Auto-commit transactions are intended to be used for simple use cases such as when learning Cypher or writing one-off scripts.

> ⚠️ It is not recommended to use auto-commit transactions in production environments.

> ℹ️ Unlike other kinds of Cypher Query, `PERIODIC COMMIT` queries do not participate in the causal chain.
>
> Therefore, the only way to execute them from a driver is to use auto-commit transactions.
>
> Please refer to the Cypher Manual → `PERIODIC COMMIT` query hint.

Example 65. Asynchronous auto-commit transactions

Auto-commit transaction

```csharp
public async Task<List<string>> ReadProductTitles()
{
    var records = new List<string>();
    var session = Driver.AsyncSession();

    try
    {
        // Send cypher query to the database.
        // The existing IResult interface implements IEnumerable
        // and does not play well with asynchronous use cases. The replacement
        // IResultCursor interface is returned from the RunAsync
        // family of methods instead and provides async capable methods.
        var reader = await session.RunAsync(
            "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
            new {id = 0} // Parameters in the query, if any
        );

        // Loop through the records asynchronously
        while (await reader.FetchAsync())
        {
            // Each current read in buffer can be reached via Current
            records.Add(reader.Current[0].ToString());
        }
    }
    finally
    {
        // asynchronously close session
        await session.CloseAsync();
    }

    return records;
}
```

Consuming results

The asynchronous session API provides language-idiomatic methods to aid integration with asynchronous applications and frameworks.

Example 66. Asynchronous consuming results

Consuming results

```csharp
public async Task<List<string>> GetPeopleAsync()
{
    var session = Driver.AsyncSession();
    try
    {
        return await session.ReadTransactionAsync(async tx =>
        {
            var result = await tx.RunAsync("MATCH (a:Person) RETURN a.name ORDER BY a.name");

            return await result.ToListAsync(r => r[0].As<string>());
        });
    }
    finally
    {
        await session.CloseAsync();
    }
}
```

Reactive sessions

Available in:

- .NET Driver

- Java Driver

- JavaScript Driver

Starting with **Neo4j 4.0**, the reactive processing of queries is supported. This can be achieved through reactive sessions. Reactive sessions allow for dynamic management of the data that is being exchanged between the driver and the server.

Typical of reactive programming, consumers control the rate at which they consume records from queries and the driver in turn manages the rate at which records are requested from the server. Flow control is supported throughout the entire Neo4j stack, meaning that the query engine responds correctly to the flow control signals. This results in far more efficient resource handling and ensures that the receiving side is not forced to buffer arbitrary amounts of data.

For more information about reactive stream, please see the following:

- The Reactive Manifesto

- Reactive Streams for JVM and JavaScript

- Project Reactor reference documentation

- Introduction to Reactive Extensions to .NET

> **ℹ** Reactive sessions will typically be used in a client application that is already oriented towards the reactive style; it is expected that a reactive dependency or framework is in place.
>
> Refer to Get started for more information on recommended dependencies.

Lifecycle

Session lifetime begins with session construction. A session then exists until it is closed, which is typically set to occur after its contained query results have been consumed.

Transaction functions

This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic. Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are in a single instance environment, this routing has no impact but it does give you the flexibility should you choose to later adopt a clustered environment.

Before writing a transaction function it is important to ensure that any side-effects carried out by a transaction function should be designed to be idempotent. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

> **ℹ** When a transaction fails, the driver retry logic is invoked. For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors. Retry logic can be configured when creating a session.

Example 67. Reactive transaction functions

> **Reactive transaction function**
>
> ```
> public IObservable<string> PrintAllProducts()
> {
>     var session = Driver.RxSession();
>
>     return session.ReadTransaction(tx =>
>     {
>         return tx.Run(
>                 "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
>                 new {id = 0} // Parameters in the query, if any
>             )
>             .Records()
>             .Select(record => record[0].ToString());
>     }).OnErrorResumeNext(session.Close<string>());
> }
> ```

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Auto-commit transactions

An auto-commit transaction is a basic but limited form of transaction. Such a transaction consists of only one Cypher Query and is not automatically replayed on failure. Therefore any error scenarios will need to be handled by the client application itself.

Auto-commit transactions are intended to be used for simple use cases such as when learning **Cypher** or writing one-off scripts.

⚠️	It is not recommended to use auto-commit transactions in production environments.

ℹ️	The only way to execute `PERIODIC COMMIT` Cypher Queries is to auto-commit the transaction. Unlike other kinds of Cypher query, `PERIODIC COMMIT` queries do not participate in the causal chain. Please refer to the Cypher Manual → PERIODIC COMMIT query hint.

Example 68. Reactive auto-commit transactions

> **Auto-commit transactions**
>
> ```csharp
> public IObservable<string> ReadProductTitles()
> {
>     var session = Driver.RxSession();
>
>     return session.Run(
>             "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
>             new {id = 0} // Parameters in the query, if any
>         )
>         .Records()
>         .Select(record => record[0].ToString())
>         .OnErrorResumeNext(session.Close<string>());
> }
> ```

Consuming results

To consume data from a query in a reactive session, a subscriber is required to handle the results that are being returned by the publisher.

Each transaction corresponds to a data flow which supplies the data from the server. Result processing begins when records are pulled from this flow. Only one subscriber may pull data from a given flow.

Example 69. Reactive consuming results

> **Consuming results**
>
> ```csharp
> public IObservable<string> GetPeople()
> {
>     var session = Driver.RxSession();
>     return session.ReadTransaction(tx =>
>     {
>         return tx.Run("MATCH (a:Person) RETURN a.name ORDER BY a.name")
>             .Records()
>             .Select(record => record[0].As<string>());
>     }).OnErrorResumeNext(session.Close<string>());
> }
> ```

Cancellation

As per the reactive stream specification, a reactive data flow can be cancelled part way through. This prematurely commits or rolls back the transaction and stops the query engine from producing any more records.

Session configuration

Bookmarks

The mechanism which ensures causal consistency between transactions within a session. Bookmarks are implicitly passed between transactions within a single session to meet the causal consistency

requirements. There may be scenarios where you might want to use the bookmark from one session in a different new session.

Default: None (Sessions will initially be created without a bookmark)

Default Access Mode

A fallback for the access mode setting when transaction functions are not used. Typically, access mode is set per transaction by calling the appropriate transaction function method. In other cases, this setting is inherited. Note that transaction functions will ignore/override this setting.

Default: Write

Database

The database with which the session will interact. When you are working with a database which is not the default (i.e. the `system` database or another database in Neo4j 4.0 Enterprise Edition), you can explicitly configure the database which the driver is executing transactions against. See Operations Manual → The default database for more information on databases.

Default: the default database as configured on the server.

`Fetch Size`

The number of records to fetch in each batch from the server. Neo4j 4.0 introduces the ability to pull records in batches, allowing the client application to take control of data population and apply back pressure to the server. This `FetchSize` applies to simple sessions and async-sessions whereas reactive sessions can be controlled directly using the request method of the subscription.

Default: 1000 records

The session API

This chapter describes the following:

- Simple sessions
  - ° Lifecycle
  - ° Transaction functions
  - ° Auto-commit transactions
  - ° Consuming results
  - ° Retaining results
- Asynchronous sessions
  - ° Transaction functions
  - ° Auto-commit transaction
  - ° Consuming results
- Reactive sessions
  - ° Lifecycle

Simple sessions

ℹ	Available in: • .NET Driver • Go Driver • Java Driver • Python Driver

Simple sessions provide a "classic" *blocking style* API for Cypher execution. In general, simple sessions provide the easiest programming style to work with since API calls are executed in a strictly sequential fashion.

Lifecycle

The session lifetime extends from session construction to session closure. In languages that support them, simple sessions are usually scoped within a context block; this ensures that they are properly closed and that any underlying connections are released and not leaked.

Example 70. Session construction and closure

```
using (var session = driver.Session(...)) {
    // transactions go here
}
```

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Transaction functions

Transaction functions are used for containing transactional units of work. This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic.

Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Transaction functions can be called as either read or write operations. This choice will route the transaction

to an appropriate server within a clustered environment. If you are operating in a single instance environment, this routing has no impact. It does give you flexibility if you choose to adopt a clustered environment later on.

Before writing a transaction function *it is important to ensure that it is designed to be idempotent*. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

> **ℹ** *Transaction functions are the recommended form for containing transactional units of work.*
>
> *When a transaction fails, the driver retry logic is invoked.* For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors.

Example 71. Transaction function

```
public void AddPerson(string name)
{
    using (var session = Driver.Session())
    {
        session.WriteTransaction(tx => tx.Run("CREATE (a:Person {name: $name})", new {name}));
    }
}
```

Auto-commit transactions

An *auto-commit transaction* is a *basic but limited form of transaction*. Such a transaction consists of only one Cypher query and is not automatically replayed on failure. Therefore any error scenarios must be handled by the client application itself.

Auto-commit transactions are *intended to be used for simple use cases* such as when learning **Cypher** or writing one-off scripts.

> **⚠** It is not recommended to use auto-commit transactions in production environments.

> **ℹ** Unlike other kinds of Cypher Query, `PERIODIC COMMIT` queries do not participate in the causal chain.
>
> Therefore, the only way to execute `PERIODIC COMMIT` from a driver is to use auto-commit transactions.
>
> Please refer to the Cypher Manual → `PERIODIC COMMIT` query hint.

Example 72. Simple auto-commit transactions

```csharp
public void AddPerson(string name)
{
    using (var session = Driver.Session())
    {
        session.Run("CREATE (a:Person {name: $name})", new {name});
    }
}
```

Consuming results

Query results are typically consumed as a stream of records. The drivers provide a language-idiomatic way to iterate through that stream.

Example 73. Consuming results

```csharp
public List<string> GetPeople()
{
    using (var session = Driver.Session())
    {
        return session.ReadTransaction(tx =>
        {
            var result = tx.Run("MATCH (a:Person) RETURN a.name ORDER BY a.name");
            return result.Select(record => record[0].As<string>()).ToList();
        });
    }
}
```

Retaining results

Within a session, *only one result stream can be active at any one time*. Therefore, if the result of one query is not fully consumed before another query is executed, the remainder of the first result will be automatically buffered within the result object.

Figure 10. Result buffer

This buffer provides a staging point for results, and divides result handling into *fetching* (moving from the network to the buffer) and *consuming* (moving from the buffer to the application).

> For large results, the result buffer may require a significant amount of memory.
>
> For this reason, *it is recommended to consume results in order wherever possible*.

Client applications can choose to take control of more advanced query patterns by explicitly retaining results. Such explicit retention may also be useful when a result needs to be saved for future processing. The drivers offer support for this process, as per the example below:

Example 74. Retain results for further processing

```csharp
public int AddEmployees(string companyName)
{
    using (var session = Driver.Session())
    {
        var persons =
            session.ReadTransaction(tx => tx.Run("MATCH (a:Person) RETURN a.name AS name"
).ToList());
        return persons.Sum(person => session.WriteTransaction(tx =>
        {
            tx.Run("MATCH (emp:Person {name: $person_name}) " +
                    "MERGE (com:Company {name: $company_name}) " +
                    "MERGE (emp)-[:WORKS_FOR]->(com)",
                new {person_name = person["name"].As<string>(), company_name = companyName});
            return 1;
        }));
    }
}
```

Asynchronous sessions

	Available in: • .NET Driver • Java Driver • JavaScript Driver

Asynchronous sessions provide an API wherein function calls typically return available objects such as futures. This allows client applications to work within asynchronous frameworks and take advantage of cooperative multitasking.

Lifecycle

Session lifetime begins with session construction. A session then exists until it is closed, which is typically set to occur after its contained query results have been consumed.

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Transaction functions

Transaction functions are *the recommended form for containing transactional units of work*. This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic. Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are in a single instance environment, this routing has no impact but it does give you the flexibility should you choose to later adopt a clustered environment.

Before writing a transaction function it is important to ensure that any side-effects carried out by a transaction function should be designed to be idempotent. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

	When a transaction fails, the driver retry logic is invoked. For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors. Retry logic can be configured when creating a session.

Example 75. Asynchronous transaction functions

Transaction function

```csharp
public async Task<List<string>> PrintAllProducts()
{
    List<string> result = null;
    var session = Driver.AsyncSession();

    try
    {
        // Wrap whole operation into an managed transaction and
        // get the results back.
        result = await session.ReadTransactionAsync(async tx =>
        {
            var products = new List<string>();

            // Send cypher query to the database
            var reader = await tx.RunAsync(
                "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
                new {id = 0} // Parameters in the query, if any
            );

            // Loop through the records asynchronously
            while (await reader.FetchAsync())
            {
                // Each current read in buffer can be reached via Current
                products.Add(reader.Current[0].ToString());
            }

            return products;
        });
    }
    finally
    {
        // asynchronously close session
        await session.CloseAsync();
    }

    return result;
}
```

Auto-commit transactions

An auto-commit transaction is a basic but limited form of transaction. Such a transaction consists of only one Cypher query and is not automatically replayed on failure. Therefore any error scenarios must be handled by the client application itself.

Auto-commit transactions are intended to be used for simple use cases such as when learning Cypher or writing one-off scripts.

⚠️	It is not recommended to use auto-commit transactions in production environments.

> ℹ️ Unlike other kinds of Cypher Query, `PERIODIC COMMIT` queries do not participate in the causal chain.
>
> Therefore, the only way to execute them from a driver is to use auto-commit transactions.
>
> Please refer to the Cypher Manual → `PERIODIC COMMIT` query hint.

Example 76. Asynchronous auto-commit transactions

Auto-commit transaction

```csharp
public async Task<List<string>> ReadProductTitles()
{
    var records = new List<string>();
    var session = Driver.AsyncSession();

    try
    {
        // Send cypher query to the database.
        // The existing IResult interface implements IEnumerable
        // and does not play well with asynchronous use cases. The replacement
        // IResultCursor interface is returned from the RunAsync
        // family of methods instead and provides async capable methods.
        var reader = await session.RunAsync(
            "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
            new {id = 0} // Parameters in the query, if any
        );

        // Loop through the records asynchronously
        while (await reader.FetchAsync())
        {
            // Each current read in buffer can be reached via Current
            records.Add(reader.Current[0].ToString());
        }
    }
    finally
    {
        // asynchronously close session
        await session.CloseAsync();
    }

    return records;
}
```

Consuming results

The asynchronous session API provides language-idiomatic methods to aid integration with asynchronous applications and frameworks.

Example 77. Asynchronous consuming results

Consuming results

```csharp
public async Task<List<string>> GetPeopleAsync()
{
    var session = Driver.AsyncSession();
    try
    {
        return await session.ReadTransactionAsync(async tx =>
        {
            var result = await tx.RunAsync("MATCH (a:Person) RETURN a.name ORDER BY a.name");

            return await result.ToListAsync(r => r[0].As<string>());
        });
    }
    finally
    {
        await session.CloseAsync();
    }
}
```

Reactive sessions

Available in:

- .NET Driver

- Java Driver

- JavaScript Driver

Starting with **Neo4j 4.0**, the reactive processing of queries is supported. This can be achieved through reactive sessions. Reactive sessions allow for dynamic management of the data that is being exchanged between the driver and the server.

Typical of reactive programming, consumers control the rate at which they consume records from queries and the driver in turn manages the rate at which records are requested from the server. Flow control is supported throughout the entire Neo4j stack, meaning that the query engine responds correctly to the flow control signals. This results in far more efficient resource handling and ensures that the receiving side is not forced to buffer arbitrary amounts of data.

For more information about reactive stream, please see the following:

- The Reactive Manifesto

- Reactive Streams for JVM and JavaScript

- Project Reactor reference documentation

- Introduction to Reactive Extensions to .NET

Lifecycle

Session lifetime begins with session construction. A session then exists until it is closed, which is typically set to occur after its contained query results have been consumed.

Transaction functions

This form of transaction requires minimal boilerplate code and allows for a clear separation of database queries and application logic. Transaction functions are also desirable since they encapsulate retry logic and allow for the greatest degree of flexibility when swapping out a single instance of server for a cluster.

Functions can be called as either read or write operations. This choice will route the transaction to an appropriate server within a clustered environment. If you are in a single instance environment, this routing has no impact but it does give you the flexibility should you choose to later adopt a clustered environment.

Before writing a transaction function it is important to ensure that any side-effects carried out by a transaction function should be designed to be idempotent. This is because a function may be executed multiple times if initial runs fail.

Any query results obtained within a transaction function should be consumed within that function, as connection-bound resources cannot be managed correctly when out of scope. To that end, transaction functions can return values but these should be derived values rather than raw results.

> ℹ️ When a transaction fails, the driver retry logic is invoked. For several failure cases, the transaction can be immediately retried against a different server. These cases include connection issues, server role changes (e.g. leadership elections) and transient errors. Retry logic can be configured when creating a session.

Example 78. Reactive transaction functions

Reactive transaction function

```csharp
public IObservable<string> PrintAllProducts()
{
    var session = Driver.RxSession();

    return session.ReadTransaction(tx =>
    {
        return tx.Run(
                "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
                new {id = 0} // Parameters in the query, if any
            )
            .Records()
            .Select(record => record[0].ToString());
    }).OnErrorResumeNext(session.Close<string>());
}
```

Sessions can be configured in a number of different ways. This is carried out by supplying configuration inside the session constructor. See Session configuration for more details.

Auto-commit transactions

An auto-commit transaction is a basic but limited form of transaction. Such a transaction consists of only one Cypher Query and is not automatically replayed on failure. Therefore any error scenarios will need to be handled by the client application itself.

Auto-commit transactions are intended to be used for simple use cases such as when learning **Cypher** or writing one-off scripts.

⚠️	It is not recommended to use auto-commit transactions in production environments.

ℹ️	The only way to execute `PERIODIC COMMIT` Cypher Queries is to auto-commit the transaction. Unlike other kinds of Cypher query, `PERIODIC COMMIT` queries do not participate in the causal chain. Please refer to the Cypher Manual → PERIODIC COMMIT query hint.

Example 79. Reactive auto-commit transactions

Auto-commit transactions

```csharp
public IObservable<string> ReadProductTitles()
{
    var session = Driver.RxSession();

    return session.Run(
            "MATCH (p:Product) WHERE p.id = $id RETURN p.title", // Cypher query
            new {id = 0} // Parameters in the query, if any
        )
        .Records()
        .Select(record => record[0].ToString())
        .OnErrorResumeNext(session.Close<string>());
}
```

Consuming results

To consume data from a query in a reactive session, a subscriber is required to handle the results that are being returned by the publisher.

Each transaction corresponds to a data flow which supplies the data from the server. Result processing begins when records are pulled from this flow. Only one subscriber may pull data from a given flow.

Example 80. Reactive consuming results

Consuming results

```csharp
public IObservable<string> GetPeople()
{
    var session = Driver.RxSession();
    return session.ReadTransaction(tx =>
    {
        return tx.Run("MATCH (a:Person) RETURN a.name ORDER BY a.name")
            .Records()
            .Select(record => record[0].As<string>());
    }).OnErrorResumeNext(session.Close<string>());
}
```

Cancellation

As per the reactive stream specification, a reactive data flow can be cancelled part way through. This prematurely commits or rolls back the transaction and stops the query engine from producing any more records.

Session configuration

Bookmarks

The mechanism which ensures causal consistency between transactions within a session. Bookmarks are implicitly passed between transactions within a single session to meet the causal consistency

requirements. There may be scenarios where you might want to use the bookmark from one session in a different new session.

Default: None (Sessions will initially be created without a bookmark)

Default Access Mode

A fallback for the access mode setting when transaction functions are not used. Typically, access mode is set per transaction by calling the appropriate transaction function method. In other cases, this setting is inherited. Note that transaction functions will ignore/override this setting.

Default: Write

Database

The database with which the session will interact. When you are working with a database which is not the default (i.e. the `system` database or another database in Neo4j 4.0 Enterprise Edition), you can explicitly configure the database which the driver is executing transactions against. See Operations Manual → The default database for more information on databases.

Default: the default database as configured on the server.

`Fetch Size`

The number of records to fetch in each batch from the server. Neo4j 4.0 introduces the ability to pull records in batches, allowing the client application to take control of data population and apply back pressure to the server. This `FetchSize` applies to simple sessions and async-sessions whereas reactive sessions can be controlled directly using the request method of the subscription.

Default: 1000 records

Appendix A: Driver terminology

access mode

The mode in which a transaction is executed, either read or write.

acquire (connection)

To borrow a driver connection that is not currently in use from a connection pool.

auto-commit

A single query which is wrapped in a transaction and committed automatically.

Bolt

Bolt is a Neo4j proprietary, binary protocol used for communication between client applications and database servers. Bolt is versioned independently from the database and the drivers.

Bolt Routing Protocol

The steps required for a driver to obtain a routing table from a cluster member.

Bolt server

A Neo4j instance that can accept incoming Bolt connections.

bookmark

A marker for a point in the transactional history of Neo4j.

causal chaining

A mechanism to ensure that the transactions carried out in a session are executed in order, even when each transaction may be carried out on different cluster members.

client application

A piece of software that interacts with a database server via a driver.

connection

A persistent communication channel between a client application and a database server.

connection pool

A set of connections maintained for quick access, that can be acquired and released as required.

driver (object)

A globally accessible controller for all database access.

driver (package)

A software library that provides access to Neo4j from a particular programming language. The Neo4j drivers implement the Bolt protocol.

query result

The stream of records that are returned on execution of a query.

release (connection)

To return a connection back into a connection pool after use.

routing driver

A driver that can route traffic to multiple members of a cluster using the routing protocol.

routing table

A set of server addresses that identify cluster members associated with roles.

server address

A combination of host name and port or IP address and port that targets a server.

session

A causally linked sequence of transactions.

thread safety

See https://en.wikipedia.org/wiki/Thread_safety.

transaction

A transaction comprises a unit of work performed against a database. It is treated in a coherent and reliable way, independent of other transactions. A transaction, by definition, must be atomic, consistent, isolated, and durable.

transaction function

The method of grouping a number of queries together which, when run in a session, are retried on failure.

transaction manager

The component/code responsible for deciding what to do if a transaction fails, i.e to retry, give up or do something else.

License