

Para entender los índices vamos a practicar sobre una base de datos ya creada con numerosas entradas. Para ello nos clonamos el repositorio `mongoDB_seeder` de mi github ([https://github.com/pabSec/mongoDB\\_seeder](https://github.com/pabSec/mongoDB_seeder)) en un contenedor o máquina con `node.js` (`docker run -ti -d node:latest`). Este paso podemos saltarlo y simplemente restaurar la base de datos log completa que se encuentra en el mismo repositorio (`mongorestore`).

Instalamos los paquetes `mongoose` y `faker` para que el script `log.js` se ejecute correctamente (`npm install faker` y `npm install mongoose`).

```
> db.logs.count()
40000
Tenemos un total de 40000 documentos con datos relativos a un log de un servidor de este estilo:
> db.logs.findOne()
{
  "_id" : ObjectId("58a358ab700eee0079c8e543"),
  "host" : "2.20.64.134",
  "user" : "Sandra_Glover",
  "time" : ISODate("2017-02-14T19:21:15.207Z"),
  "path" : "https://www.forocoches.com",
  "request" : "POST",
  "status" : "404",
  "response_size" : 9677,
  "user_agent" : "IE",
  "__v" : 0
}
```

Vamos a comprobar en primer lugar cuanto tarda en ejecutar varios tipos de búsqueda para hacernos a la idea del rendimiento de `mongoDB` con y sin índices definidos:

```
> db.logs.find({path:"https://www.forocoches.com"}).explain("executionStats")
{
  "queryPlanner" : {
    "plannerVersion" : 1,
    "namespace" : "log.logs",
    "indexFilterSet" : false,
    "parsedQuery" : {
      "path" : {
        "$eq" : "https://www.forocoches.com"
      }
    },
    "winningPlan" : {
      "stage" : "COLLSCAN",
      "filter" : {
        "path" : {
          "$eq" : "https://www.forocoches.com"
        }
      }
    },
    "direction" : "forward"
  },
  "rejectedPlans" : [ ]
}
```

```

    },
    "executionStats" : {
      "executionSuccess" : true,
      "nReturned" : 5655,
      "executionTimeMillis" : 42,
      "totalKeysExamined" : 0,
      "totalDocsExamined" : 40000,
      "executionStages" : {
        "stage" : "COLLSCAN",
        "filter" : {
          "path" : {
            "$eq" : "https://www.forocoches.com"
          }
        },
        "nReturned" : 5655,
        "executionTimeMillisEstimate" : 50,
        "works" : 40002,
        "advanced" : 5655,
        "needTime" : 34346,
        "needYield" : 0,
        "saveState" : 312,
        "restoreState" : 312,
        "isEOF" : 1,
        "invalidates" : 0,
        "direction" : "forward",
        "docsExamined" : 40000
      }
    },
    "serverInfo" : {
      "host" : "7ea9850856d3",
      "port" : 27017,
      "version" : "3.4.2",
      "gitVersion" : "3f76e40c105fc223b3e5aac3e20dcd026b83b38b"
    },
    "ok" : 1
  }
}
> db.logs.find({user:"Sandra_Glover"}).hint({"$natural":1}).explain("executionStats")
{
  "queryPlanner" : {
    "plannerVersion" : 1,
    "namespace" : "log.logs",
    "indexFilterSet" : false,
    "parsedQuery" : {
      "user" : {
        "$eq" : "Sandra_Glover"
      }
    }
  },
  "winningPlan" : {
    "stage" : "COLLSCAN",
    "filter" : {
      "user" : {
        "$eq" : "Sandra_Glover"
      }
    }
  },
  "direction" : "forward"
},
"rejectedPlans" : [ ]
}

```

```

    },
    "executionStats" : {
      "executionSuccess" : true,
      "nReturned" : 337,
      "executionTimeMillis" : 16,
      "totalKeysExamined" : 0,
      "totalDocsExamined" : 40000,
      "executionStages" : {
        "stage" : "COLLSCAN",
        "filter" : {
          "user" : {
            "$eq" : "Sandra_Glover"
          }
        },
        "nReturned" : 337,
        "executionTimeMillisEstimate" : 10,
        "works" : 40002,
        "advanced" : 337,
        "needTime" : 39664,
        "needYield" : 0,
        "saveState" : 312,
        "restoreState" : 312,
        "isEOF" : 1,
        "invalidates" : 0,
        "direction" : "forward",
        "docsExamined" : 40000
      }
    },
    "serverInfo" : {
      "host" : "7ea9850856d3",
      "port" : 27017,
      "version" : "3.4.2",
      "gitVersion" : "3f76e40c105fc223b3e5aac3e20dcd026b83b38b"
    },
    "ok" : 1
  }
}

```

Y a continuación jugaremos con los índices para comprobar su correcto funcionamiento. Supongamos que queremos realizar una aplicación cuyo panel de control nos permita visualizar los datos para cada path así como los datos de navegación para cada usuario.

La lógica nos dice que deberíamos probar con índices simples para cada uno de esos campos así como un índice compuesto para cada uno. A continuación probaremos con esas combinaciones y comprobaremos cuál resulta más eficiente.

```

> db.logs.createIndex({user:1})
{
  "createdCollectionAutomatically" : false,
  "numIndexesBefore" : 1,
  "numIndexesAfter" : 2,
  "ok" : 1
}
> db.logs.createIndex({path:1})

```

```

{
  "createdCollectionAutomatically" : false,
  "numIndexesBefore" : 2,
  "numIndexesAfter" : 3,
  "ok" : 1
}
> db.logs.createIndex({user:1,path:1})
{
  "createdCollectionAutomatically" : false,
  "numIndexesBefore" : 3,
  "numIndexesAfter" : 4,
  "ok" : 1
}
> db.logs.createIndex({path:1,user:1})
{
  "createdCollectionAutomatically" : false,
  "numIndexesBefore" : 4,
  "numIndexesAfter" : 5,
  "ok" : 1
}

```

Y vamos a realizar las queries anteriores para ver qué índices utiliza mongo a la hora de buscar por defecto.

```

> db.logs.find({user:"Sandra_Glover"}).hint({user:1}).explain("executionStats")
{
  "queryPlanner" : {
    "plannerVersion" : 1,
    "namespace" : "log.logs",
    "indexFilterSet" : false,
    "parsedQuery" : {
      "user" : {
        "$eq" : "Sandra_Glover"
      }
    },
    "winningPlan" : {
      "stage" : "FETCH",
      "inputStage" : {
        "stage" : "IXSCAN",
        "keyPattern" : {
          "user" : 1
        },
        "indexName" : "user_1",
        "isMultiKey" : false,
        "multiKeyPaths" : {
          "user" : [ ]
        },
        "isUnique" : false,
        "isSparse" : false,
        "isPartial" : false,
        "indexVersion" : 2,
        "direction" : "forward",
        "indexBounds" : {
          "user" : [
            ["Sandra_Glover", "Sandra_Glover"]

```

```

        ]
    }
}
},
"rejectedPlans" : [ ]
},
"executionStats" : {
    "executionSuccess" : true,
    "nReturned" : 337,
    "executionTimeMillis" : 0,
    "totalKeysExamined" : 337,
    "totalDocsExamined" : 337,
    "executionStages" : {
        "stage" : "FETCH",
        "nReturned" : 337,
        "executionTimeMillisEstimate" : 0,
        "works" : 338,
        "advanced" : 337,
        "needTime" : 0,
        "needYield" : 0,
        "saveState" : 2,
        "restoreState" : 2,
        "isEOF" : 1,
        "invalidates" : 0,
        "docsExamined" : 337,
        "alreadyHasObj" : 0,
        "inputStage" : {
            "stage" : "IXSCAN",
            "nReturned" : 337,
            "executionTimeMillisEstimate" : 0,
            "works" : 338,
            "advanced" : 337,
            "needTime" : 0,
            "needYield" : 0,
            "saveState" : 2,
            "restoreState" : 2,
            "isEOF" : 1,
            "invalidates" : 0,
            "keyPattern" : {
                "user" : 1
            },
            "indexName" : "user_1",
            "isMultiKey" : false,
            "multiKeyPaths" : {
                "user" : [ ]
            },
            "isUnique" : false,
            "isSparse" : false,
            "isPartial" : false,
            "indexVersion" : 2,
            "direction" : "forward",
            "indexBounds" : {
                "user" : [
                    ["\Sandra_Glover\","Sandra_Glover\"]]
                ]
            },
            "keysExamined" : 337,
            "seeks" : 1,

```

```

        "dupsTested" : 0,
        "dupsDropped" : 0,
        "seenInvalidated" : 0
      }
    }
  },
  "serverInfo" : {
    "host" : "7ea9850856d3",
    "port" : 27017,
    "version" : "3.4.2",
    "gitVersion" : "3f76e40c105fc223b3e5aac3e20dcd026b83b38b"
  },
  "ok" : 1
}

```

Para una base de datos de 40000 documentos, definiendo un índice simple, reducimos la operación de 16 milisegundos a 0. Del mismo modo, para el path “<https://www.forocoches.com>” reducimos la query de 42 milisegundos a 6:

```

> db.logs.find({path:"https://www.forocoches.com"}).hint({path:1}).explain("executionStats")
{
  "queryPlanner" : {
    "plannerVersion" : 1,
    "namespace" : "log.logs",
    "indexFilterSet" : false,
    "parsedQuery" : {
      "path" : {
        "$eq" : "https://www.forocoches.com"
      }
    },
    "winningPlan" : {
      "stage" : "FETCH",
      "inputStage" : {
        "stage" : "IXSCAN",
        "keyPattern" : {
          "path" : 1
        },
        "indexName" : "path_1",
        "isMultiKey" : false,
        "multiKeyPaths" : {
          "path" : [ ]
        },
        "isUnique" : false,
        "isSparse" : false,
        "isPartial" : false,
        "indexVersion" : 2,
        "direction" : "forward",
        "indexBounds" : {
          "path" : [
            ["https://www.forocoches.com", "https://ww
w.forocoches.com\"]]
          ]
        }
      },
      "rejectedPlans" : [ ]
    }
  },
  "rejectedPlans" : [ ]
}

```

```

    },
    "executionStats" : {
      "executionSuccess" : true,
      "nReturned" : 5655,
      "executionTimeMillis" : 6,
      "totalKeysExamined" : 5655,
      "totalDocsExamined" : 5655,
      "executionStages" : {
        "stage" : "FETCH",
        "nReturned" : 5655,
        "executionTimeMillisEstimate" : 10,
        "works" : 5656,
        "advanced" : 5655,
        "needTime" : 0,
        "needYield" : 0,
        "saveState" : 44,
        "restoreState" : 44,
        "isEOF" : 1,
        "invalidates" : 0,
        "docsExamined" : 5655,
        "alreadyHasObj" : 0,
        "inputStage" : {
          "stage" : "IXSCAN",
          "nReturned" : 5655,
          "executionTimeMillisEstimate" : 10,
          "works" : 5656,
          "advanced" : 5655,
          "needTime" : 0,
          "needYield" : 0,
          "saveState" : 44,
          "restoreState" : 44,
          "isEOF" : 1,
          "invalidates" : 0,
          "keyPattern" : {
            "path" : 1
          },
          "indexName" : "path_1",
          "isMultiKey" : false,
          "multiKeyPaths" : {
            "path" : [ ]
          },
          "isUnique" : false,
          "isSparse" : false,
          "isPartial" : false,
          "indexVersion" : 2,
          "direction" : "forward",
          "indexBounds" : {
            "path" : [
              ["https://www.forocoche
w.forocoche
.com\""],
              ]
            }
          },
          "keysExamined" : 5655,
          "seeks" : 1,
          "dupsTested" : 0,
          "dupsDropped" : 0,
          "seenInvalidated" : 0
        }
      }
    }
  }
}

```

```
    },  
    },  
    "serverInfo" : {  
      "host" : "7ea9850856d3",  
      "port" : 27017,  
      "version" : "3.4.2",  
      "gitVersion" : "3f76e40c105fc223b3e5aac3e20dcd026b83b38b"  
    },  
    "ok" : 1  
  }  
}
```