



Tackling Your Data Challenges

Getting Started with Apache Kafka, Splunk, and GraphQL

Solving your data challenges

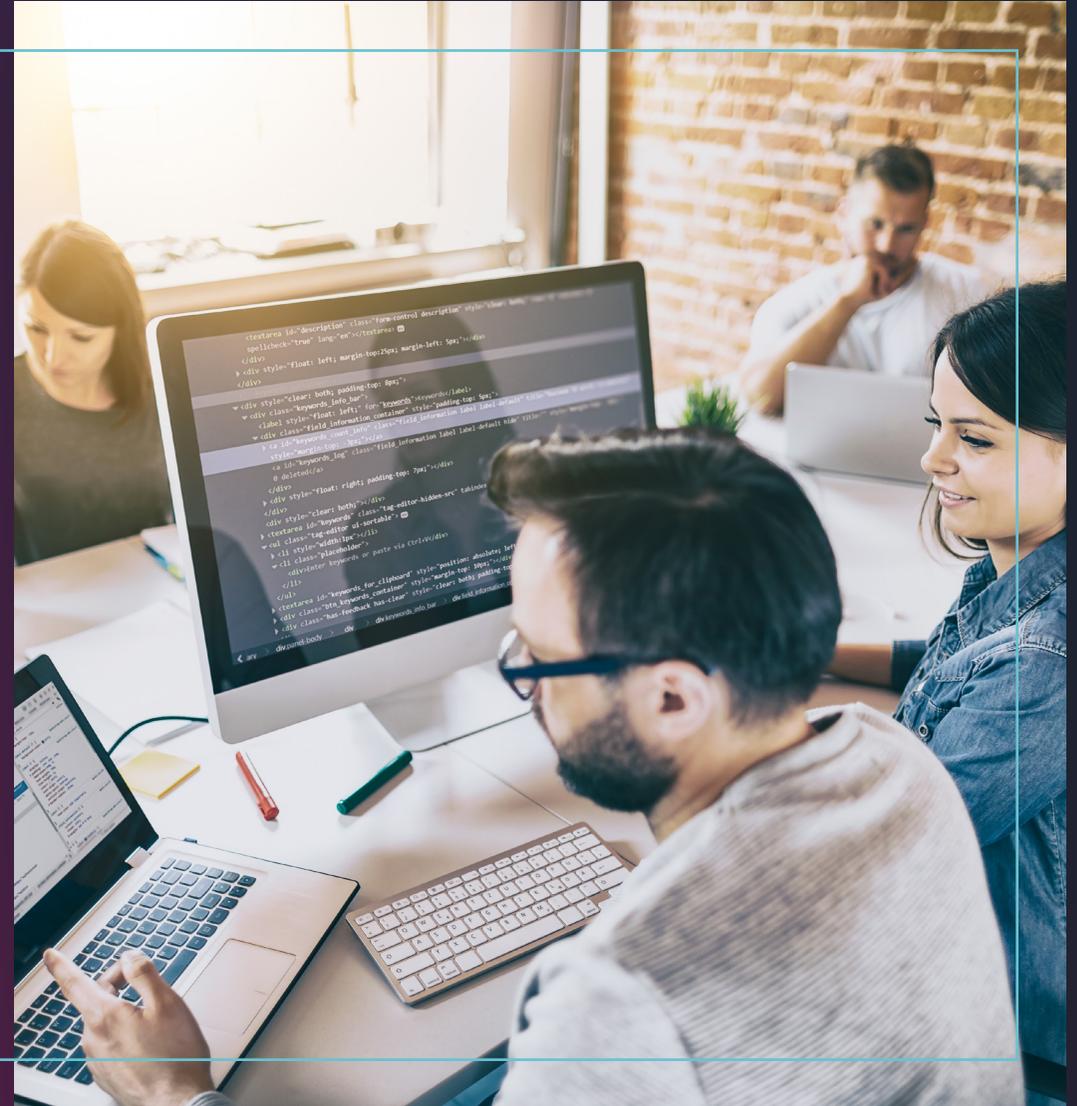
In our era of big data, your IT infrastructure faces many new demands from the need to display real-time data feeds to transferring and loading data quickly from your database to mobile devices. At the same time, your IT infrastructure collects all kinds of machine-generated data that could be useful to help you prevent the next cybersecurity attack or system meltdown. Turning all this unstructured data into predictive insights can help your IT team become more proactive, rather than reactive.

Here are 3 emerging IT technologies that can help solve your data issues and digitally transform your organization. These **3 hot tech skills** are trending on **Udemy**—a global online learning marketplace with over 80,000 courses, 24+ million learners, and thousands of businesses.

- 1 **Apache Kafka**
- 2 **Splunk**
- 3 **GraphQL**

In this guide, IT pros who teach tech courses on **Udemy for Business**—the learning platform for organizations around the world—offer advice on how to get started with Apache Kafka, Splunk, and GraphQL:

- **How these 3 technologies solve key data issues from real-time data feeds to relational data loading**
- **The benefits over alternatives**
- **Steps to deploy them at your organization**
- **The skills your IT or Developer team will need to know**





Apache Kafka



By Stephane Maarek

Udemy for Business instructor and
CEO of DataCumulus

How Apache Kafka solves your real-time data issues

In our era of big data, your IT infrastructure may be taxed by the influx of data from a wide variety of sources. On top of that, customers demand to see their data in “real time” without any lag time, so your servers need to process and display data quickly. Apache Kafka is a new technology, developed in 2011, that allows you to do just that.

Apache Kafka is a distributed streaming platform that enables companies to create real-time data feeds. Apache Kafka is used by companies like Uber, Twitter, Airbnb, Yelp and over 30% of today's Fortune 500 companies. For example, by integrating diverse kinds of data such as likes, page clicks, searches, orders, shopping carts, and inventory, Apache Kafka can help feed data in real time into a predictive analytics engine to help you analyze customer behavior.



Apache Kafka

Now that Apache Kafka has reached a stable 2.0 version, more companies are adopting the technology as the backbone of their IT infrastructure. Increasingly, CTOs are prioritizing more real-time architecture and reducing the wait time on data availability. Apache Kafka is the **#1 hottest IT skill** trending on **Udemy** in 2018. Apache Kafka-related questions on Google Search and tech forums like StackOverflow and Github have also skyrocketed in recent years—all signaling that it's a trending hot topic.

So what are the benefits of Apache Kafka, why should your company adopt it, and what skills will your IT and engineering team need to successfully implement it?

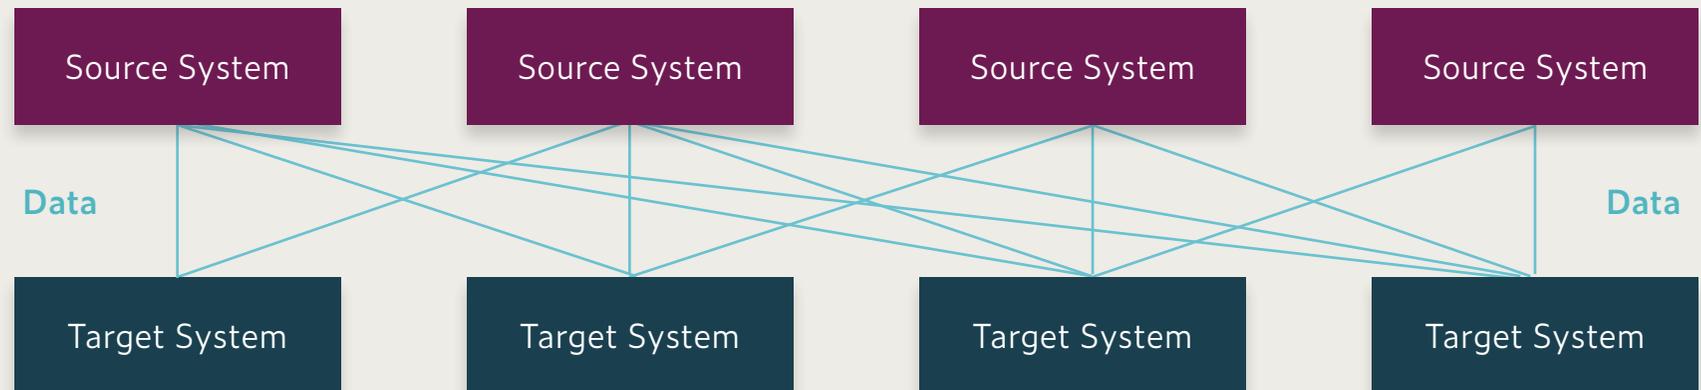


Apache Kafka is the **#1 hottest IT skill** trending on Udemy in 2018



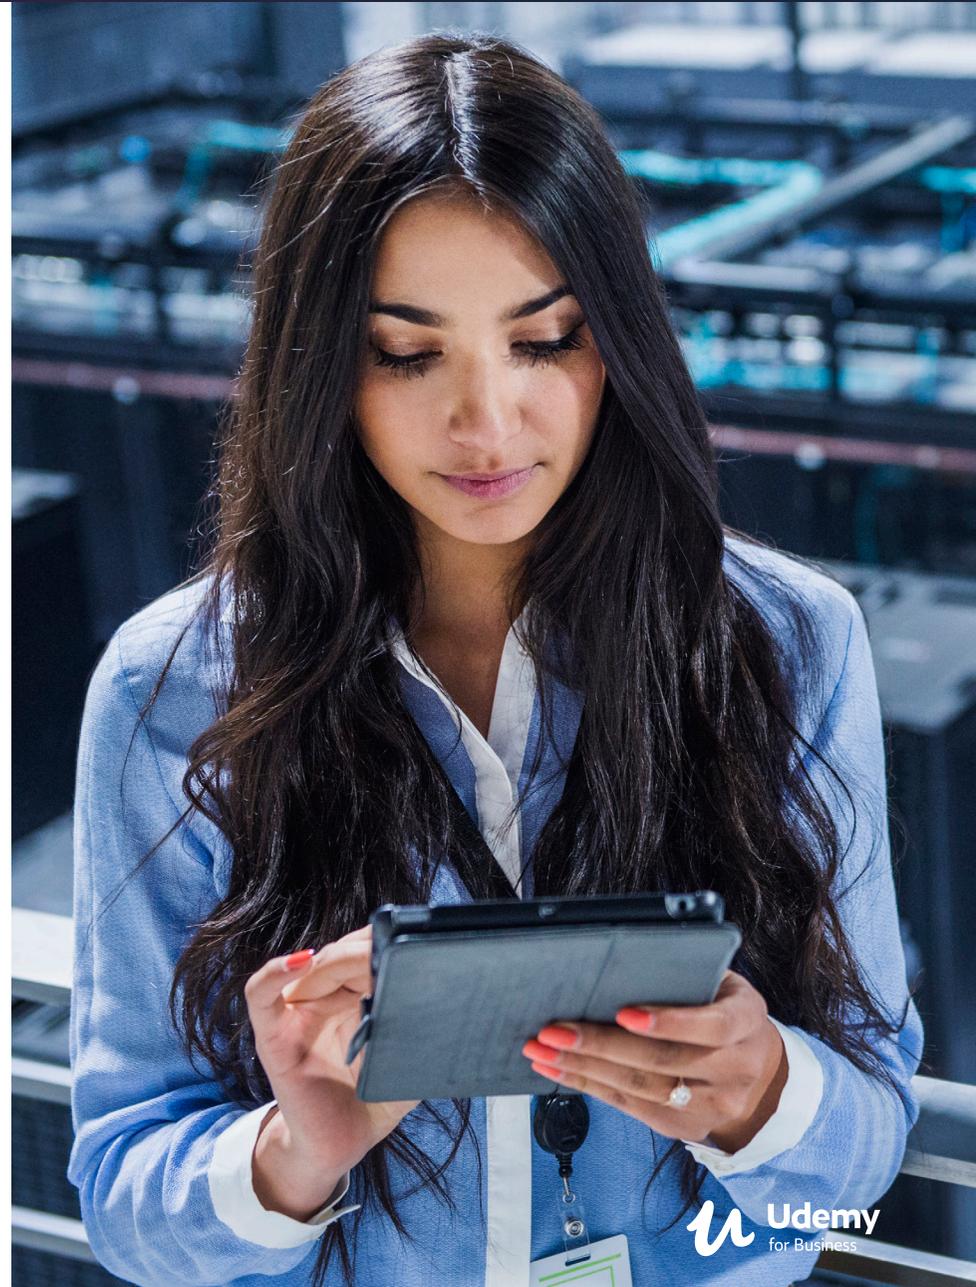
Benefits of Apache Kafka

As companies deliver an increasing amount of data from different sources (e.g. website, user interactions, financial transactions) to a wide range of target systems (e.g. databases, analytics, email systems), developers have to write integrations for each one. For example, if you have 4 source systems and 6 target systems, your developers would have to write code for 24 integrations. This is a cumbersome process, not to mention, a slow and error-prone way to deliver data. Here are four ways that Apache Kafka solves a lot of these issues.



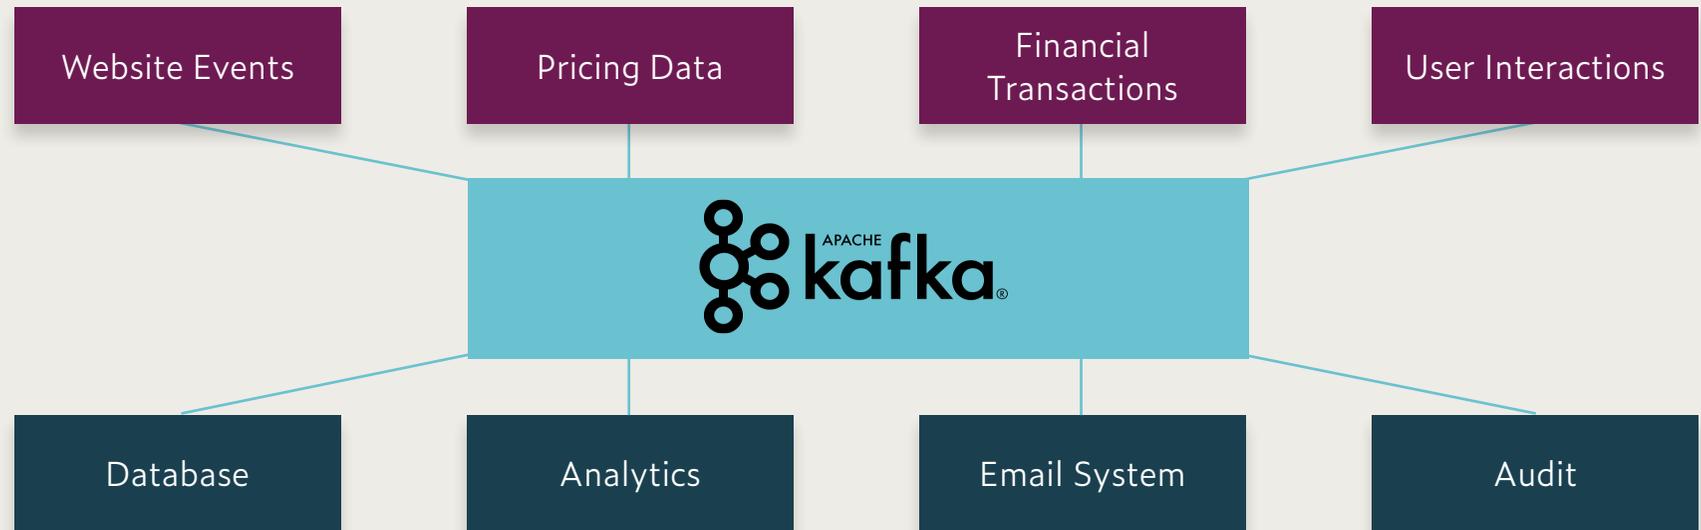
1 | Apache Kafka acts as a buffer so your systems won't crash

Previously, data transformations from external source systems were done in batches often at night. Apache Kafka solves this slow, multi-step process by acting as an intermediary receiving data from source systems and then making this data available to target systems in real time. What's more, your systems won't crash because Apache Kafka is its own separate set of servers (called an Apache Kafka cluster).



2 | Apache Kafka reduces the need for multiple integrations

Essentially, Apache Kafka reduces the need for multiple integrations—as all your data goes through Apache Kafka. Rather than your developers coding multiple integrations so you can harvest data from different systems, you only have to create one integration with Apache Kafka for each producing system and each consuming system (see diagram below).





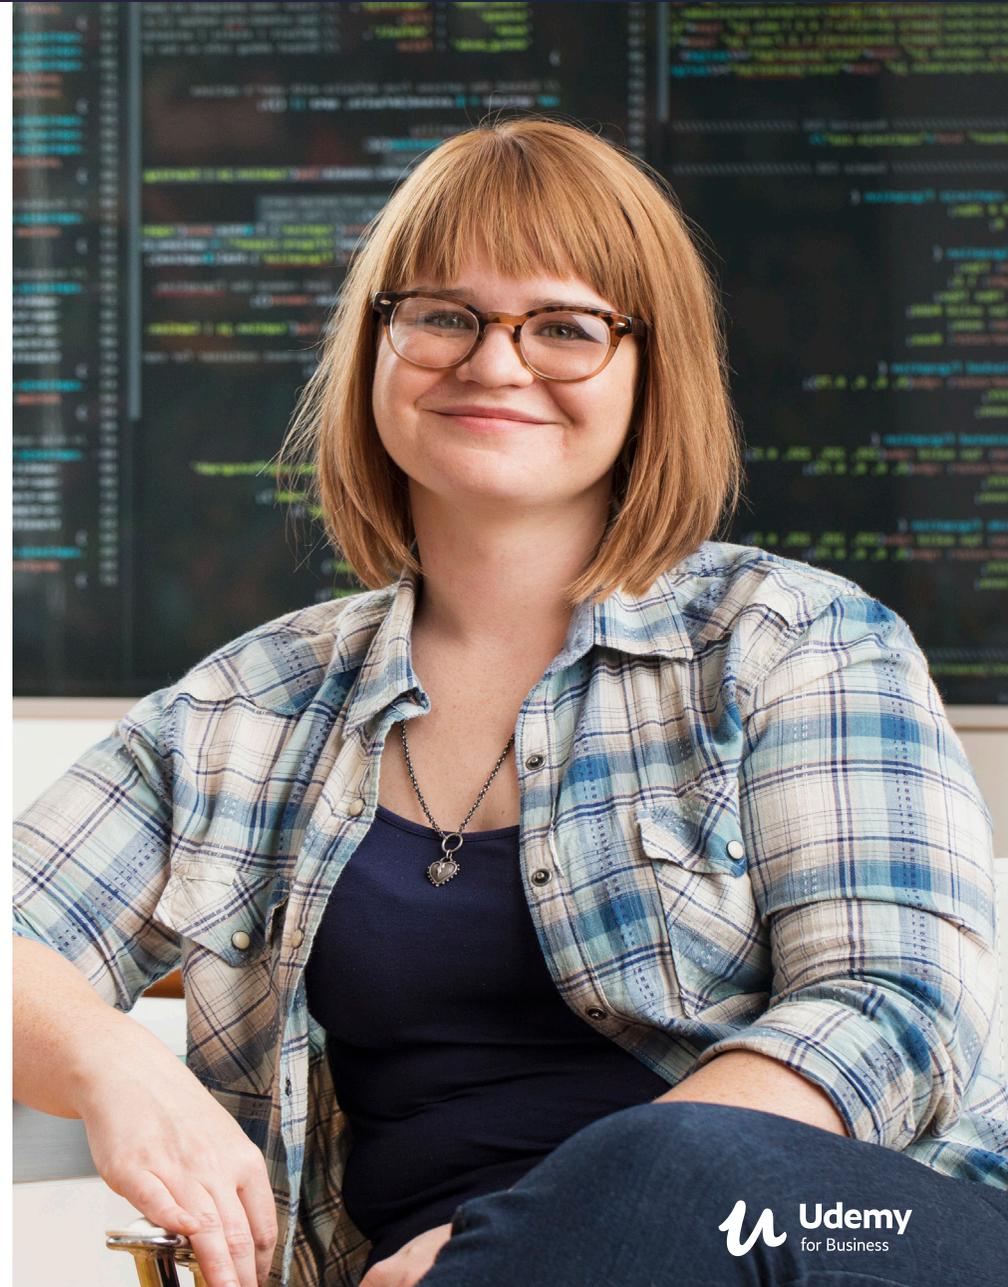
3 | Apache Kafka has low latency and high throughput

By decoupling your data streams, Apache Kafka lets you consume data when you want it. Without the need for slow integrations, Apache Kafka decreases latency (or how long it takes for each data point to load) to a mere 10 milliseconds (~10x decrease or more compared to other integrations). This means you can deliver data quickly and in real time. Apache Kafka can also horizontally scale to hundreds of brokers (or servers) within a cluster to manage big data.

Some companies have a high load of millions of data points per second going through Kafka. For example, Uber uses Kafka to feed car position data into their surge pricing computation model in real time.

4 | Everyone can access data with Apache Kafka

Apache Kafka centralizes your data so that, access to data for any team becomes easier. For example, in the past, your fraud team may have had to engage with the web team to get a specific type of user data--since they were run on different target systems. Now your fraud team will be able to access the user data directly via Apache Kafka, alongside other feeds such as financial data or website interactions. Simple, right?



How to deploy Apache Kafka

Apache Kafka skills your IT team needs

Once you understand the benefits and decide to adopt Apache Kafka at your organization, your IT team will need to acquire key skills to set up and manage it.

How to learn, set up, and configure Apache Kafka. Apache Kafka is already built, open source, and free. So it's more about first learning about it, then setting it up and configuring it for your systems. A good place to start is my Udemy for Business course on [Apache Kafka Series: Learn Apache Kafka for Beginners](#). The course covers the Apache Kafka ecosystem, how some target architectures can look, as well as fundamental concepts of Kafka like topics, partitions, replication, brokers, producers, consumer groups, Zookeeper, delivery semantics, and more. The course also offers hands-on practice so your team can gain some practical experience using Apache Kafka.

Recommended course



Apache Kafka Series - Learn Apache Kafka for Beginners

Stephane Maarek

★★★★★ 4.5 (3,793)

[Apache Kafka Series: Learn Apache Kafka for Beginners](#)

Apache Kafka

Once you're ready, I recommend a more advanced Udemy for Business course called [Kafka Cluster Setup and Administration](#).

Java Programming. While clients libraries exist to interact with Apache Kafka using most programming languages (Python, Go, Javascript, etc.), Apache Kafka works best when using a Java Virtual Machine type of language, such as Java or Scala. Therefore it is important for your team to be able to **code in Java** when writing their integrations.

Kafka Streams and Kafka Connect. If you want to simplify integrations, your team will also need some Kafka-specific skills like Kafka Streams and Kafka Connect. These are the more advanced Kafka concepts and frameworks your team will need to build reliable and production ready integrations over time. As a consultant, I usually show how to build one or

Recommended courses



Apache Kafka Series: Kafka Cluster Setup & Administration
Stephane Maarek
★★★★★ 4.6 (592)

Complete Java Masterclass - updated for Java 10
Tim Buchalka, Tim Buchalka's Lear...
★★★★★ 4.6 (50,047)

[Apache Kafka Series: Kafka Cluster Setup and Administration](#)

[Complete Java Masterclass - updated for Java 10](#)

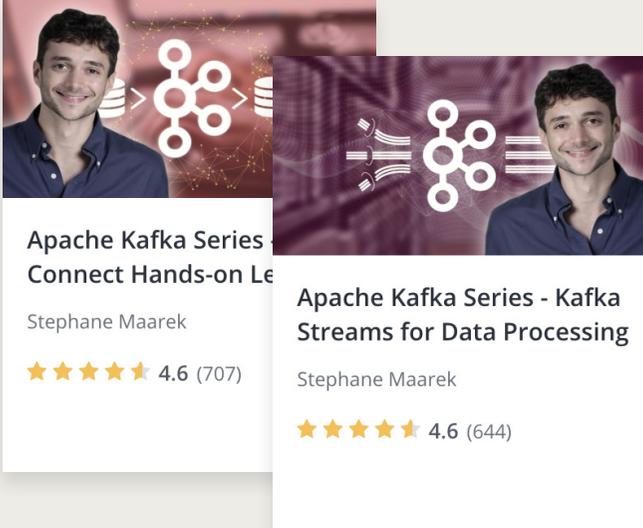
Apache Kafka

two integrations, but your team would have to scale that for the rest of the integrations.

Kafka Connect is a tool for scalable and reliable streaming data between Apache Kafka and other data systems. You can already leverage tons of existing connectors written for you at: confluent.io/product/connectors. Udemy for Business course [Kafka Connect](#) teaches you all the skills you will need to implement and leverage these connectors.

The Kafka Streams library is used to process, aggregate, and transform your data within Kafka. The course [Kafka Streams for Data Processing](#) teaches how to use this data processing library on Apache Kafka, through several examples that demonstrate the range of possibilities.

Recommended courses



Apache Kafka Series - Kafka Connect Hands-on Learning
Stephane Maarek
★★★★★ 4.6 (707)

Apache Kafka Series - Kafka Streams for Data Processing
Stephane Maarek
★★★★★ 4.6 (644)

[Apache Kafka Series: Kafka Connect Hands-on Learning](#)

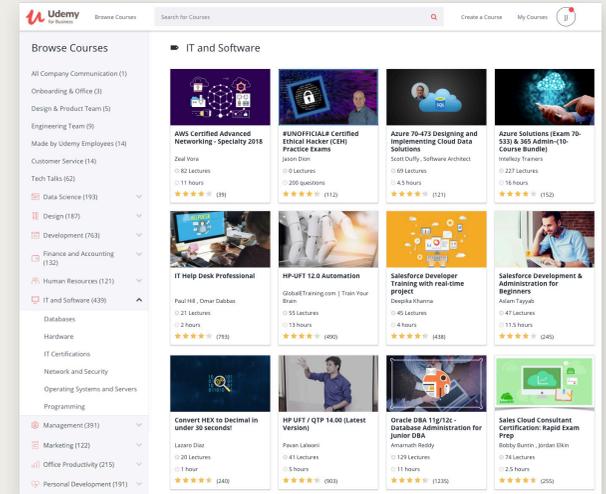
[Apache Kafka Series: Kafka Streams for Data Processing](#)

Migrating to Apache Kafka: start small

Don't migrate your whole system to Apache Kafka at once. Instead, start with a small non-critical project. For example, don't change the backbone of your financial systems, but change something less important such as your email notification system. Second, one of the biggest mistakes I see is companies spend months trying to build a reliable Apache Kafka cluster. Instead, I would recommend starting with managed services or hiring a consultant to set up a small project on Apache Kafka. This enables you to get started right away.



Stephane Maarek is a Udemy for Business instructor and CEO of Data Cumulus. As an experienced Apache Kafka consultant, he has helped companies transform their business through the appropriate use of real-time data. He has a Master's degree in Engineering from Cornell University.



Upskill your IT Team on Apache Kafka with Udemy for Business.

REQUEST DEMO



Splunk



By Adam Frisbee

Udemy for Business and
University of Utah instructor

How Splunk can help IT turn machine data into insights

Machines generate a lot of data. A single Linux or Windows server alone generates more logs than anyone has the time to examine manually. But within that data lies important insights and action items that can save your organization from the next cyber attack, prevent a mission-critical system from crashing, make predictions about future capacity needs, or increase sales by understanding user behavior on your website.

Sadly, much of the time this valuable lode goes unmined. It is only after a serious event that IT pros feverishly search through this data in an attempt to figure out what exactly happened.



Splunk is the #8 hottest IT skill trending on Udemy in 2018



Splunk makes mining this data easy and automatic. By using its powerful engine, intuitive user interface, and robust query language, IT professionals can create reports, dashboards, and alerts based on data generated from machines, databases, IoT sensors, HTTP outputs, syslogs, or almost any other source. Instead of being reactive, IT pros can now be proactive and even predictive. Splunk makes machine data work for IT pros, and in turn, is helping make IT a better, more capable business function.

How data moves through Splunk



Source data

WMI	Sensors
Performance	Logs
AD	APIs
SNMP	Files/directories
TCP/UDP	WMI
Metrics	HTTP/S
Behaviors	Scripts



Forwarders

- Parsers
- Processing
- Metadata
- Compression
- Security



Indexers

- Event processing
- Transforming
- Timestamping
- Storing



Search heads

- Data navigation
- Reporting/alerting/dashboarding
- User interaction
- Prediction
- Data mining
- Machine learning



Benefits of Splunk

Splunk is a veritable “Swiss Army knife” of the IT world. Splunk can ingest data from nearly any source imaginable. (I haven’t encountered a data source that Splunk cannot consume, but it might exist.)

1 | Splunk can capture and analyze unstructured data

Structured data from relational databases, Hadoop clusters, CSV files and the like are straightforward and easily consumed by Splunk. But any common business intelligence tool can capture and analyze structured data (I’m looking at you, *Tableau*).

Where Splunk excels is in the capturing and analyzing of unstructured, machine-generated data—the type of data that is immensely important to IT professionals. This data comes from many sources, and much of it from logs generated by information systems: servers, hypervisors, containers, firewalls, routers, sensors, databases, files/directories, and many more.

2 | Splunk can use machine learning to look for patterns in the data

Machines are great at documenting everything that happens to them or by them. They are not great at generating actionable knowledge from that data, or even making that data readable to humans.

Splunk can use machine learning to look for patterns in the data. This means that Splunk can help IT pros predict an imminent cyber attack, alert you when an action will make the disk reach its maximum storage capacity, or predict when a backup will fail. It can also, of course, help you make standard business intelligence predictions based on data from databases and other business sources.





3 | Splunk has an active user community resulting in a wealth of add-ons and apps

Part of the appeal of Splunk is its highly active community, which includes Splunk employees, customers, and enthusiasts. I visited their trendy headquarters in San Francisco in 2017, and it felt like I walked into an episode of HBO's *Silicon Valley*. The headquarters had a quirky, laid-back feel, and the people were focused on making a great product and having fun while doing it.

Because of this attitude, Splunk has enabled a great and highly participative user community (a fanbase would perhaps be more accurate). The open-source extensibility of Splunk through APIs and SDKs means that thousands of add-ons and apps are available (most of them are free), with more being added frequently. These apps are created by vendors, users, and Splunk engineers. They extend the ability of Splunk by offering

pre-defined data models, reports, dashboards, and much more. Anyone can make an app, and therefore Splunk is much more than just its core products—it's an essential and evolving tool for any IT organization.

Why choose Splunk over alternatives?

Splunk is not a mere business intelligence (BI) tool like Tableau. BI tools rely on mostly structured data and particular types of data sources to derive business insights. With Splunk, IT professionals can combine BI functionality with the powerful insight of unstructured machine data to get a more holistic view of the organization. ***This mix of structured and unstructured data is, in my opinion, the future of data.*** Splunk can be the single data analytics tool or complement the existing BI tools, and like most products, Splunk is not the only solution that does this.





Splunk vs. Elastic Stack

A popular, open source alternative to Splunk is Elastic Stack (formerly ELK). Elastic Stack is really four open source products combined: Kibana, Elasticsearch, Beats, and Logstash. When these products are working together, they can replicate much of the functionality of Splunk, but Elastic Stack requires installation, configuration, and integration of a modular system, whereas Splunk bundles its core functionality as a complete package.

Splunk vs. Sumo Logic

Sumo Logic is another popular alternative and leader of the “unseat Splunk’s throne” club. Sumo Logic has the advantage of being born in the cloud. Splunk’s cloud-native offering is a fairly recent development. Sumo Logic also has the advantage of being less expensive than Splunk. Sumo Logic is still new, so the availability of plugins and apps do not compare with the size of Splunk’s “Splunkbase” (app store). In addition, as a newer startup, they seem to be struggling with their support options. With Splunk, IT pros can take advantage of Splunk Answers (A Stack Overflow-like Q&A site), and best-of-class premium support.

How to deploy Splunk

Get started with a trial of Splunk cloud

Planning and design can go a long way in making sure you have a successful and productive experience with Splunk. I recommend getting started with a trial of Splunk cloud and analyzing a small subset of data that represents the type of data for which you intend to use Splunk.

Consider your use case

The next step is to consider your use case, and plan the appropriate resources and architecture. Splunk has best practice deployment models depending on the size of the user-base and the volume of expected data. Since Splunk is high performing, each search head (primary interface machine) will need 16 CPU cores and 16GB RAM. A small business may only have one search head, but large firms can have dozens.



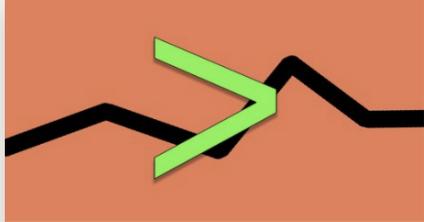
Splunk is not cheap. For a medium enterprise deployment, in addition to licensing costs, to realize ROI from Splunk, you will likely need the equivalent resources of a full-time employee. This is because of the changing nature of data—cyber-attacks, systems, log file formats, and other variables. The needs of the IT organization and business will change as well and require new dashboards, alerts, and reports.

The skills your IT team will need to know

Training is essential for any Splunk implementation. Splunk is one of those “easy to get started but challenging to master” products. It’s critical that your IT team learns the ins and outs of Splunk.

As a first step, my Udemy for Business course, [The Complete Splunk Beginner Course](#), can help get your IT team ready for a Splunk deployment. This course covers the basics of SPL, which is Splunk’s “Search Processing Language.” SPL is a basic query and manipulation language that is like a cross between SQL and

Recommended course



The Complete Splunk Beginner Course
Adam Frisbee
★★★★☆ 4.5 (1,575)

[The Complete Splunk Beginner Course](#)

Linux shell, though not nearly as complicated as either. You don't need to be a software developer to learn and understand SPL, and once you learn the basics, you will likely realize how simple and straightforward it is.

Architecture planning is important because there are many different ways to deploy a Splunk environment. In the course, we will deploy Splunk in a few different environments, analyze data from a provided dataset, build data models, design dashboards, and create reports and alerts. An **understanding of Linux is helpful**, even if you are working in a Windows environment. Splunk was written in C/C++ and uses many of the conventions of Linux (SPL uses the pipe | function frequently, for example).

Recommended course



Learn Linux in 5 Days and Level Up Your Career

Jason Cannon

★★★★☆ 4.5 (10,952)

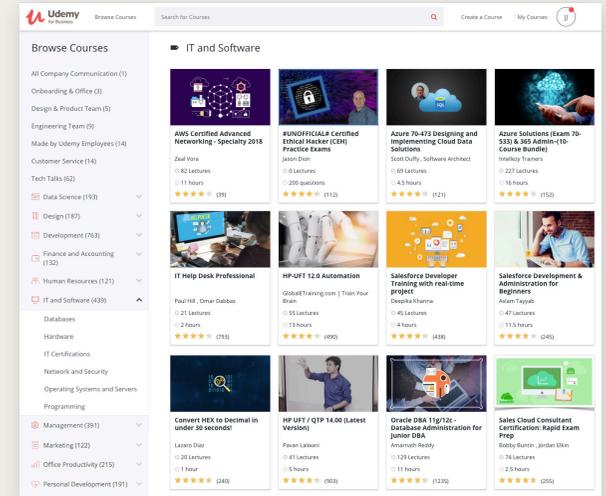
[Learn Linux in 5 Days and Level Up Your Career](#)

Becoming an IT hero with Splunk

Splunk is one of the most powerful and useful tools I have encountered in my travels, and I've been in IT for a long time. By planning ahead, instituting the proper training curriculum, and employing your background of basic IT infrastructure knowledge, you can become an IT hero.

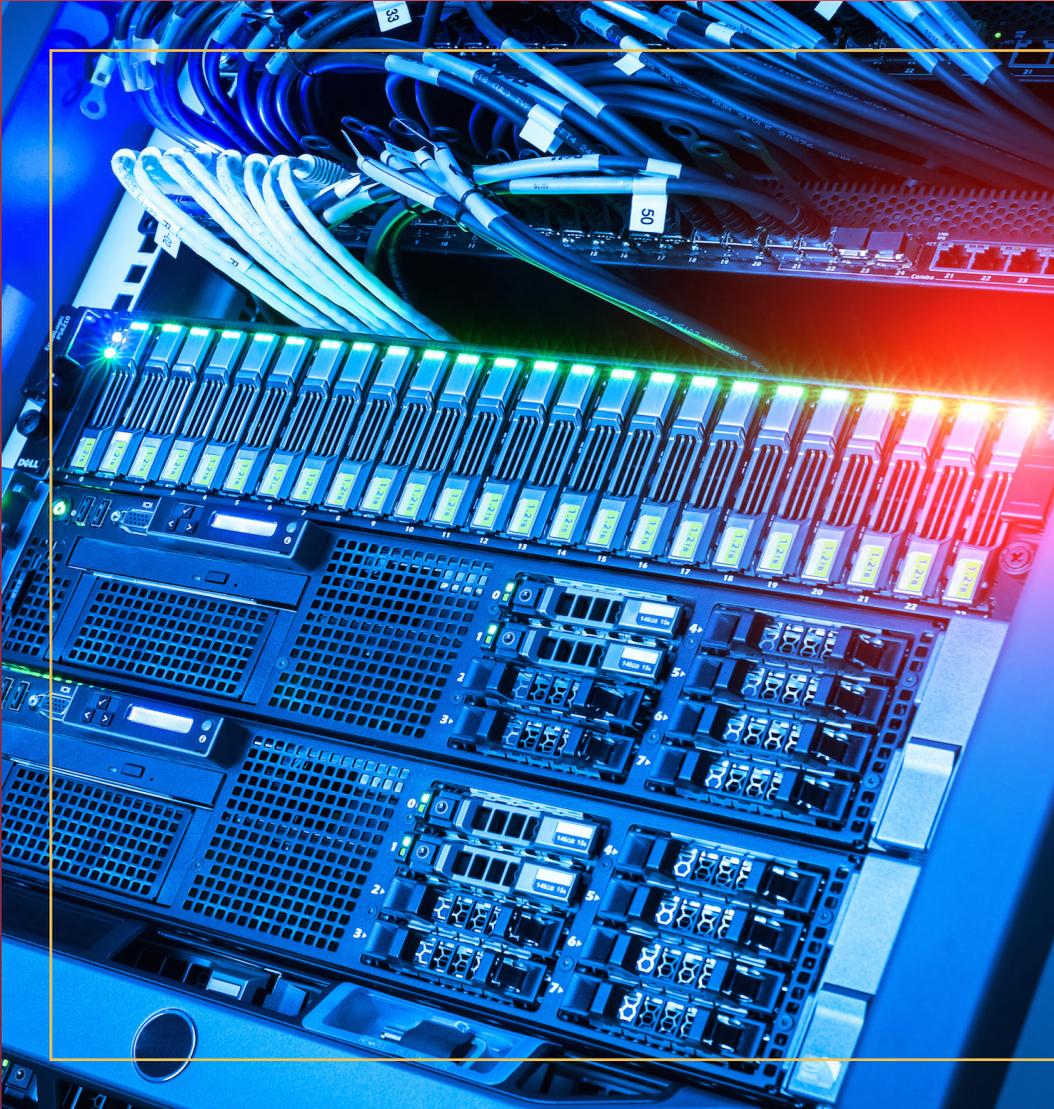


Adam Frisbee is a Udemey for Business and University of Utah information systems instructor with over 15 years of IT experience. His technical expertise is in cloud computing, Splunk, data warehousing, and systems analysis, design, and administration. Adam is also passionate about education, as is shown by his many years of successful university-level teaching. In addition to teaching on Udemey for Business and at the university, he is currently the IT director of a large biotechnology firm.



Find out how Udemey for Business can train your IT team for a Splunk deployment.

REQUEST DEMO



GraphQL



By Stephen Grider

Udemy for Business instructor and
engineering architect

What is GraphQL and why does it matter to your business?

Demand for GraphQL courses on the **Udemy** platform grew this year, making it the **#2 hottest IT skill trending in 2018**. Excitement around GraphQL has been building over the last few years, but is it warranted? GraphQL is a data query language **developed internally by Facebook** in 2012 to improve how data is retrieved and loaded on mobile phones. GraphQL was publicly released in 2015 and provides **an alternative to REST** and ad-hoc web service architectures.

Anytime you "like" a post on Facebook, or check out your latest messages from a friend on Snapchat, information gets exchanged between your device and a database connected to the internet. This exchange of information



GraphQL is the **#2 hottest IT skill** trending on Udemy in 2018



Graph QL



between your device and a database is frequently referred to as “data loading.” Virtually all mobile apps and websites use some sort of data loading procedure to make their content come to life.

Given the importance of this data loading function, engineers have developed several different systems to more efficiently and quickly transfer information. One of the currently popular systems, known as REST, has been used in tens of thousands of different products in the last decade, but it’s not a perfect solution. In particular, using REST to transfer data often results in trouble loading “relational data,” which can sometimes load more information than required, and imposes stricter requirements and limited flexibility for front end developers.

As a solution, Facebook developed a competing technique for transferring data between devices and databases called GraphQL. In recent years, this newer loading methodology has been adopted by hundreds of companies to solve issues with older REST approaches.

Benefits of GraphQL

Here are some of the issues with REST, how GraphQL solves these problems, and some of the key benefits of GraphQL over REST.

1 | Increased flexibility

REST became popular over the last decade due to its predictable implementation. An engineer who learns how to use REST at one engineering organization can usually transfer that knowledge easily to another organization. But this blessing is also a curse when it comes to flexibility.

One of the core ideas of REST is the practice of defining a very limited set of ways that a developer can interact with data. While creating standards, it also limits flexibility and slows performance. For example, if a blogging application implemented a REST interface, a developer would be restricted to interacting with blog posts in one of





four (sometimes more) ways: retrieving a single blog post, retrieving a list of blog posts, updating a single blog post, or creating a single blog post. These methods compose the only way to interact with blog posts when following REST conventions. More advanced interactions, such as listing all blog posts a user has created and the number of comments each blog post creates problems for REST and affects how quickly the data loads. GraphQL solves these problems by creating more flexibility for developers and allowing for more efficient loading of relational data (more on this below).

2 | Better loading of relational data

GraphQL dramatically simplifies the process of loading relational data. An example of relational data might be a blog post that has some associated comments. The comments only make sense when applied to the blog post they were made on, or we might say that the comments 'belong' to the blog post.

Graph QL

GraphQL improves the retrieval of relational data by allowing developers to fetch multiple related pieces of information, such as a blog post and its related comments, in a single network request over the Internet. In general, developers try to minimize the number of requests being made by an app to better accommodate users with lower quality network connections. In theory, there is no limit to the amount of information that can be retrieved in a single GraphQL request. A GraphQL server could share a list of blogs posts, comments that belong to them, the authors of each comment, and the total number of comments created by each author.

In a REST implementation, loading a blog post and an associated list of comments would be a two-step process. First, the app would have to make a request over the Internet for a blog post, then make a follow-up, or sequential request to fetch the related list of comments. This represents two separate requests being made over the internet which slows performance, especially on mobile devices.





3 | Increased performance

Besides allowing a developer to more easily access related data in a single request, GraphQL also allows developers to reduce the size of each request—enhancing performance and speed. For example, imagine how a blog post is displayed to a user in a blogging app. The post might have a title, date of publication, and some amount of content – this is all information that the app would need to function properly. However, an application might also store many other pieces of information attached to that blog post, such as the number of views the blog post has, the topics it mentions, or even a list of related blog posts. These extra pieces of information might not be needed by the application at all to get the blog post to appear to the user, so if that information is communicated from the database to the app, more data is being shared than needed.

GraphQL allows developers to precisely define the properties they need to access on a particular record, such as a blog post. The benefit is clear: there is less of a chance of 'oversharing' data, and less data is transferred from a back end server to the user's device, resulting in a performance gain.

This is in stark comparison to following REST conventions, which do not easily provide a developer the ability to tailor the amount of information that is served on each record.

4 | Streamline front end and back end developer collaboration

Whereas REST implementations define a strict set of principles and conventions for fetching data, GraphQL allows front end developers to change the information they request on the fly to better suit the application they are trying to build.

This can be done at the sole discretion of a front end developer with little to no input from a back end developer. In contrast, REST implementations traditionally require closer collaboration between front end and back end developers, which is commonly seen as a drain on productivity due to the need for one party to complete a programming task before the other can begin working on their feature.

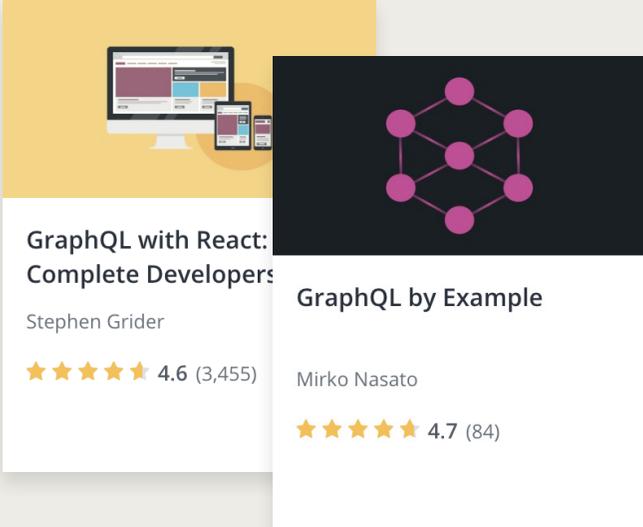
How to deploy GraphQL

Roadmap for implementation

Now that you understand the benefits of GraphQL, how can you implement this new technology in your organization and apply GraphQL to your next app? What kind of skills will your developer team need? What steps should you take to implement? Here are some key steps you'll need to take.

Train your developers on GraphQL. [Udemy for Business](#) offers a [collection of GraphQL courses](#) that will get your development team up and running with this next-gen interface. My course [GraphQL with React: The Complete Developer's Guide](#) can help you and your team get started on how to implement GraphQL at your organization.

Recommended courses



GraphQL with React: The Complete Developer's Guide
Stephen Grider
★★★★★ 4.6 (3,455)

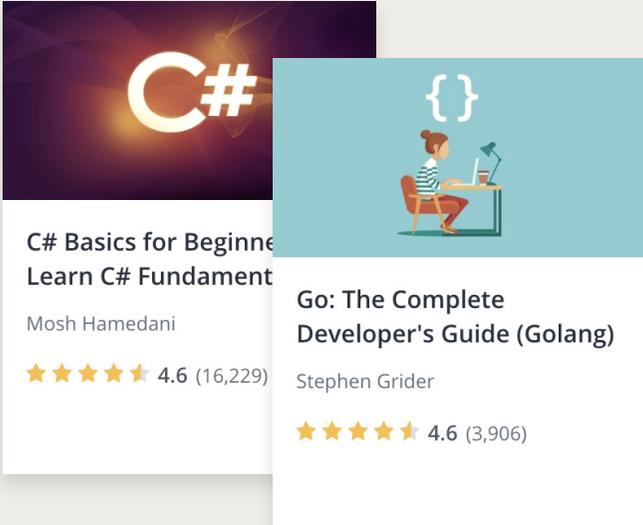
GraphQL by Example
Mirko Nasato
★★★★★ 4.7 (84)

[GraphQL with React: The Complete Developer's Guide](#)

[GraphQL by Example](#)

Bring your developers up to speed on programming languages supported by GraphQL. Facebook maintains an official reference of a GraphQL server written with Javascript, but open source community-led implementations are available for the majority of popular back end languages. For example, first-class **GraphQL implementations** are available for C#, Go, Java, Scala, Python, and Ruby, to name just a few. Bring your developers up to speed on these programming languages on UdeMy for Business with courses like **C# Basics for Beginners**, **Go: The Complete Developer's Guide**, **The Complete Java Masterclass**, or **The Complete Python Bootcamp**.

Recommended courses



C# Basics for Beginners
Learn C# Fundamentals
Mosh Hamedani
★★★★★ 4.6 (16,229)

Go: The Complete Developer's Guide (Golang)
Stephen Grider
★★★★★ 4.6 (3,906)

[C# Basics for Beginners](#)

[Go: The Complete Developer's Guide](#)

Focus on developing the back end first. Once you've picked out a GraphQL implementation to follow, you'll want to focus on developing the back end experience of your app first. Traditionally, the development of an app that uses GraphQL starts on the back end, to better allow flexible access to information stored in the database.

Move to the front end next. Once the back end server has integrated GraphQL support, move your engineers on to the front end. You'll find that your front end developers can move along at a vastly improved pace, as they can more flexibly define the information required to implement each feature in your app. Gone are the days of front end developers waiting weeks for a database engineer to implement a particular interface.

If you get stuck along the way, remember that there are [countless sources and guides](#) available to help you get started with GraphQL.

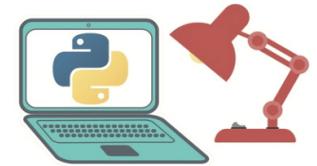
Recommended courses



Complete Java Masterclass updated for Java 10

Tim Buchalka, Tim Buchalka

★★★★★ 4.6 (50,047)



Complete Python Bootcamp: Go from zero to...

Jose Portilla, Pierian Data Internati...

★★★★★ 4.5 (78,218)

[The Complete Java Masterclass](#)

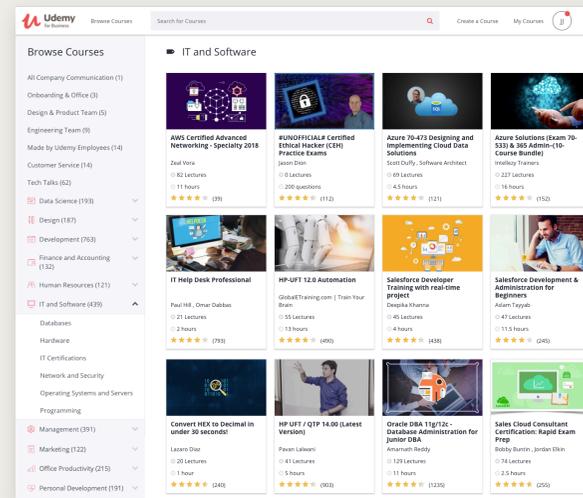
[The Complete Python Bootcamp](#)

GraphQL: Safe and solid bet for any project

Although REST has served us well for many years, its many limitations have made GraphQL more of a joy to work with. By using GraphQL, you'll have an easier time working with relational data, enjoy improved performance, and increase your front end developer productivity. Thanks to the support of Facebook and the open source community, development of GraphQL has been evolving over the years--making it a safe and solid bet for any new professional project.



Stephen Grider is a Udemey for Business instructor and engineering architect. He has built complex Javascript front ends for top corporations in the San Francisco Bay Area. With over 200,000 students, Stephen teaches a wide variety of courses from **GraphQL** and **Ethereum and Solidity** to **Webpack** on Udemey for Business.



Upskill your IT team on GraphQL with Udemey for Business.

REQUEST DEMO

Empower your IT team to solve data challenges

As businesses stay competitive by providing real-time data to customers faster and harnessing data insights better, your IT infrastructure will need to handle these new data challenges. Empower your IT team to enhance your data infrastructure by investing in Udemy for Business.

[REQUEST DEMO](#)



About Udemy for Business

Udemy for Business helps companies stay competitive in the digital transformation of the workplace by offering fresh, on-demand learning content powered by a content marketplace. Our mission is to help employees do whatever comes next—whether that’s tackling the next project, learning a new skill, or mastering a role. We do this by curating 2,500+ of Udemy’s top-rated courses taught by the world’s leading experts. Our content covers key business and technical topics ranging from development and IT to leadership, marketing, design, stress management, and much more. In addition to a curated content collection, organizations can host and distribute their own proprietary content on Udemy. Leading organizations including PayPal, Lyft, Booking.com, Pinterest, Adidas, and HSBC choose Udemy for Business to upskill their workforce and drive learning forward.

PayPal



Lyft



SIEMENS

Learn more: business.udemy.com

© 2018 Udemy. All Rights Reserved.

