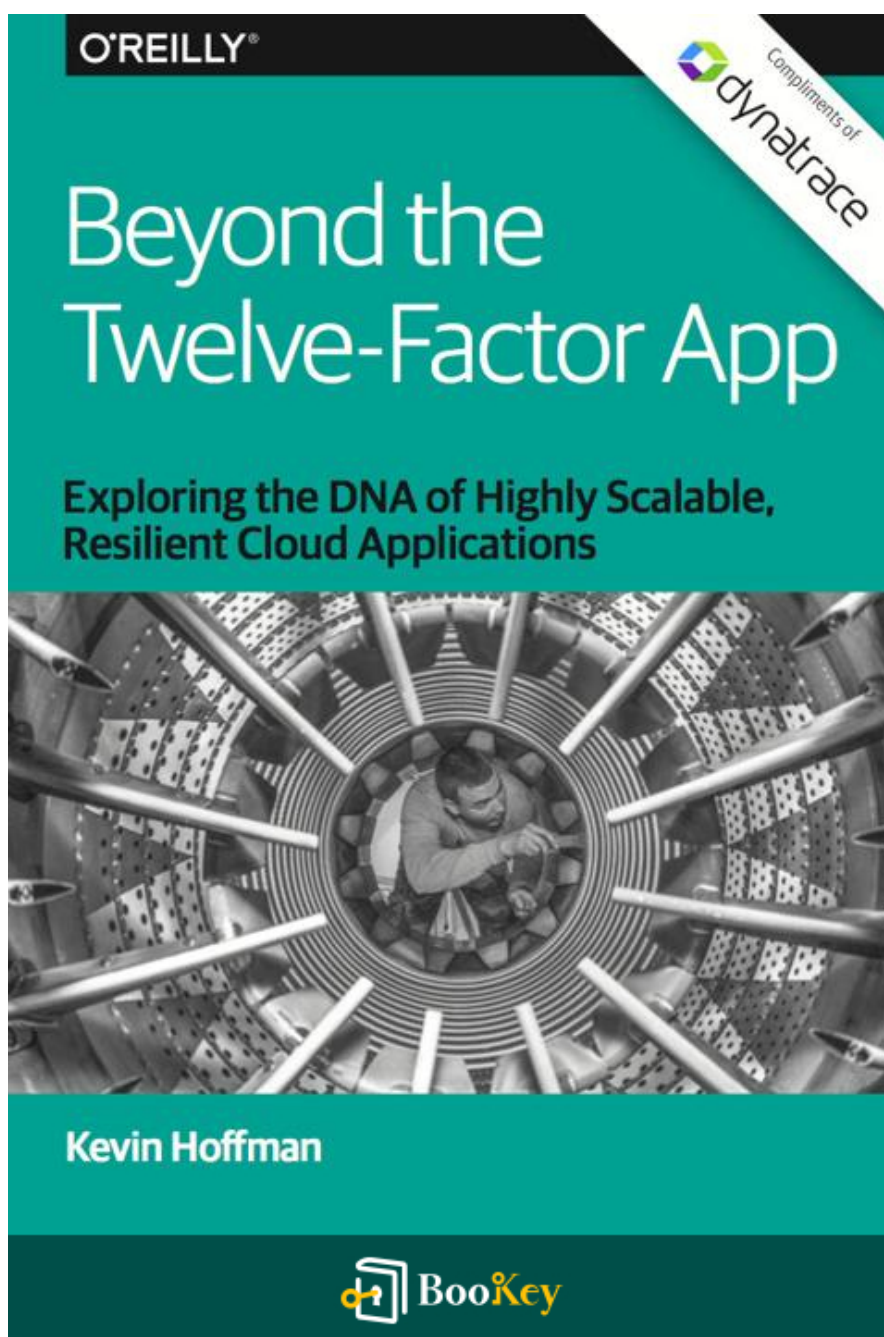


Beyond The Twelve-factor App PDF (Limited Copy)

Kevin Hoffman



More Free Book



Scan to Download

Beyond The Twelve-factor App Summary

Mastering Cloud Architecture for Competitive Edge in Today's
Marketplace.

Written by New York Central Park Page Turners Books Club

More Free Book



Scan to Download

About the book

In "Beyond the Twelve-Factor App," Kevin Hoffman delves into the evolution of cloud computing, illustrating how it has transitioned from a specialized solution embraced primarily by startups to a fundamental component of enterprise infrastructure. The book emphasizes the importance of mastering cloud system design principles as organizations strive to navigate the complexities of a rapidly changing technology landscape.

Hoffman begins by outlining the core tenets of cloud architecture, building upon the foundational concepts introduced in the Twelve-Factor App methodology, which serves as a model for developing scalable and maintainable web applications. He enriches this discussion by introducing new concepts and practices that are vital for organizations of all sizes to implement in order to harness the full potential of cloud technologies.

Throughout the chapters, the author provides insights into various cloud design patterns and explains how they can lead to more resilient and efficient systems. He elaborates on the significance of automation, microservices, and continuous integration/continuous deployment (CI/CD), drawing attention to how these methodologies enable teams to innovate rapidly while maintaining high-quality outputs.

As the narrative progresses, Hoffman also addresses common challenges that

More Free Book



Scan to Download

enterprises face during their cloud transformation journeys, such as managing legacy systems and ensuring security and compliance. He offers strategies for overcoming these obstacles, emphasizing a culture of collaboration and adaptability within teams, which are crucial for success in a competitive landscape.

In conclusion, "Beyond the Twelve-Factor App" serves as a comprehensive guide for organizations aiming to remain at the forefront of technological advancements. By empowering teams with practical knowledge and forward-thinking strategies, Hoffman ensures that they are well-equipped to thrive in an era where cloud computing is not just an option, but a necessity for sustainable growth and innovation.

More Free Book



Scan to Download

About the author

In the chapters of "Beyond the Twelve-Factor App," Kevin Hoffman delves into the evolving landscape of cloud-native application development, providing a clear framework for modern software engineering. As a seasoned software architect, Hoffman's insights stem from his extensive two-decade experience in the field, which allows him to effectively address the challenges developers face in transitioning from traditional monolithic architectures to agile, cloud-based systems.

The narrative begins by establishing the foundational elements of cloud-native applications, emphasizing the importance of scalability and resilience. Hoffman introduces core concepts such as microservices and containerization, explaining how they enable teams to build applications that can effortlessly adapt to changing demands. He underscores the need for robust design patterns and architecture that support continuous delivery, pragmatic testing, and automated deployments.

As the chapters progress, Hoffman critiques the original Twelve-Factor App methodology, providing a modern interpretation that incorporates emerging trends and technologies. He addresses the growing complexity of applications and the necessity for best practices in managing dependencies, data storage, and configuration. This includes insights on how to effectively orchestrate services using platforms like Kubernetes, which has become the

More Free Book



Scan to Download

industry standard for managing containerized applications.

Moreover, Hoffman emphasizes the significance of monitoring and observability in cloud-native environments, encouraging developers to adopt a proactive approach to application performance management. By integrating logging, metrics, and tracing, teams can gain real-time visibility into their systems, allowing for quicker identification and resolution of issues.

Throughout the text, Hoffman also highlights the cultural shift necessary for successful cloud adoption, advocating for collaborative practices that break down silos within organizations. By fostering a culture of innovation and continuous learning, teams can harness the full potential of cloud-native methodologies.

In conclusion, Hoffman's exploration in "Beyond the Twelve-Factor App" serves as a vital guide for practitioners seeking to refine their cloud delivery strategies. With a blend of practical advice, real-world anecdotes, and a focus on actionable insights, he equips readers with the tools needed to navigate the complexities of modern software development in a cloud-centric world.

More Free Book



Scan to Download



Try Bookey App to read 1000+ summary of world best books

Unlock **1000+** Titles, **80+** Topics

New titles added every week

- Brand
- Leadership & Collaboration
- Time Management
- Relationship & Communication
- Business Strategy
- Creativity
- Public
- Money & Investing
- Know Yourself
- Positive Psychology
- Entrepreneurship
- World History
- Parent-Child Communication
- Self-care
- Mind & Spirituality

Insights of world best books



Free Trial with Bookey

Summary Content List

Chapter 1:

Chapter 2: One Codebase, One Application

Chapter 3: API First

Chapter 4: Dependency Management

Chapter 5: Design, Build, Release, Run

Chapter 6: Configuration, Credentials, and Code

Chapter 7: Logs

Chapter 8: Disposability

Chapter 9: Backing Services

Chapter 10: Environment Parity

Chapter 11: Administrative Processes

Chapter 12: Port Binding

Chapter 13: Stateless Processes

Chapter 14: Telemetry

Chapter 15: Authentication and Authorization

Chapter 16: A Word on Cloud Native

More Free Book



Scan to Download

Chapter 1 Summary:

Preface

In an era where complex digital discussions require a shared vocabulary, buzzwords have become essential. "Twelve-Factor Application" is one such phrase gaining traction in the realm of application development. However, the interpretation of what constitutes a "Twelve-Factor" application can vary widely among developers. This book endeavors to demystify the Twelve-Factor methodology, providing expanded insights to foster modern cloud-based application development.

The Original 12 Factors

The rise of cloud computing has transformed how applications are built and deployed, introducing complexities that many developers have struggled to navigate. Heroku, a leading platform for cloud application deployment, played a pivotal role in this transition by pioneering a framework that simplified the process. In 2012, Heroku's team formalized the Twelve Factors—a set of foundational principles designed to guide developers in creating cloud-native applications.

More Free Book



Scan to Download

These Twelve Factors are:

1. **One codebase tracked in revision control, many deploys** - Maintain a single codebase that can be deployed in multiple environments.
2. **Explicitly declare and isolate dependencies** - Clearly specify and manage external libraries your application depends on.
3. **Store configuration in the environment** - Use environment variables to keep configuration separate from code.
4. **Treat backing services as attached resources** - View services such as databases and caches as interchangeable components.
5. **Strictly separate build and run stages** - Distinguish between the processes that build the application and those that run it.
6. **Execute the app as one or more stateless processes** - Design applications to handle any request without relying on stored information.
7. **Export services via port binding** - Make services accessible over a network by binding them to a specific port.
8. **Scale out via the process model** - Enable scaling through the addition of multiple processes rather than scaling up a single instance.
9. **Maximize robustness with fast startup and graceful shutdown** - Ensure quick initialization and orderly termination of processes.
10. **Keep development, staging, and production as similar as possible** - Minimize discrepancies between environments to avoid deployment issues.
11. **Treat logs as event streams** - Handle logs as real-time data streams for effective monitoring and troubleshooting.

More Free Book



Scan to Download

12. Run admin/management tasks as one-off processes - Manage administrative tasks in isolated execution environments.

These principles form the bedrock of developing cloud-ready applications. As technology evolves, the need arises to adapt these tenets to fit more contemporary practices.

Beyond the Twelve-Factor Application

This book introduces an enhanced framework that builds on the original Twelve Factors, adapting to modern demands such as telemetry, security, and API-first design. The updated guidelines reflect a prioritized approach that highlights the importance of these considerations in contemporary cloud-native applications. The revised principles include:

- 1. One codebase, one application** - Simplifying the structure to streamline development.
- 2. API first** - Prioritizing APIs as the foundation of application interaction and integration.
- 3. Dependency management** - Emphasizing clear and organized control of dependencies.
- 4. Design, build, release, and run** - Harmonizing the development lifecycle for increased efficiency.

More Free Book



Scan to Download

5. **Configuration, credentials, and code** - Keeping configuration details secure and separate from the application logic.
6. **Logs** - Utilizing comprehensive logging for better insights into application performance.
7. **Disposability** - Designing applications to be ephemeral and easily replaced.
8. **Backing services** - Treating additional services as seamless extensions.
9. **Environment parity** - Ensuring congruence across development and production environments.
10. **Administrative processes** - Clearly defining administrative tasks as self-contained operations.
11. **Port binding** - Maintaining clear communication pathways.
12. **Stateless processes** - Encouraging simplicity in state management.
13. **Concurrency** - Facilitating parallel processing to improve performance.
14. **Telemetry** - Incorporating monitoring mechanisms to observe application behavior.
15. **Authentication and authorization** - Ensuring secure access and user management.

These updated principles offer a robust framework for developing

More Free Book



Scan to Download

cloud-native applications, fostering a cycle of continuous improvement that benefits both new and legacy developers as they transition into the cloud-centric landscape of modern software development.

More Free Book



Scan to Download

Chapter 2 Summary: One Codebase, One Application

In this summary, we delve into the concepts of immutable releases and the guiding principle of "One Codebase, One Application," essential for efficient software development and deployment.

Immutable Releases

Immutable releases refer to build artifacts that remain unchanged once created. These artifacts are critical for several reasons: they facilitate effective testing, ensure consistency between development and production environments, and help teams predict the outcomes of deployments accurately.

One Codebase, One Application

The notion of "One Codebase, One Application" emphasizes the importance of maintaining a single, centralized codebase within a version control system. This approach allows for the generation of multiple immutable releases tailored for various environments, thereby enhancing deployment efficiency and reducing lead times in production.

- **Understanding Codebase:** A codebase is essentially a repository, or a set of repositories, that has a unified root directory. When multiple codebases exist chaotically, it often signals the need for restructuring or decomposition, suggesting that the application should be viewed as a

More Free Book



Scan to Download

singular unit instead of fragmented components.

- **Common Violations:** Issues arise when applications are unnecessarily split across numerous repositories or when multiple applications are erroneously derived from a single codebase. Such violations complicate automation and hinder seamless deployment processes.

- **Conway's Law:** An insightful principle known as Conway's Law posits that the organization of a development team directly influences the architecture of the product they create. If team structures are disorganized or not aligned with the product's needs, coding efficiency can suffer. To counter this, it may be beneficial to organize smaller teams to focus on individual applications or microservices, fostering better productivity.

- **Shared Code Strategies:** While it is acceptable to share code among multiple applications, this necessitates treating the shared code as an independent codebase. It is prudent to assess whether this shared code warrants its own separation as a vendored product, which could provide additional benefits.

In summary, maintaining a well-organized and cohesive structure of codebases is vital. This organization not only streamlines development processes but also promotes effective deployment strategies, ultimately leading to improved software quality and reliability.

More Free Book



Scan to Download

Chapter 3 Summary: API First

API First in Modern Application Development

Introduction to API First

In today's cloud-native application development landscape, the traditional hierarchical dependency model is giving way to complex service ecosystems that are more interconnected and less rigid. The "API First" philosophy acknowledges this shift, emphasizing the critical role of APIs (Application Programming Interfaces) as key enablers for seamless integration and interaction within this network of services. By prioritizing APIs from the onset, developers can create applications that effectively leverage the benefits of this new ecosystem.

Why API First?

Adopting an API First approach enhances collaboration among development teams, as it encourages focus on public contracts (the agreements on how different services will communicate). This emphasis decreases the chance of integration failures, which are common when multiple teams are working simultaneously on various services. With APIs treated as primary artifacts, the development process remains fluid, allowing teams to integrate external

More Free Book



Scan to Download

services efficiently while maintaining internal workflow continuity.

Building Services with API First

The API First approach parallels the "mobile first" development mindset, where the initial focus is on creating robust APIs for consumption by diverse client applications. This philosophy is particularly impactful in cloud-native environments, as it encourages teams to plan and design APIs before writing code. Engaging in discussions with stakeholders early in the process helps clarify requirements and expectations, ensuring that APIs meet the needs of all potential users.

Tools and Standards for API First

A variety of tools and standards play a crucial role in supporting API First development, including API Blueprint for documentation and CI servers for continuous integration. These resources facilitate the documentation, testing, and integration of APIs across multiple services, allowing organizations to prototype quickly while ensuring that their APIs are rigorously tested. By leveraging these tools, teams can maintain high quality and consistency in their API offerings.

Benefits of API First Development

More Free Book



Scan to Download

Embracing an API First mentality fosters a more flexible and adaptable architecture, steering clear of tightly coupled monolithic systems. This flexibility enables organizations to respond to new demands and shifting consumer needs without significant overhauls to their existing systems. As a result, companies can evolve and scale their services more efficiently, gaining a substantial competitive edge in a rapidly changing market.

Conclusion

Implementing the API First approach is essential for modern development teams, offering significant advantages in investment and resource management. By cultivating a resilient and scalable infrastructure through this methodology, organizations position themselves to thrive in the ever-evolving landscape of software development.

More Free Book



Scan to Download

Chapter 4: Dependency Management

Chapter 4: Dependency Management

This chapter delves into the critical aspect of dependency management in modern software development, particularly as applications transition to cloud environments.

Introduction to Bootstrapping

At the heart of the chapter is the concept of "bootstrapping," which denotes the ability of applications to be self-sufficient by including all necessary dependencies within their own framework. This eliminates reliance on external resources, allowing for more robust and portable applications.

The Concept of the Mommy Server

Traditionally, enterprise applications have relied on a centralized server, often referred to as the "mommy server," which hosted and managed all necessary resources and dependencies. In stark contrast, the bootstrapped approach advocates for integrating all requirements into a single build artifact. This shift promotes independence from centralized servers, enabling applications to function autonomously.

More Free Book



Scan to Download

Cloud Maturation and Application Adaptation

As cloud technologies mature, applications must adapt to operate in these environments without the support of mommy servers. This shift requires developers to understand how to manage and transport dependencies independently, ensuring that applications remain functional and resilient.

Isolated Dependencies

The chapter emphasizes the importance of isolated dependencies—libraries and modules stored within the application or its immediate environment, rather than relying on shared repositories. Contemporary programming languages and frameworks frequently provide built-in tools for this kind of management, making it easier for developers to ensure that all dependencies are packaged within the application.

Modern Dependency Management Tools

Various tools cater to different programming languages, streamlining the dependency management process. For instance, Maven and Gradle serve the Java ecosystem, NuGet supports .NET applications, Bundler is used in Ruby, and godeps assists developers working with Go. These tools enable developers to specify and organize dependencies, ensuring that applications

More Free Book



Scan to Download

can function without prior installations on the system.

Risks of Improper Dependency Isolation

The chapter highlights the risks associated with neglecting proper dependency isolation. When dependencies are inadequately managed, conflicts between different library versions can cause runtime failures, leading to data loss or system breakdowns. This underscores the necessity of meticulous dependency management to prevent such issues.

Simplicity vs. Complexity in Applications

The balance between simplicity and complexity in application development is scrutinized. Although complex applications pose more daunting challenges in troubleshooting, prioritizing simplicity can lead to easier maintenance. Through effective dependency management, developers can automate deployments, reducing the assumptions that could lead to errors.

Practicality and Current Industry Practices

While embedding servers directly within application artifacts is an ideal approach for cloud deployment, practical constraints often call for a hybrid solution. Tools like buildpacks can bridge gaps by integrating applications with necessary server components. A disciplined approach to dependency

More Free Book



Scan to Download

management is essential for ensuring that applications not only survive but thrive in cloud environments, ultimately enhancing their efficiency and reliability.

Overall, this chapter presents a comprehensive view of how dependency management is evolving in tandem with technological advancements, stressing the need for self-sufficiency and robust practices in a cloud-centric future.

Install Bookey App to Unlock Full Text and Audio

Free Trial with Bookey





Why Bookey is must have App for Book Lovers



30min Content

The deeper and clearer interpretation we provide, the better grasp of each title you have.



Text and Audio format

Absorb knowledge even in fragmented time.



Quiz

Check whether you have mastered what you just learned.



And more

Multiple Voices & fonts, Mind Map, Quotes, IdeaClips...

Free Trial with Bookey



Chapter 5 Summary: Design, Build, Release, Run

Design, Build, Release, Run: Summary

In the realm of modern software development, the delineation of the design, build, release, and run stages is essential to efficiently manage applications. This structured yet iterative approach emboldens teams to achieve continuous integration and deployment, resulting in more frequent and reliable application releases.

Design Phase

The design stage prioritizes flexibility and incremental feature implementation over extensive pre-planning. Developers outline application dependencies and strategize how these components will be organized for deployment. A high-level design blueprint serves as a roadmap, with the understanding that adjustments will naturally arise during development.

Build Phase

In the build phase, the raw code is transformed into a versioned binary artifact, a crucial step that incorporates the previously declared dependencies into a structured format. Continuous Integration (CI) servers facilitate this

More Free Book



Scan to Download

automation, creating a streamlined process that connects multiple builds to their corresponding deployments. The concept of immutability in build artifacts reassures developers that their code will behave consistently across various environments.

Release Phase

In a cloud-native setting, the release phase focuses on amalgamating built artifacts with configurations tailored to specific environments, leading to unique and immutable releases. Each release receives a unique identifier, which is critical for tracking changes and reverting to previous versions if necessary. By separating the build and release processes, teams can improve historical audits and address potential deployment issues effectively.

Run Phase

During the run phase, applications are deployed using services provided by cloud vendors, while also enabling local development and testing.

Application containers play a pivotal role here, managing deployment and handling essential tasks such as scaling and monitoring system health. The overarching goal at this stage is to increase delivery speed while ensuring robust performance through automated testing and deployment practices.

Conclusion

More Free Book



Scan to Download

By following the structured pipeline of design, build, release, and run, teams can significantly enhance the speed of application delivery without sacrificing quality or dependability. This framework enables them to remain agile and responsive to evolving needs in a cloud-native environment, embodying the core principles of modern software development.

More Free Book



Scan to Download

Chapter 6 Summary: Configuration, Credentials, and Code

CHAPTER 5: Configuration, Credentials, and Code

In the world of software development, particularly as applications approach production, the handling of configuration, credentials, and code is critical. Missteps can lead to significant complications. Configuration elements are those values that can change depending on the environment where the application is deployed—such as a developer's workstation, the quality assurance (QA) stage, or the final production environment.

Key elements of configuration include crucial data like URLs connecting to various backing services (such as web services and email servers), database connection information, and credentials for third-party services (e.g., AWS, Google Maps, Twitter, Facebook). This information is typically found in properties files or formats like XML or YAML.

One fundamental principle in managing these configurations is to ensure a clear separation between configuration and code. Internal information that remains static across all deployments does not qualify as configuration. It is essential that credentials are never embedded within the codebase, as this practice can inadvertently expose sensitive information.

More Free Book



Scan to Download

A useful litmus test for assessing whether configuration and credentials are properly externalized is to consider if the source code could be made public, for example, by uploading it to a platform like GitHub. If doing so compromises sensitive or environment-specific details, then there is a flaw in the approach.

The practice of maintaining an external configuration allows for the deployment of immutable builds across various environments, striking a balance between development and production setups. Achieving this involves not just conceptualizing the need for externalization but also implementing it effectively.

Traditionally, developers have relied on property files or other configuration formats, but these methods are now seen as outdated, especially in cloud contexts where flexibility is paramount. A more robust solution is to eliminate these configuration files altogether, leading to code that pulls configuration data from environment variables. This approach is widely recognized as best practice in cloud platforms like Cloud Foundry or Heroku.

Additionally, leveraging specialized configuration management tools, such as the Spring Cloud Configuration Server, can greatly enhance the management of external configurations. These tools not only provide a way

More Free Book



Scan to Download

to expose configurations effectively but often come with revision control capabilities, which help in tracking changes and maintaining a history of user modifications. This structured approach to managing configuration ensures both security and adaptability in dynamic development environments.

More Free Book



Scan to Download

Chapter 7 Summary: Logs

Chapter 6: Logs

In this chapter, we delve into the concept of logging within the framework of cloud-native applications, aligning our discussion with the 11th factor of the Twelve-Factor App methodology: logs. Logs are essentially sequences of events captured and emitted by an application, organized in chronological order. They serve as critical documentation of an application's internal workings and behavior.

A fundamental principle for cloud-native applications is the shift away from traditional log management practices, where developers had direct control over log storage and configurations. Instead, in a cloud environment, applications should write log entries exclusively to standard output (stdout) and standard error (stderr). This transition may induce concerns about relinquishing control, but it is vital for embracing cloud-native methodologies that prioritize flexibility and scalability.

Furthermore, it's essential to categorize log management as a nonfunctional requirement typically handled by cloud providers or specialized external tools rather than being embedded within the application's code. Tools like the ELK (Elasticsearch, Logstash, Kibana) stack, Splunk, and Sumo Logic

More Free Book



Scan to Download

are invaluable for aggregating, processing, and storing logs efficiently. These solutions offer robust capabilities for log analysis, augmenting the application's performance without burdening the developers.

The benefits of adopting a decoupled logging strategy are significant. By separating log management from the application's internal operations, developers can streamline their code, utilize established industry tools, and quickly adapt to evolving log processing requirements without the need for extensive application modifications. This decoupling is especially advantageous in dynamic environments where application instances can scale elastically in response to demand.

In conclusion, embracing a cloud-native logging approach fosters cleaner codebases, allows developers to focus on core business functionalities, and enhances the ability to leverage cloud solutions for comprehensive log management. By recognizing and implementing these principles, development teams can ensure their applications are better equipped for the complexities of modern cloud environments.

More Free Book



Scan to Download

Chapter 8: Disposability

Disposability in Cloud-Native Applications

The concept of disposability, known as the ninth factor in the Twelve-Factor App methodologies, emphasizes the transient nature of applications within cloud environments. Just like the infrastructure that supports them, cloud-native applications are designed for temporary existence, allowing for seamless launching and halting of processes. This flexibility is essential for efficiently scaling, deploying, and recovering applications, making it a cornerstone of modern software development.

Rapid Startup and Shutdown

To fully utilize cloud capabilities, applications must be engineered for rapid startup and shutdown. Traditionally, developers accustomed to legacy enterprise systems may face challenges, as these applications often have long startup times that can extend to several minutes. This issue isn't limited to outdated software; even applications built with interpreted languages or poorly optimized code can experience sluggish startup sequences.

Impact of Slow Startup Times

More Free Book



Scan to Download

Prolonged startup times can have detrimental effects, especially in high-traffic environments. When an application takes too long to become operational, it risks rejecting numerous incoming requests, which can lead to poor user experiences and increased downtime. This can trigger alerts or failures in health checks, complicating the application's ability to successfully start up in cloud settings.

Challenges with Resource Management

Slow startup processes are not just inconvenient; they can significantly impact resource management during peak loads. When an application is slow to start, deploying additional instances to handle increased traffic becomes problematic. Additionally, if an application doesn't shut down quickly after a failure, it can hamper recovery efforts and lead to complications in resource disposal, posing risks such as data corruption.

Separating Long-Running Activities

To align with the principles of cloud-native design, applications should avoid lengthy tasks during startup, such as populating caches or preparing

More Free Book



Scan to Download

dependencies. Instead, these operations can be handled by dedicated backing services. By offloading such tasks, applications can ensure faster lifecycles and maintain optimal performance without overwhelming their front-end processes during startup, thus enhancing the overall efficiency and reliability of cloud-native applications.

Install Bookey App to Unlock Full Text and Audio

Free Trial with Bookey





Positive feedback

Sara Scholz

...tes after each book summary
...erstanding but also make the
...and engaging. Bookey has
...ling for me.

Fantastic!!!



I'm amazed by the variety of books and languages Bookey supports. It's not just an app, it's a gateway to global knowledge. Plus, earning points for charity is a big plus!

Masood El Toure

Fi



Ab
bo
to
my

José Botín

...ding habit
...o's design
...ual growth

Love it!



Bookey offers me time to go through the important parts of a book. It also gives me enough idea whether or not I should purchase the whole book version or not! It is easy to use!

Wonnie Tappkx

Time saver!



Bookey is my go-to app for summaries are concise, ins curated. It's like having acc right at my fingertips!

Awesome app!



I love audiobooks but don't always have time to listen to the entire book! bookey allows me to get a summary of the highlights of the book I'm interested in!!! What a great concept !!!highly recommended!

Rahul Malviya

Beautiful App



This app is a lifesaver for book lovers with busy schedules. The summaries are spot on, and the mind maps help reinforce wh I've learned. Highly recommend!

Alex Walk

Free Trial with Bookey

Chapter 9 Summary: Backing Services

Backing Services: A Comprehensive Overview

In the realm of cloud-native applications, backing services play a crucial role in ensuring application functionality. These services encompass a variety of essential components, including data storage solutions (like databases), messaging systems, caching mechanisms, and security functionalities. Given the ephemeral nature of cloud filesystems, even file storage should be viewed through the lens of a backing service rather than as a traditional filesystem.

Understanding Bound Resources

A bound resource is the essential link that connects an application to its backing services. This connection is facilitated through binding elements such as usernames, passwords, and URLs. By using bound resources, applications can consume these services while remaining decoupled from their specific implementations. This architectural choice enhances modularity and maintainability.

Guidelines for Effective Resource Binding

More Free Book



Scan to Download

To optimize the use of bound resources, several key rules must be adhered to:

1. Applications should express their requirement for backing services, allowing the cloud environment to handle the specifics of resource binding.
2. Resource bindings should be managed using external configuration to ensure adaptability and ease of updates.
3. It ought to be feasible to attach or detach backing services from applications without necessitating a full redeployment. This flexibility is essential for maintaining uptime and performance.

Illustrative Example of Binding

Consider an application that relies on an Oracle database. The application must declare its dependency on the database but should keep all connection configurations—such as credentials and connection strings—external. This prevents tight coupling between the application and any specific database instance, allowing for easy modifications or replacements if necessary.

Benefits of Treating Backing Services as Bound Resources

More Free Book



Scan to Download

Approaching backing services as bound resources brings significant advantages, notably in flexibility and resilience. For instance, if a database experiences downtime, it can be restarted or replaced without affecting the application's overall operation. This decoupling supports a more robust infrastructure, enabling applications to withstand service disruptions while maintaining functionality.

Implementing the Circuit Breaker Pattern

An effective strategy for managing communication with malfunctioning backing services is the circuit breaker pattern. This design pattern allows applications to temporarily halt interactions with a failing service, offering a fallback alternative. By defining proper backing services and integrating this pattern, administrators can redirect application bindings as necessary, enhancing resilience and preventing cascading failures.

In conclusion, modern cloud-native applications should adopt the perspective of treating backing services as bound resources. This strategic approach not only facilitates flexibility and resilience but also empowers effective error handling through mechanisms like the circuit breaker. By adhering to these principles, developers can build applications that are robust and well-equipped to handle the complexities of cloud environments.

More Free Book



Scan to Download

Chapter 10 Summary: Environment Parity

Chapter 10 Summary: Environment Parity

In the context of software development, "environment parity" refers to the practice of ensuring that all approved target environments—such as development, quality assurance (QA), and production—are as similar as possible. This similarity is crucial for minimizing deployment issues and enhancing the confidence of teams in the functioning of applications across different settings.

The Significance of Environment Parity

Organizations often grapple with inconsistencies in their environments, which can lead to anxiety about deployments and uncertainty regarding whether applications will operate as intended. The chapter outlines three primary challenges that contribute to these environment gaps:

1. **Time Delays:** When there are lengthy intervals between code check-ins and actual deployments, changes may be overlooked or forgotten, diminishing the overall confidence in the release process.
2. **Team Dynamics:** Deployment practices often vary by the size of the



organization. Smaller firms may allow developers to manage the entire deployment process, fostering direct responsibility, while larger entities tend to create multiple handoffs, complicating the deployment procedure.

3. **Resource Disparities:** Compromises made during the development phase, such as opting for simpler local setups, can widen the gap between development and production environments. This discrepancy can lead to reliability issues when applications are eventually deployed.

Strategies for Enhancing Environment Parity

1. Streamlining Deployment Timelines

Modern methodologies advocate for drastically reducing the time between code check-in and the deployment process from weeks or months to mere minutes or hours. This can be achieved through automated testing and deployment, with an emphasis on implementing zero-downtime deployments. Such practices significantly bolster team confidence in the stability and reliability of code changes.

2. Integrating Development and Deployment Roles

To simplify the deployment process, it is recommended that the same individuals who develop the code also manage its deployment, especially in



cloud-native settings. Automation plays a critical role here, enabling deployments to occur with minimal manual intervention, thus streamlining the workflow.

3. Consistent Resource Usage

Maintaining environment parity requires the use of standardized resource types across both development and production environments. Tools like Docker can facilitate this uniformity, enhancing the predictability and reliability of applications in production scenarios.

Continuous Delivery Mindset

The chapter emphasizes that every code commit should be viewed as a potential candidate for deployment, encouraging a robust continuous delivery pipeline. By ensuring that environments remain similar, development teams can accurately predict application behavior, thereby fostering successful deployments in cloud-first applications.

In summary, by addressing the challenges of time, team dynamics, and resource consistency, organizations can cultivate a climate of reliability and confidence across all their development and production environments.

More Free Book



Scan to Download

Chapter 11 Summary: Administrative Processes

Administrative Processes

Overview of REPLs and Administrative Tasks

This chapter introduces the concept of REPLs, or Read-Eval-Print Loops, which are environments that facilitate interactive programming through immediate feedback. It aligns with the twelfth principle of the Twelve-Factor App methodology: "Run administrative/management tasks as one-off processes." However, this principle has faced scrutiny for potentially leading developers astray, especially in effectively managing administrative processes.

Critique of Administrative Processes

While the Twelve-Factor App approach accentuates the role of administrative tasks, it might inadvertently promote practices that are not conducive to all scenarios. Certain situations exist where the implementation of administrative processes fails, signaling the need for a reevaluation of alternatives to streamline operations.

Examples of Administrative Processes to Avoid

The chapter outlines several common administrative practices that can be problematic:



- **Database migrations** can complicate deployments.
- **Interactive programming consoles (REPLs)** may hinder the stateless nature of applications.
- **Timed scripts**, such as nightly batch jobs, often introduce unpredictability.
- **One-off jobs** that run custom code just once can lead to inefficiencies.

Challenges with Timers and Interactive Shells

Using internal timers for batch processes in cloud settings can result in chaotic behaviors, including data duplication. Interactive shells, while useful in some contexts, offer limited benefits in well-designed stateless applications, which can lead to more complications than solutions.

Proposed Solutions for Administrative Processes

To address the challenges posed by traditional administrative methods, the chapter suggests more reliable strategies. One such approach is to leverage external commands, like cron or Autosys, to trigger administrative tasks; however, their reliability in cloud environments is questionable. A more effective solution would involve exposing a RESTful endpoint for invoking ad hoc functionalities. Additionally, developing a separate microservice designated for batch operations can streamline the process further.



Advantages of Revised Approaches

Adopting these refined approaches allows for:

- Consistent, one-time execution of batch operations, preventing chaos from multiple instances.
- Enhanced security through controlled access to batch processing endpoints.
- Utilization of cloud scalability and flexibility, optimizing the performance of administrative tasks.

Conclusion and Recommendations

The chapter concludes with a strong recommendation to critically assess the role of administrative processes within one's architecture. It suggests exploring cloud-native functions, such as AWS Lambdas, for executing one-off tasks. Furthermore, it advocates for evaluating architectural modifications that can eliminate the dependency on traditional administrative processes, leading to more efficient and reliable operations overall.

More Free Book



Scan to Download

Chapter 12: Port Binding

Chapter 11: Port Binding

In Chapter 11, the focus shifts to the foundational concept of port binding in cloud-native applications, which play a crucial role in modern software development. Port binding refers to the practice of exposing application services through designated ports, a technique essential for seamless communication between applications and users.

The chapter begins by establishing that modern web applications, particularly in enterprise settings, typically operate within server containers such as Tomcat, JBoss, or IIS. These containers automatically manage port assignments at startup, a process that simplifies deployment but can sometimes lead to complexity if not managed correctly.

A common pattern observed in enterprise environments involves hosting multiple applications within a single container. By designating different port numbers for each application and utilizing DNS for user-friendly naming (e.g., routing app1 through port 8080 and app2 through port 8081), enterprises can streamline user access and enhance operational efficiency. This organizational structure allows users to interact with applications without having to memorize numerical port assignments.

More Free Book



Scan to Download

The chapter also emphasizes the benefits of adopting a Platform-as-a-Service (PaaS) model. PaaS significantly reduces the burden on developers, as the cloud provider assumes responsibility for managing critical infrastructure aspects, including port assignments, routing, scaling, availability, and fault tolerance. This allows developers to focus on building and refining applications instead of getting bogged down by infrastructure concerns.

Turning to practical guidelines, the chapter references the original 12-factor app principles, which advocate for a self-contained application structure. The recommendation here is that cloud-native applications should avoid being injected into external containers to maintain purity in their design. Nevertheless, given the complexities of existing enterprise applications, the chapter advises maintaining a 1:1 correlation between applications and their corresponding application servers whenever possible.

The implications of port binding for developers are significant. By allowing applications to operate under different URLs across various environments—such as development on localhost, quality assurance with a specific IP, and production on a corporate domain—developers can achieve greater flexibility without necessitating code changes. This adaptability is especially critical in the fast-paced development cycles common in cloud-native environments.

More Free Book



Scan to Download

Additionally, the chapter introduces the concept of externalized runtime port binding, which can serve as a backing service for other applications. This capability highlights the overall flexibility and robustness that cloud environments provide, allowing applications to interact seamlessly with one another while maintaining the advantages of cloud-native architecture.

In summary, Chapter 11 provides an insightful overview of port binding within the realm of cloud-native applications, showcasing its importance in application delivery, operational efficiency, and development flexibility.

Install Bookey App to Unlock Full Text and Audio

Free Trial with Bookey

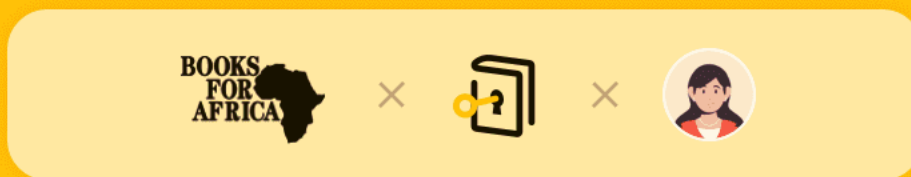




Read, Share, Empower

Finish Your Reading Challenge, Donate Books to African Children.

The Concept



This book donation activity is rolling out together with Books For Africa. We release this project because we share the same belief as BFA: For many children in Africa, the gift of books truly is a gift of hope.

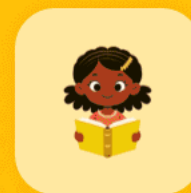
The Rule



Earn 100 points



Redeem a book



Donate to Africa

Your learning not only brings knowledge but also allows you to earn points for charitable causes! For every 100 points you earn, a book will be donated to Africa.

Free Trial with Bookey

Chapter 13 Summary: Stateless Processes

Chapter 13 Summary: Stateless Processes and the Share-Nothing Pattern

Single Stateless Process Concept

In the realm of cloud-native applications, the concept of a "single stateless process" emerges as a foundational principle. This approach emphasizes the need for applications to operate without retaining any state information between requests. While the original Twelve-Factor principles support the existence of multiple processes, the single stateless process model streamlines operations, ensuring that applications are inherently more resilient and easier to scale.

Understanding Statelessness

Statelessness is a critical characteristic of modern applications, wherein each request is treated independently, without relying on any prior interactions. This means that while applications can create temporary, transient states during processing—such as those involved in transaction handling—all enduring state information must be managed externally. For instance, databases and other backing services are essential for storing data that needs to persist beyond individual request cycles.

More Free Book



Scan to Download

The Share-Nothing Pattern

Central to maintaining a robust architecture in cloud environments is the Share-Nothing Pattern. This strategy advocates for processes to operate independently without sharing resources. Shared resources can introduce vulnerabilities and can lead to cascading failures, particularly when processes are ephemeral and can be disposed of at any time. For instance, relying on a shared filesystem as a backing service is discouraged; instead, applications should utilize external services to manage data sharing, thereby reinforcing the stateless nature of the processes.

Data Caching Best Practices

Although caching can enhance application performance by storing frequently accessed data, it must be implemented judiciously to avoid bloating and slow startup times. For long-running applications, utilizing external caching solutions, such as Redis or Gemfire, becomes a best practice. These external services help manage data effectively while allowing the application to maintain its stateless architecture, ensuring quick response times and high availability.

Overall, Chapter 13 underscores the significance of stateless processes and the Share-Nothing Pattern as essential elements for developing resilient,

More Free Book



Scan to Download

scalable cloud-native applications, ultimately guiding developers toward effective data management and system design.

More Free Book



Scan to Download

Chapter 14 Summary: Telemetry

Telemetry

Introduction to Telemetry

Though telemetry is not one of the original twelve factors of modern software development, it serves a vital function in the realm of cloud-native applications. Essentially, telemetry involves collecting and transmitting measurement data remotely, which enables developers to monitor applications operating in the cloud—situations where direct access is limited compared to traditional local environments.

Viewing Applications as Space Probes

To illustrate the importance of telemetry, the author draws an analogy between deploying cloud applications and launching scientific instruments into space. Just as telemetry is necessary to monitor a space probe, which may be light-years away, cloud applications require similar monitoring principles. Effective telemetry ensures developers have access to critical data and remote control mechanisms essential for evaluating the health and performance of their applications.

More Free Book



Scan to Download

Categories of Telemetry Data

Telemetry can be categorized into three main types:

1. Application Performance Monitoring (APM)

APM is crucial for tracking the overall performance of an application. It involves collecting performance-related events and metrics, which developers must define and standardize to ensure effective monitoring.

2. Domain-Specific Telemetry

This type collects event streams that are directly relevant to the business, providing valuable insights for analytics and reporting. Often, this data is integrated into big data systems to garner deeper business intelligence.

3. Health and System Logs

Health and system logs, typically provided by cloud service providers, document various application events, including start-ups, shut-downs, scaling actions, and routine health checks.

Challenges of Cloud Monitoring

More Free Book



Scan to Download

Monitoring cloud applications poses distinct challenges that can be more intricate than those associated with traditional enterprise applications. The sheer volume of data generated from health checks, audits, and performance metrics requires thoughtful strategies for data aggregation and storage. Without adequate planning, the risk of information overload rises, complicating the monitoring tasks.

Conclusion

In conclusion, effective telemetry is fundamental to the successful deployment of cloud applications. Similar to the mission control needed for a satellite launch, having robust monitoring capabilities is critical to ensuring the reliability and performance of applications in production. Neglecting telemetry can lead to significant operational failures, highlighting its indispensable role in maintaining application health in cloud environments.

More Free Book



Scan to Download

Chapter 15 Summary: Authentication and Authorization

Chapter 15: Authentication and Authorization

In the realm of application development, particularly within cloud environments, security emerges as a fundamental pillar rather than a mere afterthought. This chapter underscores the significance of embedding robust security measures into applications, ensuring they are designed with protection in mind from the very beginning.

The Importance of Security in Applications

Given the complexities and vulnerabilities inherent in cloud infrastructures, security must be prioritized throughout the development lifecycle. It is crucial to recognize that rushing to meet functional requirements without adequate security measures can expose both the application and its users to significant risks.

The Necessity of Secure Code

Cloud-native applications, which operate across diverse data centers and interact with multiple clients, rely heavily on secure coding practices. This proactive approach not only protects sensitive data but also facilitates the

More Free Book



Scan to Download

establishment of an audit trail. Such a trail is vital for tracing user actions and monitoring data modifications, thereby enhancing accountability.

Role-Based Access Control (RBAC)

One effective method for safeguarding cloud-native applications is through Role-Based Access Control (RBAC). This framework necessitates that each request for access to application resources is accompanied by an identification of the requester and an understanding of their designated roles. These roles define the level of permissions granted, ensuring that users can only access information and perform actions aligned with their responsibilities.

Integration of Security in Development

To achieve optimal security, it is crucial to weave security protocols into the development process from the outset. Leveraging authentication and authorization frameworks such as OAuth2 and OpenID Connect is essential for this integration. By treating security as a foundational element rather than an afterthought, developers can create more resilient applications that protect user data and maintain trust.

Overall, Chapter 15 emphasizes that security should be an integral part of cloud-native application development. By adopting RBAC and embedding

More Free Book



Scan to Download

security measures early in the process, developers can safeguard their applications while ensuring compliance and accountability.

More Free Book



Scan to Download

Chapter 16: A Word on Cloud Native

A Word on Cloud Native - Summary

What Is Cloud Native?

The term "cloud native" is often misinterpreted, particularly when confused with cloud computing. To clarify, "cloud" in this context typically refers to Platform as a Service (PaaS) solutions like Google App Engine and Heroku. These platforms abstract the complexities of infrastructure, empowering developers to focus on crafting applications without needing to manage hardware specifics.

Importantly, being cloud-native is not limited to new applications built exclusively for the cloud. Instead, it pertains to applications designed to operate efficiently within a PaaS environment, embracing features such as horizontal elastic scaling. However, as definitions of cloud-native evolve, a significant debate persists between purists—who adhere strictly to ideal principles—and pragmatists—who recognize the need for flexibility and adaptation in real-world scenarios.

Why Cloud Native?

More Free Book



Scan to Download

Historically, applications were created for physical servers, which frequently led to limitations in scalability and susceptibility to hardware failures. The advent of virtualization responded to these challenges, enabling easier deployment and facilitating horizontal scaling. Today's competitive landscape requires businesses across various industries to deliver reliable software rapidly, making software development a critical focus.

Adopting a cloud-native architecture allows companies to channel their efforts into their core competencies while leveraging cloud services to manage the infrastructure. In this evolving landscape, organizations should be considering why they haven't yet embraced cloud-native practices rather than questioning their inherent value.

The Purist vs. the Pragmatist

Within the realm of software development, there exists a dynamic tension between adhering to lofty ideals and addressing the practical constraints found in real projects. This chapter underscores the need to differentiate between non-negotiable principles and areas where some flexibility can be applied to meet specific project requirements. Striking a compromise between purism and pragmatism is vital for the successful delivery of cloud-native applications, ensuring they are completed on time and within budget. Understanding when and how to navigate this balance will greatly contribute to effective implementation and ongoing success in the

More Free Book



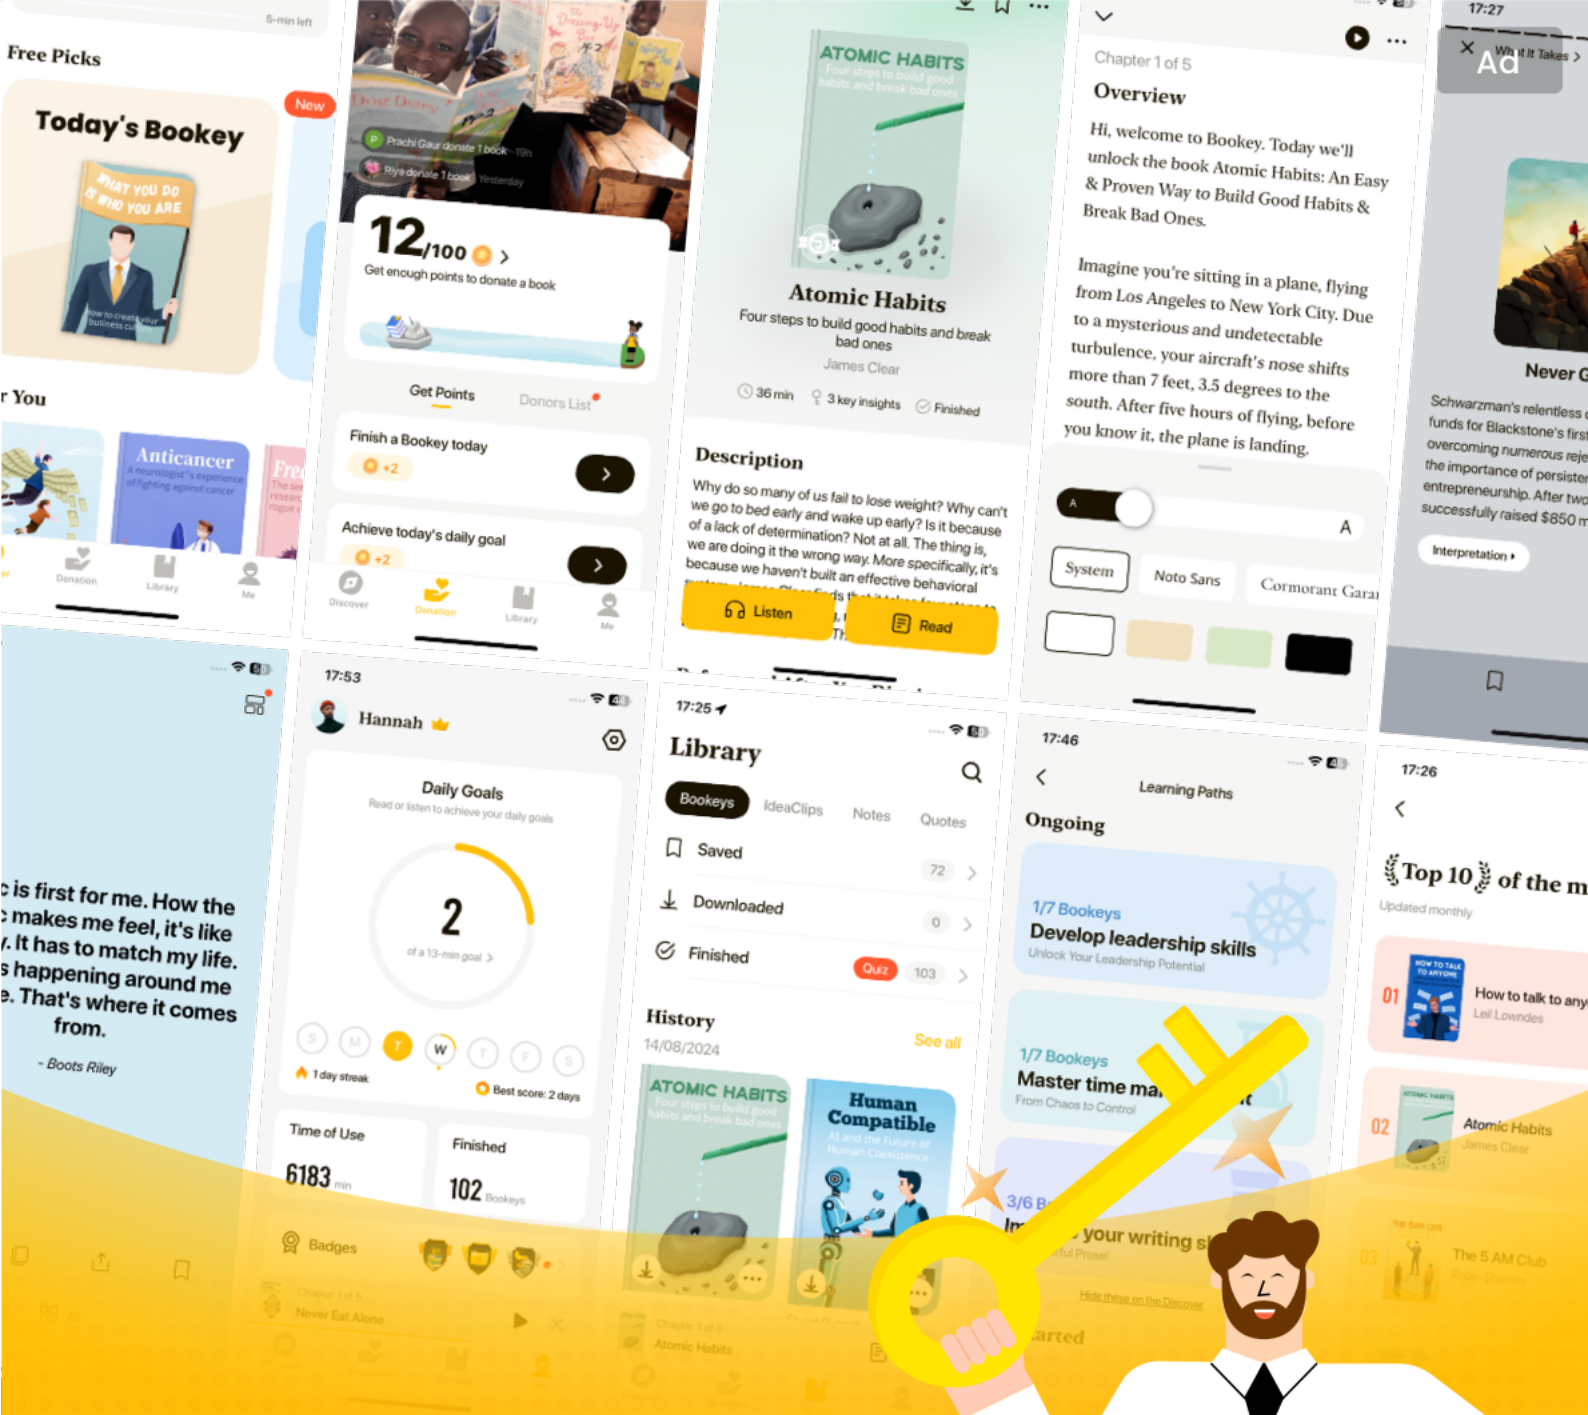
Scan to Download

cloud-native landscape.

Install Bookey App to Unlock Full Text and Audio

Free Trial with Bookey





World' best ideas unlock your potential

Free Trial with Bookey



Scan to download

