



The Neo4j Drivers Manual v1.7 for C#

Table of Contents

Get started	2
About the official drivers	2
Driver versions and installation	2
A "Hello World" example	3
Driver API docs	4
Client applications	5
The driver object	5
Connection URIs, encryption and trust.	6
Authentication.	13
Configuration.	15
Service unavailable.	17
Sessions and transactions.	18
Sessions	18
Transactions	18
Causal chaining.	20
Access modes	24
Asynchronous programming	25
Working with Cypher values	29
The Cypher type system	29
Statement results	31
Statement result summaries	34
Appendix A: Driver terminology.	35

Copyright © 2022 Neo4j, Inc.

License: [Creative Commons 4.0](#)

This is the driver manual for Neo4j official drivers version 1.7, and the Go driver version 1.8, authored by the Neo4j Team.

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.
- [Sessions and transactions](#) — How to create units of work and provide a logical context for that work.
- [Working with Cypher values](#) — 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

This section gives an overview of the official Neo4j drivers and how to connect to a Neo4j database with a "Hello World" example.

About the official drivers

Neo4j provides official drivers for many popular programming languages. These drivers are supported by Neo4j and come with full cluster routing support.

Community drivers also exist for many languages, but vary greatly in terms of feature sets, maturity, and support. To find out 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 1.x driver series

Language/framework	Versions supported
Java	Java 8+ (latest patch releases).
Python	CPython 2.7, 3.4, 3.5, 3.6 and 3.7 (latest patch releases). Note that Python 2 support is deprecated and will be discontinued in the 2.x series drivers.
JavaScript	All LTS versions of Node.JS, specifically the 4.x and 6.x series runtimes (https://github.com/nodejs/LTS).
.NET	.NET standard 1.3 (https://github.com/dotnet/standard/blob/master/docs/versions.md) and .NET 4.5.2 and above.
Go	Go 1.10+

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.



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

The driver *major.minor* version number describes the feature set available within that driver. Across languages, drivers are generally released at similar times, although patches are driver-dependent. The 1.x series drivers have been built for Neo4j 3.x.



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 to Drivers 4.0. For information, see [Upgrade and Migration Guide](#) → [Drivers changes](#).



The Go 1.8 driver does support multiple databases.

It is recommended to always use the latest driver release available. This will ensure that all server functionality is made available to client applications. To install a driver or to find out more about which driver versions are available, use the relevant language distribution system, described below.

Example 3. Acquire the driver

To find out the latest version of the driver, visit <https://www.nuget.org/packages/Neo4j.Driver/>.

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

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

You can also choose to install a certain version of the driver.

Below is the syntax for installing a certain version of the driver.

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

In the following example we are installing driver version 1.7.2.

```
PM> Install-Package Neo4j.Driver -Version 1.7.2
```

You can review the release notes for this driver [here](#).

A "Hello World" example

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

Example 4. Hello World

```
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 5. API docs

<https://neo4j.com/docs/api/dotnet-driver/1.7/>

Client applications

This section describes how to manage database connections within an application.

This section contains the following:

- [The driver object](#)
- [Connection URIs, encryption and trust](#)
 - [Connection URIs](#)
 - [Encryption](#)
 - [Trust](#)
 - [Examples](#)
- [Authentication](#)
 - [Basic authentication](#)
 - [Kerberos authentication](#)
 - [Custom authentication](#)
- [Configuration](#)
 - [Connection pool management](#)
 - [Connection timeout](#)
 - [Max retry time](#)
- [Service unavailable](#)

The driver object

A Neo4j client application will require a driver instance in order to provide access to the database. This driver should be made available to all parts of the application that need to interact with Neo4j. In languages where [thread safety](#) is an issue, the driver can be considered thread-safe.



A note on lifecycle

Applications will typically construct a driver instance on startup and destroy it on exit. Destroying a driver instance will immediately shut down any connections previously opened via that driver; for drivers that contain a connection pool, the entire pool will be shut down.

To construct a driver instance, a [connection URI](#) and [authentication information](#) must be supplied. Additional configuration details can be supplied if required. All of these details are immutable for the lifetime of the driver. Therefore, if multiple configurations are required (such as when working with multiple database users) then multiple driver instances must be used.

An example of driver construction and destruction can be seen below:

Example 6. 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, encryption and trust

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.

Connection URIs

The official Neo4j drivers currently support the following URI schemes and driver object types:

Table 2. Available URI schemes

URI Scheme	Driver object type
bolt	Direct driver
bolt+routing	Routing driver

Direct drivers (bolt)

A direct driver is created via a **bolt** URI, for example: **bolt://localhost:7687**. This kind of driver is used to maintain connections to a single database server and is typically used when working with a single Neo4j instance or when targeting a specific member of a cluster. Note that a routing driver is generally preferable when working with a Causal Cluster.

Routing drivers (bolt+routing)

A routing driver is created via a **bolt+routing** URI, for example: **bolt+routing://graph.example.com:7687**. The address in the URI must be that of a Core Server. This kind of driver uses the [Bolt Routing Protocol](#) and works in tandem with the cluster to route transactions to available cluster members.

Routing drivers with routing context

Routing drivers with routing context are an available option when using drivers of version 1.3 or above together with a Neo4j Causal Cluster of version 3.2 or above. In such a setup, a routing driver can include a

preferred routing context via the query part of the `bolt+routing` URI.

In the standard Neo4j configuration, routing contexts are defined on the server side by means of server policies. Thus the driver communicates the routing context to the cluster in the form of a server policy. It then obtains refined routing information back from the cluster, based on the server policy.

The address in the URI of a routing driver with routing context must be that of a Core Server.

Example 7. 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`:

```
bolt+routing://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](#).

Cluster bootstrap address resolution

Drivers have the ability to intercept cluster bootstrap address resolution. This is generally performed by a routing driver on construction or after complete contact has been lost with a cluster. This hook can be useful to augment regular DNS resolution and can be used to initialize a routing driver with multiple server addresses.

The intercept hook is provided as a driver configuration option and takes the form of a callback function that accepts a single input address and yields or returns multiple output addresses. This function may hard code the output addresses or may draw them from another configuration source, as required. Output addresses can contain host names or IP addresses and are forwarded to the connection pool as initial routers.




Figure 1. Cluster bootstrap address resolution

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

Example 8. Custom Address Resolver

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

public void AddPerson(string name)
{
    using (var driver = CreateDriverWithCustomResolver("bolt+routing://x.acme.com",
        AuthTokens.Basic(Username, Password),
        ServerAddress.From("a.acme.com", 7687), ServerAddress.From("b.acme.com", 7877),
        ServerAddress.From("c.acme.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);
    }
}
```

Encryption

TLS encryption is enabled for all connections by default. This can be disabled through configuration as follows:

Example 9. Unencrypted

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

The server can be modified to require encryption for all connections. Please see the [Operations Manual](#) → [Configure Neo4j connectors](#) for more information.

An attempt to connect to a server using an encryption setting not allowed by that server will result in a [Service unavailable](#) status.

Trust

During a TLS handshake, the server provides a certificate to the client application. The application can choose to accept or reject this certificate based on one of the following trust strategies:

Table 3. Trust strategies

Trust strategy	Description
<code>TRUST_ALL_CERTIFICATES</code> (default)	Accept any certificate provided by the server, regardless of CA chain.
<code>TRUST_CUSTOM_CA_SIGNED_CERTIFICATES</code>	Accept any certificate that can be verified against a custom CA.
<code>TRUST_SYSTEM_CA_SIGNED_CERTIFICATES</code>	Accept any certificate that can be verified against the system store.

[Server Name Indication \(SNI\)](#) is enabled by default, allowing proxy middleware to route encrypted Bolt connections to the correct backend server.


Example 10. Trust

```
public IDriver CreateDriverWithCustomizedTrustStrategy(string uri, string user, string password)
{
    return GraphDatabase.Driver(uri, AuthTokens.Basic(user, password),
        new Config {TrustStrategy = TrustStrategy.TrustAllCertificates});
}
```

Examples

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 11. 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	<pre>String uri = "bolt+routing://graph.example.com:7687"; IDriver driver = GraphDatabase.Driver(uri, auth, new Config {TrustStrategy = TrustStrategy.Trust SystemCaSignedCertificates});</pre>	This is the default (and only option) for Neo4j Aura
Neo4j 4.x	Unsecured	<pre>String uri = "bolt+routing://graph.example.com:7687"; IDriver driver = GraphDatabase.Driver(uri, auth, new Config {EncryptionLevel = EncryptionLevel.None});</pre>	This is the default for Neo4j 4.x series
Neo4j 4.x	Secured with full certificate	<pre>String uri = "bolt+routing://graph.example.com:7687"; IDriver driver = GraphDatabase.Driver(uri, auth, new Config {TrustStrategy = TrustStrategy.Trust SystemCaSignedCertificates});</pre>	
Neo4j 4.x	Secured with self-signed certificate	<pre>GraphDatabase.Driver("bolt+routing://graph.example.com:7687", auth)</pre>	

Product	Security	Code snippet	Comments
Neo4j 3.x	Secured with full certificate	<pre>String uri = "bolt+routing://graph.example.com:7687"; IDriver driver = GraphDatabase.Driver(uri, auth, new Config {TrustStrategy = TrustStrategy.Trust SystemCaSignedCertificates});</pre>	
Neo4j 3.x	Secured with self-signed certificate	<pre>GraphDatabase.Driver("bolt+routing://graph.example.com:7687", auth)</pre>	This is the default for Neo4j 3.x series
Neo4j 3.x	Unsecured	<pre>String uri = "bolt+routing://graph.example.com:7687"; IDriver driver = GraphDatabase.Driver(uri, auth, new Config {EncryptionLevel = EncryptionLevel.None});</pre>	



To connect to a service without routing, you can replace `bolt+routing` 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 12. Basic authentication

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



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 13. Kerberos authentication

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



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 14. 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),
        new Config {EncryptionLevel = EncryptionLevel.None});
}
```


Configuration

Connection pool management

The driver maintains a pool of connections. The pooled connections are reused by sessions and transactions to avoid the overhead added by establishing new connections for every query. The connection pool always starts up empty. New connections are created on demand by sessions and transactions. When a session or a transaction is done with its execution, the connection will be returned to the pool to be reused.

Application users can tune connection pool settings to configure the driver for different use cases based on client performance requirements and database resource consumption limits.

Detailed descriptions of connection pool settings available via driver configuration are listed below:

MaxConnectionLifetime

Pooled connections older than this threshold will be closed and removed from the pool. The actual removal happens during connection acquisition so that the new session returned is never backed by an old connection. Setting this option to a low value will cause a high connection churn, 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 value: 1h.

MaxConnectionPoolSize

This setting defines the maximum total number of connections allowed, per host, to be managed by the connection pool. In other words, for a direct driver, this sets the maximum number of connections towards a single database server. For a routing driver 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`.

Default value: This is different for different drivers, but is a number in the order of 100.

ConnectionAcquisitionTimeout

This setting limits the amount of time a session or transaction can spend waiting for a free connection to appear in the pool before throwing an exception. The exception thrown in this case is `ClientException`. Timeout only applies when connection pool is at its max capacity.

Default value: 1m.

Example 15. Connection pool management

```
public IDriver CreateDriverWithCustomizedConnectionPool(string uri, string user, string password)
{
    return GraphDatabase.Driver(uri, AuthTokens.Basic(user, password),
        new Config
        {
            MaxConnectionLifetime = TimeSpan.FromMinutes(30),
            MaxConnectionPoolSize = 50,
            ConnectionAcquisitionTimeout = TimeSpan.FromMinutes(2)
        });
}
```

Connection timeout

To configure the maximum time allowed to establish a connection, pass a duration value to the driver configuration. For example:

Example 16. Connection timeout

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

Max retry time

To configure retry behavior, supply a value for the maximum time in which to keep attempting retries of transaction functions. For example:

Example 17. Max retry time

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

Note that the time specified here does not take into account the running time of the unit of work itself, merely a limit after which retries will no longer be attempted.

Service unavailable

A Service unavailable status 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. Applications should be designed to cater for this eventuality.

Example 18. Service unavailable

```
public bool AddItem()
{
    try
    {
        using (var session = Driver.Session())
        {
            return session.WriteTransaction(
                tx =>
                {
                    tx.Run("CREATE (a:Item)");
                    return true;
                }
            );
        }
    }
    catch (ServiceUnavailableException)
    {
        return false;
    }
}
```

Sessions and transactions

This section describes how to create units of work and provide a logical context for that work.

Sessions

A session is a container for a sequence of transactions. Sessions borrow connections from a pool as required and so should be considered lightweight and disposable. In languages where [thread safety](#) is an issue, a session should not be considered thread-safe.

In languages that support them, 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 19. Session

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

Transactions

Transactions are atomic units of work consisting of one or more Cypher statement executions. A transaction is executed within a session.

To execute a Cypher statement, two pieces of information are required: the statement template and a keyed set of parameters. The template is a string containing placeholders that are substituted with parameter values at runtime. While it is possible to run non-parameterized Cypher, good programming practice is to use parameters in Cypher statements. This allows for caching of statements within the Cypher engine, which is beneficial for performance. Parameter values should adhere to [Working with Cypher values](#).

The Neo4j driver API provides for three forms of transaction:

- Auto-commit transactions
- Transaction functions
- Explicit transactions

Of these, only [transaction function](#) can be automatically replayed on failure.

Auto-commit Transactions

An auto-commit transaction is a simple but limited form of transaction. Such a transaction consists of only one Cypher statement, cannot be automatically replayed on failure, and cannot take part in a [causal chain](#).

An auto-commit transaction is invoked using the `session.run` method:

Example 20. Auto-commit transaction

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

Auto-commit transactions are sent to the network and acknowledged immediately. This means that multiple transactions cannot share network packets, thereby exhibiting a lesser network efficiency than other forms of transaction.

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 or when performance or resilience are a primary concern.

However, Auto-commit transactions are the only way to execute [USING PERIODIC COMMIT](#) Cypher statements.

Transaction functions

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

Example 21. Transaction function

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

Transaction functions are also able to handle connection problems and transient errors using an automatic retry mechanism. This retry capability can be [configured](#) on Driver construction.

Any query results obtained within a transaction function should be consumed within that function.

Transaction functions can return values but these should be derived values rather than raw results.

Explicit transactions

Explicit transactions are the longhand form of transaction functions, providing access to explicit **BEGIN**, **COMMIT** and **ROLLBACK** operations. While this form is useful for a handful of use cases, it is recommended to use transaction functions wherever possible.

Transaction configuration

Both auto-commit and explicit transactions can be assigned *metadata*, and a *timeout*.

Metadata is a map of arbitrary values that will be attached to the executing transaction. It will be visible in the output of the procedures `dbms.listQueries` and `dbms.listTransactions`. It will also get logged to the `query.log`.

A timeout will cause the related transaction to be terminated by the database if its execution time is longer than the provided value. This functionality allows to limit query/transaction execution time. A specified timeout overrides the default timeout defined by the configuration setting `config_dbms.transaction.timeout`.

For usage patterns, please refer to the [API documentation](#) for your specific language.

Cypher errors

When executing Cypher, it is possible for an exception to be thrown by the Cypher engine. Each such exception is associated with a `status code` that describes the nature of the error and a message that provides more detail.

The error classifications are listed in the table below.

Table 4. Error classifications

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

Causal chaining

When working with a Causal Cluster, transactions can be chained 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.

Causal chaining is carried out by passing **bookmarks** between transactions. Each bookmark records a point in transactional history and can be used to inform cluster members to carry out units of work in a particular sequence. Internally, a bookmark is passed from server to client on successful completion of a transaction and back from client to server on start of a new transaction, even if the transaction is an **auto-commit transaction**. On receipt of one or more bookmarks, the transaction will block until the server has fast forwarded to catch up with the latest of these.



If you use the 1.7 driver to connect to a Neo4j database that is version 3.4 or earlier, **auto-commit transactions** cannot take part in the causal chain. In that configuration, `session.run` calls should be avoided in places where causal consistency is important.

Within a session, bookmark propagation is carried out automatically 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. Propagation between sessions can be achieved by extracting the last bookmarks from one or more sessions and passing these into the construction of another. This is generally the only case in which an application will need to work with bookmarks directly.

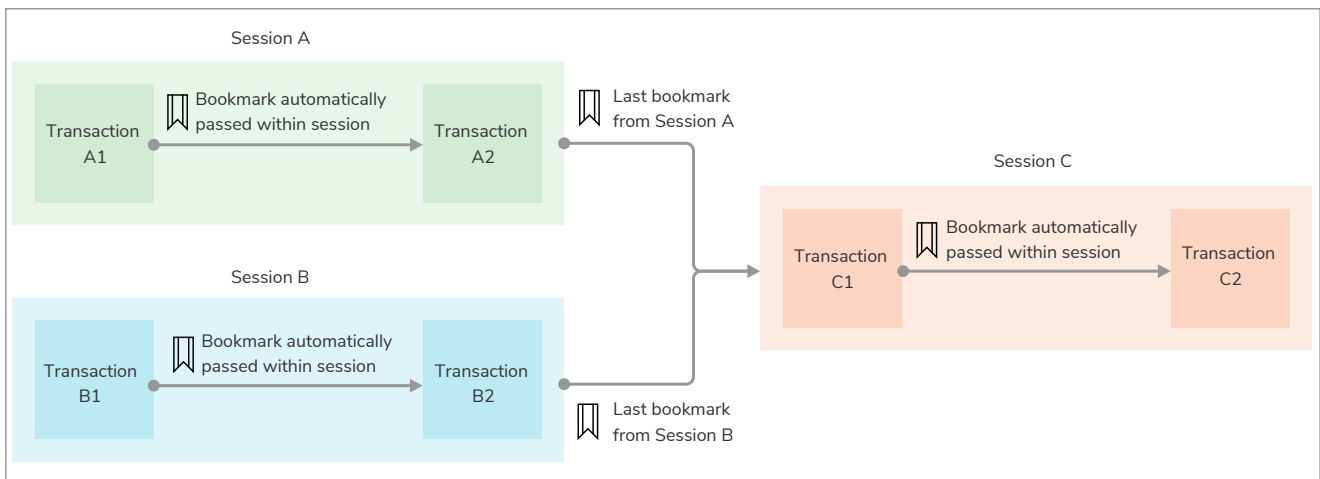


Figure 2. Passing bookmarks

The following example illustrates the passing of bookmarks between sessions. First consider the code:

Example 22. Passing bookmarks between sessions




```

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

// Create a person node
private void AddPerson(ITransaction tx, string name)
{
    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 void Employ(ITransaction tx, string personName, string companyName)
{
    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 void MakeFriends(ITransaction tx, string name1, string name2)
{
    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 void PrintFriendships(ITransaction tx)
{
    var result = tx.Run("MATCH (a)-[:KNOWS]->(b) RETURN a.name, b.name");

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

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

    // Create the first person and employment relationship.
    using (var session1 = Driver.Session(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(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(AccessMode.Write, savedBookmarks))
    {
        session3.WriteTransaction(tx => MakeFriends(tx, "Alice", "Bob"));

        session3.ReadTransaction(PrintFriendships);
    }
}

```

We are using three separate sessions: A, B and C. In Session A we run two separate transactions. In the first one we create the person `Alice`, and in the second one we record that she works at `Wayne Enterprises`. The bookmark being passed between the two transactions is handled by the session. The bookmark from the last transaction is saved into an array for future use.

In Session B we also run two separate transactions. In the first one we create the person `Bob`, and in the second one we record that he works at `LexCorp`. Again, the bookmark being passed between the two transactions is handled by the session. The bookmark from the last transaction is saved into an array for future use.

In the last session, Session C, we wish to create a friendship between Alice and Bob. This can only be done if both `Alice` and `Bob` have been created first. In order to ensure this, we pass the bookmarks from the last transactions in Session A and Session B, respectively.



If you try to extract a bookmark from a database which is not running in Causal Cluster mode, you will receive a `null` result.

Access modes

Transactions can be executed in either `read` or `write` 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: as read servers are typically more plentiful than write servers, it is beneficial to direct as much as possible of read transactions to read servers. Doing so helps keeping write servers available for write transactions.

Access modes can be supplied in two ways: per transaction or per session. An access mode specified at session creation can be overridden by the access mode of a transaction within that session. In the general case, access mode should always be specified at transaction level, using [transaction functions](#). The session-level setting is only necessary for explicit and auto-commit transactions.

Note that the driver does not parse Cypher and cannot determine whether a transaction is intended to carry out read or write operations. As a result of this, a `write` transaction tagged for `read` will be sent to a read server, but will fail on execution.

Example 23. 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 void CreatePersonNode(ITransaction tx, string name)
{
    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>();
}
```

Asynchronous programming



Java, .NET and JavaScript all support asynchronous programming. The examples here highlight specifically how Java and .NET provide for this programming model alongside their blocking API.

In addition to the methods listed in the previous sections, there also exist several asynchronous methods which allow for better integration with applications written in an asynchronous style. Asynchronous methods are named as their synchronous counterparts but with an additional `async` prefix.

Example 24. Asynchronous programming examples



Auto-commit transactions

```
var records = new List<string>();
var session = Driver.Session();

try
{
    // Send cypher statement to the database.
    // The existing IStatementResult interface implements IEnumerable
    // and does not play well with asynchronous use cases. The replacement
    // IStatementResultCursor 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 statement
        new { id = 0 } // Parameters in the statement, 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();
}
```

Transaction functions

```
var session = Driver.Session();

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

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

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

Explicit transactions

```
var records = new List<string>();
var session = Driver.Session();

try
{
    // Start an explicit transaction
    var tx = await session.BeginTransactionAsync();

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

    // Commit the transaction
    await tx.CommitAsync();
}
finally
{
    // asynchronously close session
    await session.CloseAsync();
}
```



It is always important to close the session object and it is suggested to keep `session.CloseAsync` in a *finally* block to make sure all resources (such as network connection) obtained by the session will always be cleaned up properly. The session close method also enforces rolling back of the last uncommitted or failed transaction in this session. Thus it is optional to put `tx.CommitAsync` or `tx.RollbackAsync` in a *finally* block as long as the close of session where the transaction is created will be executed in an outer *finally* block.

Working with Cypher values

This section describes the types and values used by Cypher and how they map to native language types.

The Cypher type system

Drivers translate between application language types and the Cypher type system. 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 can be potentially found in results although not all can be used as parameters.

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

* 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 driver maps 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 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.

Statement results

A statement result is comprised of a stream of records. The result is typically handled by the receiving application as it arrives, although it is also possible to retain all or part of a result for future consumption.

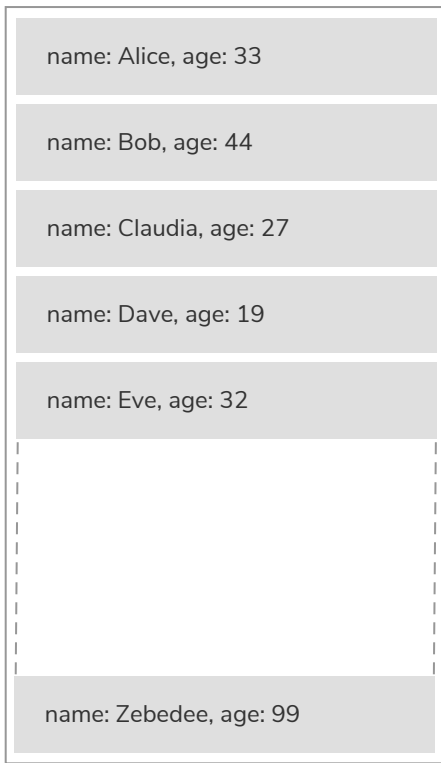


Figure 3. Result stream

Records

A record provides an immutable view of part of a result. It is an ordered map of keys and values. As values within a record are both keyed and ordered, that value can be accessed either by position (0-based integer) or by key (string).

The buffer

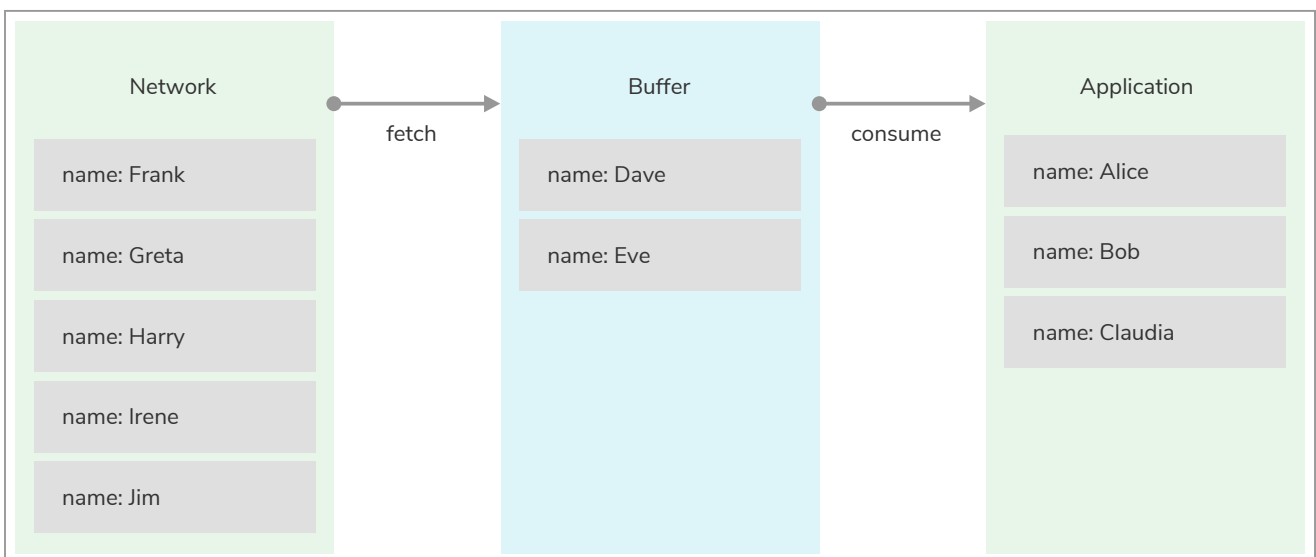


Figure 4. Result buffer

Most drivers contain a result buffer. This 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).

If results are consumed in the same order as they are produced, records merely pass through the buffer; if they are consumed out of order, the buffer will be utilized to retain records until they are consumed by the application. For large results, this may require a significant amount of memory and impact performance. For this reason, it is recommended to consume results in order wherever possible.

Consuming the stream

Query results will often be consumed as a stream. Drivers provide a language-idiomatic way to iterate through a result stream.

Example 26. Consuming the stream

```
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

The result record stream is available until another statement is run in the session, or until the current transaction is closed. To hold on to the results beyond this scope, the results need to be explicitly retained. For retaining results, each driver offers methods that collect the result stream and translate it into standard data structures for that language. Retained results can be processed while the session is free to take on the next workload. The saved results can also be used directly to run a new statement.

Example 27. Retain results for further processing

```
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;
        }));
    }
}
```

Statement result summaries

Supplementary information such as query statistics, timings and server information can be obtained from the statement result summary. If this detail is accessed before the entire result has been consumed, the remainder of the result will be buffered.

See also the [language-specific driver API documentation](#).

Appendix A: Driver terminology

This section lists the relevant terminology related to Neo4j drivers.

acquire (connection)

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

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.

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.

direct driver

A driver that can connect to a single server address.

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.

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.

statement result

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

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.

License

Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0)

You are free to

Share

copy and redistribute the material in any medium or format

Adapt

remix, transform, and build upon the material

The licensor cannot revoke these freedoms as long as you follow the license terms.

Under the following terms

Attribution

You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

NonCommercial

You may not use the material for commercial purposes.

ShareAlike

If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.

No additional restrictions

You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Notices

You do not have to comply with the license for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation.

No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material.

See <https://creativecommons.org/licenses/by-nc-sa/4.0/> for further details. The full license text is available at <https://creativecommons.org/licenses/by-nc-sa/4.0/legalcode>.