

# Angular Ivy: Revolutionizing Rendering in Angular Applications

Nikhil Kodali

Software Engineer, Tennessee Valley Authority, Chattanooga, TN

## Abstract

This paper examines the Angular introduced Ivy, a groundbreaking rendering engine that marked a significant milestone in the framework's evolution. Ivy enhances application performance and reduces bundle sizes by enabling more efficient tree-shaking and code generation. Its architecture facilitates faster compilation and improves debugging capabilities, offering clearer error messages and more straightforward code inspection. By introducing a new approach to handling templates and components, Ivy makes them more flexible and easier to understand. Additionally, it supports features like local and global components and improved dynamic component loading. This paper explores the innovations Ivy brings to Angular development, its impact on performance, and how it enhances the developer experience.

**Keywords:** Angular Ivy, Rendering Engine, Tree-Shaking, Code Optimization, Component Compilation.

## 1. Introduction

Angular has been a prominent framework for building dynamic web applications since its inception. With the web development landscape continually evolving, performance optimization and developer experience have become paramount. In response to these demands, the Angular team introduced Ivy, a next-generation rendering engine designed to address the limitations of its predecessor, View Engine.

Angular introduced Ivy, a next-generation rendering engine that revolutionized the development of web applications by enhancing performance, optimizing bundle sizes, and improving the overall developer experience. Since its inception, Angular has been a prominent framework for building dynamic single-page applications (SPAs), offering powerful tools for component-based development, dependency injection, and reactive programming. However, as the web development landscape evolved, performance optimization, smaller bundle sizes, and an improved developer experience became paramount requirements for modern applications. To address these needs, Angular's core team developed Ivy, a rendering engine designed to overcome the limitations of its predecessor, the View Engine, and usher in a new era of efficiency and flexibility in Angular development.

 [CC BY 4.0 Deed Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/)

This article is distributed under the terms of the Creative Commons CC BY 4.0 Deed Attribution 4.0 International attribution which permits copy, redistribute, remix, transform, and build upon the material in any medium or format for any purpose, even commercially without further permission provided the original work is attributed as specified on the Ninety Nine Publication and Open Access pages <https://turcomat.org>

Before Ivy, Angular used the View Engine for rendering components and templates. While the View Engine was functional, it had several limitations that affected the performance and scalability of applications. One of the main issues was the inability to effectively eliminate unused code, also known as tree-shaking. Inefficient tree-shaking led to larger bundle sizes, which in turn affected the load time and performance of applications. Additionally, the compilation process with the View Engine was complex and time-consuming, often requiring developers to perform multiple steps to compile templates and components. Debugging was also a challenge, as the error messages generated by the View Engine were often cryptic and difficult to interpret, making the debugging process cumbersome.

The introduction of Ivy aimed to address these challenges by fundamentally rethinking the rendering architecture of Angular. One of Ivy's core innovations is its efficient tree-shaking capabilities. Tree-shaking is the process of eliminating unused code from the final JavaScript bundle, which helps reduce the size of the application and improve load times. Ivy generates code that is more tree-shakable, ensuring that only the code used in the application is included in the final bundle. This results in significantly smaller bundle sizes, making applications faster and more efficient. Additionally, Ivy uses the "locality principle" for code generation, which means that each component is compiled independently of its dependencies. This approach simplifies the compilation process, reduces interdependencies between components, and makes the overall build process more efficient.

Another key feature of Ivy is its faster and more incremental compilation process. Unlike the View Engine, which required a two-phase compilation process involving both template and application compilation, Ivy compiles each component individually, reducing duplication and speeding up the build process. This incremental compilation approach is particularly beneficial during the development phase, as it allows developers to quickly see the effects of their changes without needing to recompile the entire application. Ivy also improves Ahead-of-Time (AOT) compilation, making it viable for development builds and providing faster load times and earlier error detection.

Ivy also brings significant improvements to the debugging process, which enhances the overall developer experience. One of the most notable changes is the introduction of more descriptive error messages, which help developers quickly identify and fix issues in their code. Additionally, Ivy provides better mapping between templates and generated code, allowing developers to easily inspect and debug the relationships between templates and their corresponding components. This improved debugging capability not only makes it easier for developers to troubleshoot issues but also reduces the time and effort required to maintain and update applications.

The new approach to handling templates and components is another major advancement introduced by Ivy. Ivy simplifies the syntax for writing components, making it more straightforward and easier to understand. This simplification not only makes it easier for new developers to learn Angular but also allows experienced developers to write cleaner and more maintainable code. Ivy also introduces improved support for dynamic component loading, allowing developers to load components at runtime more efficiently. This feature is particularly useful for applications that require dynamic content, such as dashboards, where

components need to be loaded based on user interactions or data changes. Additionally, Ivy supports local components, which can be declared and used within the same file without the need for NgModules, further simplifying component management and reducing boilerplate code.

The impact of Ivy on Angular development has been profound, particularly in terms of performance and developer productivity. One of the most significant benefits of Ivy is the reduction in bundle sizes, which directly translates to faster load times and improved application performance. In real-world applications, Ivy has been reported to reduce bundle sizes by up to 40%, resulting in faster startup times and a better user experience. The optimized rendering pipeline and more efficient change detection algorithm used by Ivy also contribute to improved runtime performance, reducing the memory footprint of applications and ensuring smoother interactions.

Another important aspect of Ivy is its compatibility with existing Angular applications. Ivy was designed to be backward compatible, allowing developers to migrate to the new rendering engine without making significant changes to their existing codebases. This ease of migration has been facilitated by tools and guides provided by the Angular team, which help developers transition to Ivy while maintaining the stability and functionality of their applications. The backward compatibility of Ivy has been crucial in ensuring widespread adoption within the Angular community, as it allows developers to take advantage of the new features and performance improvements without the risk of breaking their existing projects.

The introduction of Ivy has also enhanced the development workflow by providing faster builds and improved tooling support. The incremental compilation process reduces build times during development, allowing developers to iterate more quickly and focus on building features rather than waiting for long compilation times. The simplified APIs and better integration with development tools and editors have further improved the productivity of Angular developers, making it easier to create, test, and maintain components. The positive reception of Ivy within the developer community is a testament to its effectiveness in addressing the key challenges of performance, bundle size, and developer experience.

Despite the numerous benefits introduced by Ivy, there are also challenges that developers need to consider. One of the primary challenges is compatibility with third-party libraries. Some libraries needed updates to be fully compatible with Ivy, and developers may need to ensure that the libraries they are using have been updated to work with the new rendering engine. Additionally, testing existing applications may require adjustments, as the changes in component rendering introduced by Ivy can affect the way components are tested. Developers may also need to update their development environments and CI/CD pipelines to accommodate the changes introduced by Ivy.

### **Problem Statement**

The introduction of Angular Ivy aimed to address the limitations of the previous View Engine by enhancing application performance, reducing bundle sizes, and improving the developer experience. However, challenges such as compatibility with third-party libraries, testing adjustments, and development environment updates need to be addressed to fully realize the

benefits of Ivy. This study seeks to explore the innovations introduced by Angular Ivy, focusing on its impact on performance, development workflow, and the strategies required to overcome the associated challenges.

---

## 2. Methodology

The methodology for this study on Angular Ivy involved a combination of literature review, experimental implementation, and performance benchmarking. This multi-phase approach allowed for a comprehensive understanding of the innovations introduced by Ivy, its benefits over the previous rendering engine, and its impact on Angular development.

The literature review phase focused on analyzing official Angular documentation, industry publications, and academic articles to understand the motivations behind the introduction of Ivy and its intended impact on web development. This phase also involved examining the limitations of the previous View Engine and how Ivy addresses these challenges. The literature review provided a theoretical foundation for understanding the architectural changes introduced by Ivy and their implications for performance, bundle size, and developer experience.

The experimental implementation phase involved creating sample applications using both the View Engine and Ivy to compare their performance and development workflows. This phase aimed to explore the practical aspects of using Ivy, including the process of transforming existing components, optimizing bundle sizes, and managing dependencies. By building applications with both rendering engines, the study aimed to identify the specific areas where Ivy provided significant improvements and to document best practices for migrating to Ivy.

The performance benchmarking phase involved measuring key metrics such as bundle size, load time, and memory usage to assess the efficiency of Ivy compared to the View Engine. Tools like Chrome DevTools, Lighthouse, and Angular CLI's build analysis features were used to collect data on the performance of applications built with Ivy. The benchmarking focused on evaluating the impact of Ivy on both development builds and production builds, highlighting the improvements in build times, runtime performance, and resource utilization.

By combining insights from the literature review, experimental implementation, and performance benchmarking, the study aimed to provide a comprehensive evaluation of Angular Ivy's impact on web development. This multi-phase methodology allowed for a balanced analysis of both the theoretical and practical aspects of Ivy, highlighting the opportunities and challenges associated with adopting this next-generation rendering engine.

### 2.1. The Need for a New Rendering Engine

Before Ivy, Angular used the View Engine for rendering components and templates. While functional, the View Engine had limitations:

- **Large Bundle Sizes:** Inefficient tree-shaking led to larger bundles, affecting load times.
- **Complex Compilation:** The compilation process was slower and less efficient.

- **Limited Flexibility:** Handling of templates and components lacked the desired flexibility.
- **Debugging Challenges:** Error messages were often cryptic, making debugging difficult.

## 2.2. Introduction to Angular Ivy

Ivy was introduced to overcome these challenges by:

- **Optimizing Bundle Sizes:** Leveraging more efficient tree-shaking.
- **Enhancing Compilation:** Providing faster and more incremental compilation.
- **Improving Debugging:** Offering better error messages and code inspection.
- **Simplifying Components:** Introducing a new way of handling templates and components.

---

## 3. Architectural Innovations in Ivy

### 3.1. Efficient Tree-Shaking and Code Generation

#### Tree-Shaking:

- **Definition:** The process of eliminating unused code from the final bundle.
- **Ivy's Approach:** Generates code that is more tree-shakable, ensuring that only the code used in the application is included.

#### Code Generation:

- **Locality Principle:** Ivy compiles components independently, meaning each component's compilation depends only on its own code, not on that of its dependencies.
- **Benefits:**
  - Reduces interdependencies between components.
  - Simplifies the compilation process.

### 3.2. Faster Compilation

- **Incremental Compilation:** Ivy compiles only the components that have changed, speeding up the build process.
- **Ahead-of-Time (AOT) Compilation:** Improved AOT performance makes it viable for development builds, offering faster load times and earlier error detection.

### 3.3. Improved Debugging Capabilities

- **Enhanced Error Messages:** Ivy provides more descriptive error messages, helping developers quickly identify and fix issues.
- **Template Debugging:** Better mapping between templates and generated code allows for easier inspection and debugging.

### 3.4. New Template and Component Handling

- **Simplified Syntax:** Ivy introduces a more straightforward way to write and understand component code.
- **Flexibility:**
  - **Dynamic Component Loading:** Improved support for loading components at runtime.
  - **Local Components:** Components can be declared and used locally without the need for NgModules.

---

## 4. Key Features of Angular Ivy

### 4.1. Smaller Bundle Sizes

- **Per-Component Compilation:** Only the necessary code for each component is included.
- **Elimination of Unused Code:** Unreferenced components and modules are excluded from the bundle.

### 4.2. Improved Performance

- **Runtime Efficiency:** Faster rendering and change detection due to optimized code.
- **Memory Usage:** Reduced memory footprint during application execution.

### 4.3. Compatibility and Migration

- **Backward Compatibility:** Ivy is designed to be compatible with existing Angular applications.
- **Ease of Migration:** Tools and guides are provided to help developers transition to Ivy without significant code changes.

### 4.4. Enhanced Developer Experience

- **Simplified APIs:** Cleaner and more intuitive APIs for component and directive creation.
- **Better Tooling Support:** Improved integration with development tools and editors.

---

## 5. Technical Deep Dive

## 5.1. Compilation Process in Ivy

### Comparison with View Engine:

- **View Engine:** Used a two-phase compilation (template and application), leading to duplication.
- **Ivy:** Compiles components individually, reducing duplication and improving build times.

### Incremental Builds:

- **Benefit:** Only modified components are recompiled, significantly speeding up development cycles.

## 5.2. Rendering and Change Detection

### Rendering Pipeline:

- **Ivy's Renderer:** More lightweight and efficient, with less overhead.
- **Change Detection:** Uses a more efficient algorithm to detect and render changes.

## 5.3. Dynamic Components and Lazy Loading

### Dynamic Components:

- **Improved API:** Simplifies the creation and insertion of components at runtime.

### Lazy Loading:

- **Optimized Loading:** Components can be loaded on demand, reducing initial load times.

---

## 6. Impact on Angular Development

### 6.1. Performance Gains

- **Real-World Applications:** Reports indicate bundle size reductions of up to 40%.
- **Startup Times:** Faster application startup due to smaller bundles and efficient rendering.

### 6.2. Development Workflow

- **Faster Builds:** Incremental compilation reduces build times during development.
- **Easier Debugging:** Enhanced error messages and debugging tools improve productivity.

### 6.3. Adoption and Community Response

- **Positive Reception:** Developers appreciate the performance improvements and better tooling.

- **Learning Curve:** Minimal, as Ivy maintains compatibility with existing Angular syntax and patterns.
- 

## 7. Challenges and Considerations

### 7.1. Compatibility Issues

- **Third-Party Libraries:** Some libraries needed updates to be fully compatible with Ivy.
- **Testing:** Existing tests might require adjustments due to changes in component rendering.

### 7.2. Migration Efforts

- **Potential Refactoring:** Although designed to be backward compatible, some applications may require minor code changes.
  - **Tooling Updates:** Development environments and CI/CD pipelines might need configuration updates.
- 

## 8. Future Developments

### 8.1. Continued Optimization

- **Further Bundle Size Reduction:** Ongoing efforts to make applications even leaner.
- **Enhanced Features:** New capabilities leveraging Ivy's flexible architecture.

### 8.2. Integration with Other Technologies

- **Web Components:** Improved support for integrating Angular components with web components.
  - **Server-Side Rendering:** Better performance and ease of use with Angular Universal.
- 

## 9. Conclusion

Angular Ivy represents a significant advancement in Angular's evolution, addressing critical performance and developer experience challenges. By overhauling the rendering engine, Angular has provided developers with a tool that not only enhances application performance but also simplifies development and debugging. Ivy's introduction marked a new era for Angular applications, positioning the framework to meet the demands of modern web development.

---



## References

- 1) Perez, I., Bärenz, M., & Nilsson, H. (2016). "Functional reactive programming, refactored." In *Haskell*. ACM, 33–44.
- 2) Felleisen, M., Findler, R. B., Flatt, M., & Krishnamurthi, S. (2009). "A functional I/O system or, fun for freshman kids." In *ICFP*. ACM, 47–58.
- 3) Breitner, J., & Smith, C. (2017). "Lock-step simulation is child's play (experience report)." *PACMPL*, 1(ICFP), 3:1–3:15.
- 4) Almeida, J., Cunha, A., Macedo, N., Pacheco, H., & Proença, J. (2018). "Teaching how to program using automated assessment and functional Glossy games (experience report)." *PACMPL*, 2(ICFP), 82:1–82:17.
- 5) Peterson, J., Hudak, P., & Elliott, C. (1999). "Lambda in motion: Controlling robots with Haskell." In *PADL*, LNCS, vol. 1551, Springer, 91–105.
- 6) Hudak, P., Courtney, A., Nilsson, H., & Peterson, J. (2002). "Arrows, robots, and functional reactive programming." In *Advanced Functional Programming*, LNCS, vol. 2638, Springer, 159–187.
- 7) Pembeci, I., Nilsson, H., & Hager, G. D. (2002). "Functional reactive robotics: An exercise in principled integration of domain-specific languages." In *PPDP*. ACM, 168–179.
- 8) Cowley, A., & Taylor, C. J. (2011). "Stream-oriented robotics programming: The design of ROSHASK." In *IROS*. IEEE, 1048–1054.