



# Storm Blueprints: Patterns for Distributed Real-time Computation

By P. Taylor Goetz, Brian O'Neill



**Storm Blueprints: Patterns for Distributed Real-time Computation** By P. Taylor Goetz, Brian O'Neill

Use Storm design patterns to perform distributed, realtime big data processing, and analytics for realworld use cases

## About This Book

- Process high-volume log files in real time while learning the fundamentals of Storm topologies and system deployment.
- Deploy Storm on Hadoop (YARN) and understand how the systems complement each other for online advertising and trade processing.
- Follow along as each chapter presents a new problem and the architectural pattern, design, and implementation of a solution.

## Who This Book Is For

Although the book focuses primarily on Java development with Storm, the patterns are more broadly applicable and the tips, techniques, and approaches described in the book apply to architects, developers, and operations.

Additionally, the book should provoke and inspire applications of distributed computing to other industries and domains. Hadoop enthusiasts will also find this book a good introduction to Storm, providing a potential migration path from batch processing to the world of real-time analytics.

## What You Will Learn

- Learn the fundamentals of Storm
- Install and configure storm in pseudo-distributed and fully-distributed mode
- Familiarize yourself with the fundamentals of Trident and distributed state
- Design patterns for data flows in a distributed system
- Create integration patterns for persistence mechanisms such as Titan
- Deploy and run Storm clusters by leveraging YARN
- Achieve continuous availability and fault tolerance through distributed storage
- Recognize centralized logging mechanisms and processing

- Implement polyglot persistence and distributed transactions
- Calculate the effectiveness of a campaign using click-through analysis

## In Detail

Storm is the most popular framework for real-time stream processing. Storm provides the fundamental primitives and guarantees required for fault-tolerant distributed computing in high-volume, mission critical applications. It is both an integration technology as well as a data flow and control mechanism, making it the core of many big data platforms. Storm is essential if you want to deploy, operate, and develop data processing flows capable of processing billions of transactions.

"Storm: Distributed Real-time Computation Blueprints" covers a broad range of distributed computing topics, including not only design and integration patterns, but also domains and applications to which the technology is immediately useful and commonly applied. This book introduces you to Storm using real-world examples, beginning with simple Storm topologies. The examples increase in complexity, introducing advanced Storm concepts as well as more sophisticated approaches to deployment and operational concerns.

This book covers the domains of real-time log processing, sensor data analysis, collective and artificial intelligence, financial market analysis, Natural Language Processing (NLP), graph analysis, polyglot persistence and online advertising. While exploring distributed computing applications in each of those domains, the book covers advanced Storm topics such as Trident and Distributed State, as well as integration patterns for Druid and Titan. Simultaneously, the book also describes the deployment of Storm to YARN and the Amazon infrastructure, as well as other key operational concerns such as centralized logging.

By the end of the book, you will have gained an understanding of the fundamentals of Storm and Trident and be able to identify and apply those fundamentals to any suitable problem.



[Download Storm Blueprints: Patterns for Distributed Real-ti ...pdf](#)



[Read Online Storm Blueprints: Patterns for Distributed Real- ...pdf](#)

# Storm Blueprints: Patterns for Distributed Real-time Computation

By P. Taylor Goetz, Brian O'Neill

**Storm Blueprints: Patterns for Distributed Real-time Computation** By P. Taylor Goetz, Brian O'Neill

**Use Storm design patterns to perform distributed, realtime big data processing, and analytics for realworld use cases**

## About This Book

- Process high-volume log files in real time while learning the fundamentals of Storm topologies and system deployment.
- Deploy Storm on Hadoop (YARN) and understand how the systems complement each other for online advertising and trade processing.
- Follow along as each chapter presents a new problem and the architectural pattern, design, and implementation of a solution.

## Who This Book Is For

Although the book focuses primarily on Java development with Storm, the patterns are more broadly applicable and the tips, techniques, and approaches described in the book apply to architects, developers, and operations.

Additionally, the book should provoke and inspire applications of distributed computing to other industries and domains. Hadoop enthusiasts will also find this book a good introduction to Storm, providing a potential migration path from batch processing to the world of real-time analytics.

## What You Will Learn

- Learn the fundamentals of Storm
- Install and configure storm in pseudo-distributed and fully-distributed mode
- Familiarize yourself with the fundamentals of Trident and distributed state
- Design patterns for data flows in a distributed system
- Create integration patterns for persistence mechanisms such as Titan
- Deploy and run Storm clusters by leveraging YARN
- Achieve continuous availability and fault tolerance through distributed storage
- Recognize centralized logging mechanisms and processing
- Implement polyglot persistence and distributed transactions
- Calculate the effectiveness of a campaign using click-through analysis

## In Detail

Storm is the most popular framework for real-time stream processing. Storm provides the fundamental primitives and guarantees required for fault-tolerant distributed computing in high-volume, mission critical applications. It is both an integration technology as well as a data flow and control mechanism, making it the

core of many big data platforms. Storm is essential if you want to deploy, operate, and develop data processing flows capable of processing billions of transactions.

"Storm: Distributed Real-time Computation Blueprints" covers a broad range of distributed computing topics, including not only design and integration patterns, but also domains and applications to which the technology is immediately useful and commonly applied. This book introduces you to Storm using real-world examples, beginning with simple Storm topologies. The examples increase in complexity, introducing advanced Storm concepts as well as more sophisticated approaches to deployment and operational concerns.

This book covers the domains of real-time log processing, sensor data analysis, collective and artificial intelligence, financial market analysis, Natural Language Processing (NLP), graph analysis, polyglot persistence and online advertising. While exploring distributed computing applications in each of those domains, the book covers advanced Storm topics such as Trident and Distributed State, as well as integration patterns for Druid and Titan. Simultaneously, the book also describes the deployment of Storm to YARN and the Amazon infrastructure, as well as other key operational concerns such as centralized logging.

By the end of the book, you will have gained an understanding of the fundamentals of Storm and Trident and be able to identify and apply those fundamentals to any suitable problem.

### **Storm Blueprints: Patterns for Distributed Real-time Computation By P. Taylor Goetz, Brian O'Neill Bibliography**

- Sales Rank: #1546725 in Books
- Published on: 2014-04-17
- Released on: 2014-03-26
- Original language: English
- Number of items: 1
- Dimensions: 9.25" h x .76" w x 7.50" l, 1.27 pounds
- Binding: Paperback
- 336 pages



[Download Storm Blueprints: Patterns for Distributed Real-ti ...pdf](#)



[Read Online Storm Blueprints: Patterns for Distributed Real- ...pdf](#)

## Download and Read Free Online Storm Blueprints: Patterns for Distributed Real-time Computation By P. Taylor Goetz, Brian O'Neill

---

### Editorial Review

#### About the Author

##### **P. Taylor Goetz**

P. Taylor Goetz is an Apache Storm committer and release manager and has been involved with the usage and development of Storm since it was first released as open source in October of 2011. As an active contributor to the Storm user community, Taylor leads a number of open source projects that enable enterprises to integrate Storm into heterogeneous infrastructure. Presently, he works at Hortonworks where he leads the integration of Storm into Hortonworks Data Platform (HDP). Prior to joining Hortonworks, he worked at Health Market Science where he led the integration of Storm into HMS' next generation Master Data Management platform with technologies including Cassandra, Kafka, Elastic Search, and the Titan graph database.

##### **Brian O'Neill**

Brian O'Neill is a husband, hacker, hiker, and kayaker. He is a fisherman and father as well as big data believer, innovator, and distributed computing dreamer. He has been a technology leader for over 15 years and is recognized as an authority on big data. He has experience as an architect in a wide variety of settings, from startups to Fortune 500 companies. He believes in open source and contributes to numerous projects. He leads projects that extend Cassandra and integrate the database with indexing engines, distributed processing frameworks, and analytics engines. He won InfoWorld's Technology Leadership award in 2013. He authored the Dzone reference card on Cassandra and was selected as a Datastax Cassandra MVP in 2012 and 2013. In the past, he has contributed to expert groups within the Java Community Process (JCP) and has patents in artificial intelligence and contextbased discovery. He is proud to hold a B.S. in Computer Science from Brown University. Presently, Brian is Chief Technology Officer for Health Market Science (HMS), where he heads the development of their big data platform focused on data management and analysis for the healthcare space. The platform is powered by Storm and Cassandra and delivers realtime data management and analytics as a service.

### Users Review

#### From reader reviews:

##### **Ellen Farnsworth:**

Book is to be different for each grade. Book for children till adult are different content. As we know that book is very important for us. The book Storm Blueprints: Patterns for Distributed Real-time Computation seemed to be making you to know about other know-how and of course you can take more information. It doesn't matter what advantages for you. The guide Storm Blueprints: Patterns for Distributed Real-time Computation is not only giving you much more new information but also being your friend when you truly feel bored. You can spend your personal spend time to read your publication. Try to make relationship with

all the book Storm Blueprints: Patterns for Distributed Real-time Computation. You never sense lose out for everything in the event you read some books.

### **Kay Young:**

The particular book Storm Blueprints: Patterns for Distributed Real-time Computation will bring you to the new experience of reading any book. The author style to describe the idea is very unique. Should you try to find new book to see, this book very acceptable to you. The book Storm Blueprints: Patterns for Distributed Real-time Computation is much recommended to you you just read. You can also get the e-book from official web site, so you can easier to read the book.

### **Brian Kelley:**

Beside this specific Storm Blueprints: Patterns for Distributed Real-time Computation in your phone, it may give you a way to get closer to the new knowledge or info. The information and the knowledge you might got here is fresh from the oven so don't possibly be worry if you feel like an older people live in narrow commune. It is good thing to have Storm Blueprints: Patterns for Distributed Real-time Computation because this book offers for your requirements readable information. Do you often have book but you would not get what it's all about. Oh come on, that will not end up to happen if you have this with your hand. The Enjoyable set up here cannot be questionable, such as treasuring beautiful island. So do you still want to miss this? Find this book and also read it from at this point!

### **Janice Leon:**

As a student exactly feel bored in order to reading. If their teacher expected them to go to the library or make summary for some guide, they are complained. Just small students that has reading's spirit or real their pastime. They just do what the trainer want, like asked to go to the library. They go to there but nothing reading very seriously. Any students feel that looking at is not important, boring in addition to can't see colorful pics on there. Yeah, it is to be complicated. Book is very important for you. As we know that on this period, many ways to get whatever you want. Likewise word says, ways to reach Chinese's country. Therefore this Storm Blueprints: Patterns for Distributed Real-time Computation can make you experience more interested to read.

## **Download and Read Online Storm Blueprints: Patterns for Distributed Real-time Computation By P. Taylor Goetz, Brian O'Neill #BHW6K0QEX8M**

# **Read Storm Blueprints: Patterns for Distributed Real-time Computation By P. Taylor Goetz, Brian O'Neill for online ebook**

Storm Blueprints: Patterns for Distributed Real-time Computation By P. Taylor Goetz, Brian O'Neill Free PDF d0wnl0ad, audio books, books to read, good books to read, cheap books, good books, online books, books online, book reviews epub, read books online, books to read online, online library, greatbooks to read, PDF best books to read, top books to read Storm Blueprints: Patterns for Distributed Real-time Computation By P. Taylor Goetz, Brian O'Neill books to read online.

## **Online Storm Blueprints: Patterns for Distributed Real-time Computation By P. Taylor Goetz, Brian O'Neill ebook PDF download**

**Storm Blueprints: Patterns for Distributed Real-time Computation By P. Taylor Goetz, Brian O'Neill Doc**

**Storm Blueprints: Patterns for Distributed Real-time Computation By P. Taylor Goetz, Brian O'Neill MobiPocket**

**Storm Blueprints: Patterns for Distributed Real-time Computation By P. Taylor Goetz, Brian O'Neill EPub**

**BHW6K0QEX8M: Storm Blueprints: Patterns for Distributed Real-time Computation By P. Taylor Goetz, Brian O'Neill**