

The background of the slide features a dark blue grid with various financial data visualizations. On the left, there's a candlestick chart with a price change of -1.98. In the center, a black pen nib is positioned as if about to write. To the right, there's a line graph showing a fluctuating trend. The overall aesthetic is professional and data-oriented.

Lessons Learned Implementing StreamBase

Matt Fowles

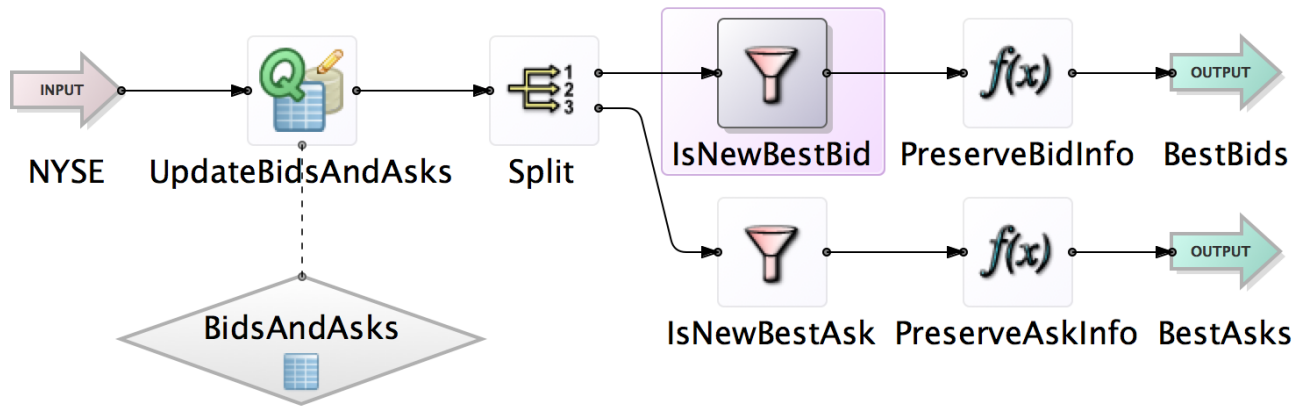
StreamBase Systems

JVM Languages 2012

Agenda

- What is StreamBase?
- Lessons Learned
 - Spaghetti code is bad; Spaghetti code generation is worse
 - Choose between Generated Code vs Runtime Library
 - Parallel APIs for generating code and for runtime code
 - Type Specialization makes code fast
 - Separate Compilation requires planning
 - Classloaders hurt the unwary
- Case Studies
 - Inherited control flow enables Disruptors
 - Pluggable table implementations and Client API Access
 - JDBC
- Conclusions and Future Work

StreamBase EventFlow



Editor | Definitions | Parameters | Dynamic Variables | **Annotations** | Metadata

StreamBase Properties (IsNewBestBid - Filter operator) | Typecheck Errors | Problems | StreamBase Extension Points | Console

Specify at least one predicate by adding and completing rows in the table below.

General | **Predicate Settings** | Concurrency

Create output port for non-matching tuples

Predicates:

Output Port	Predicate
1	best_bid == bid_price

Streams | Functions | Expression QuickRef

Input | Output

- input1 (6 fields)
 - time_int int
 - symbol string
 - bid_price double
 - ask_price double
 - best_bid double
 - best_ask double

Problem Domain for Event Processing

- Interoperate with many external systems
 - Java, C++ Libraries
 - External Systems (Reuters, 29West)
- Large Heaps (40-120 GB)
- Low Latency (20-2000 microseconds)
- Predictable performance and efficiency
- Easily Shardable
- Rapid Development
- 6+ years of experience:
 - extendability
 - debugability
 - optimizability
 - refactorability
 - understandability

Lessons Learned

- **Spaghetti code is bad; spaghetti code generation is worse**
- Choose between generated code vs runtime library
- Parallel APIs for generating code and for runtime code
- Type specialization makes code fast
- Separate compilation requires planning
- Classloaders hurt the unwary
- Improve tooling

Lessons Learned

- Spaghetti code is bad; spaghetti code generation is worse
- **Choose between generated code vs runtime library**
- Parallel APIs for generating code and for runtime code
- Type specialization makes code fast
- Separate compilation requires planning
- Classloaders hurt the unwary
- Improve tooling

Generated Code

vs

Runtime Library

"passive voice"

"active voice"

- Eliminate virtual dispatch
- Allows type specialization
- Monomorphic call sites
- Extra level of indirection
- Data can get big
- Code can get big too

- Easier to debug
- Easier to read
- Easier to test
- Faster to code
- Faster start-up
- Polymorphic call sites
- Less indirection

General Approach

- Abstract classes provide code flow within operators
- Generate type specializations
 - Heroic downcast
 - Cache according to type signature
- Generate mathematical expressions
- Generate application level control flow


```
abstract void fillPrimaryKeyFromDataClass(DiskPackage dp, DC dc);
abstract void fillDataClassFromValue(DC dc, DiskPackage sdp);
abstract DC createDataClass();
```

```
public DC getPrimaryImpl(DC pkey) {
    TableTransaction txn = null; Cursor cursor = null;
    try {
        txn = beginTransaction(); cursor = openCursor(txn);

        DiskPackage pack = new DiskPackage();
        fillPrimaryKeyFromDataClass(pack, pkey);
        cursor.getSearchKey(pack);

        if (pack.hasValue()) {
            DC dc = createDataClass();
            fillDataClassFromValue(dc, pack);
            return dc;
        } else {
            return null;
        }
    } catch (DatabaseException de) {
        throw createEvalException(de);
    } finally {
        finishTransaction(txn, cursor);
    }
}
```

Lessons Learned

- Spaghetti code is bad; spaghetti code generation is worse
- Choose between generated code vs runtime library
- **Parallel APIs for generating code and for runtime code**
- Type specialization makes code fast
- Separate compilation requires planning
- Classloaders hurt the unwary
- Improve tooling

Active voice

```
/**
 * A collection of bit flipping operations, with useful names.
 */
public final class BitOps {
    //...
    public static int roundToWord(int i) {
        return (i + 3) & ~3;
    }
    //...
}
```


Lessons Learned

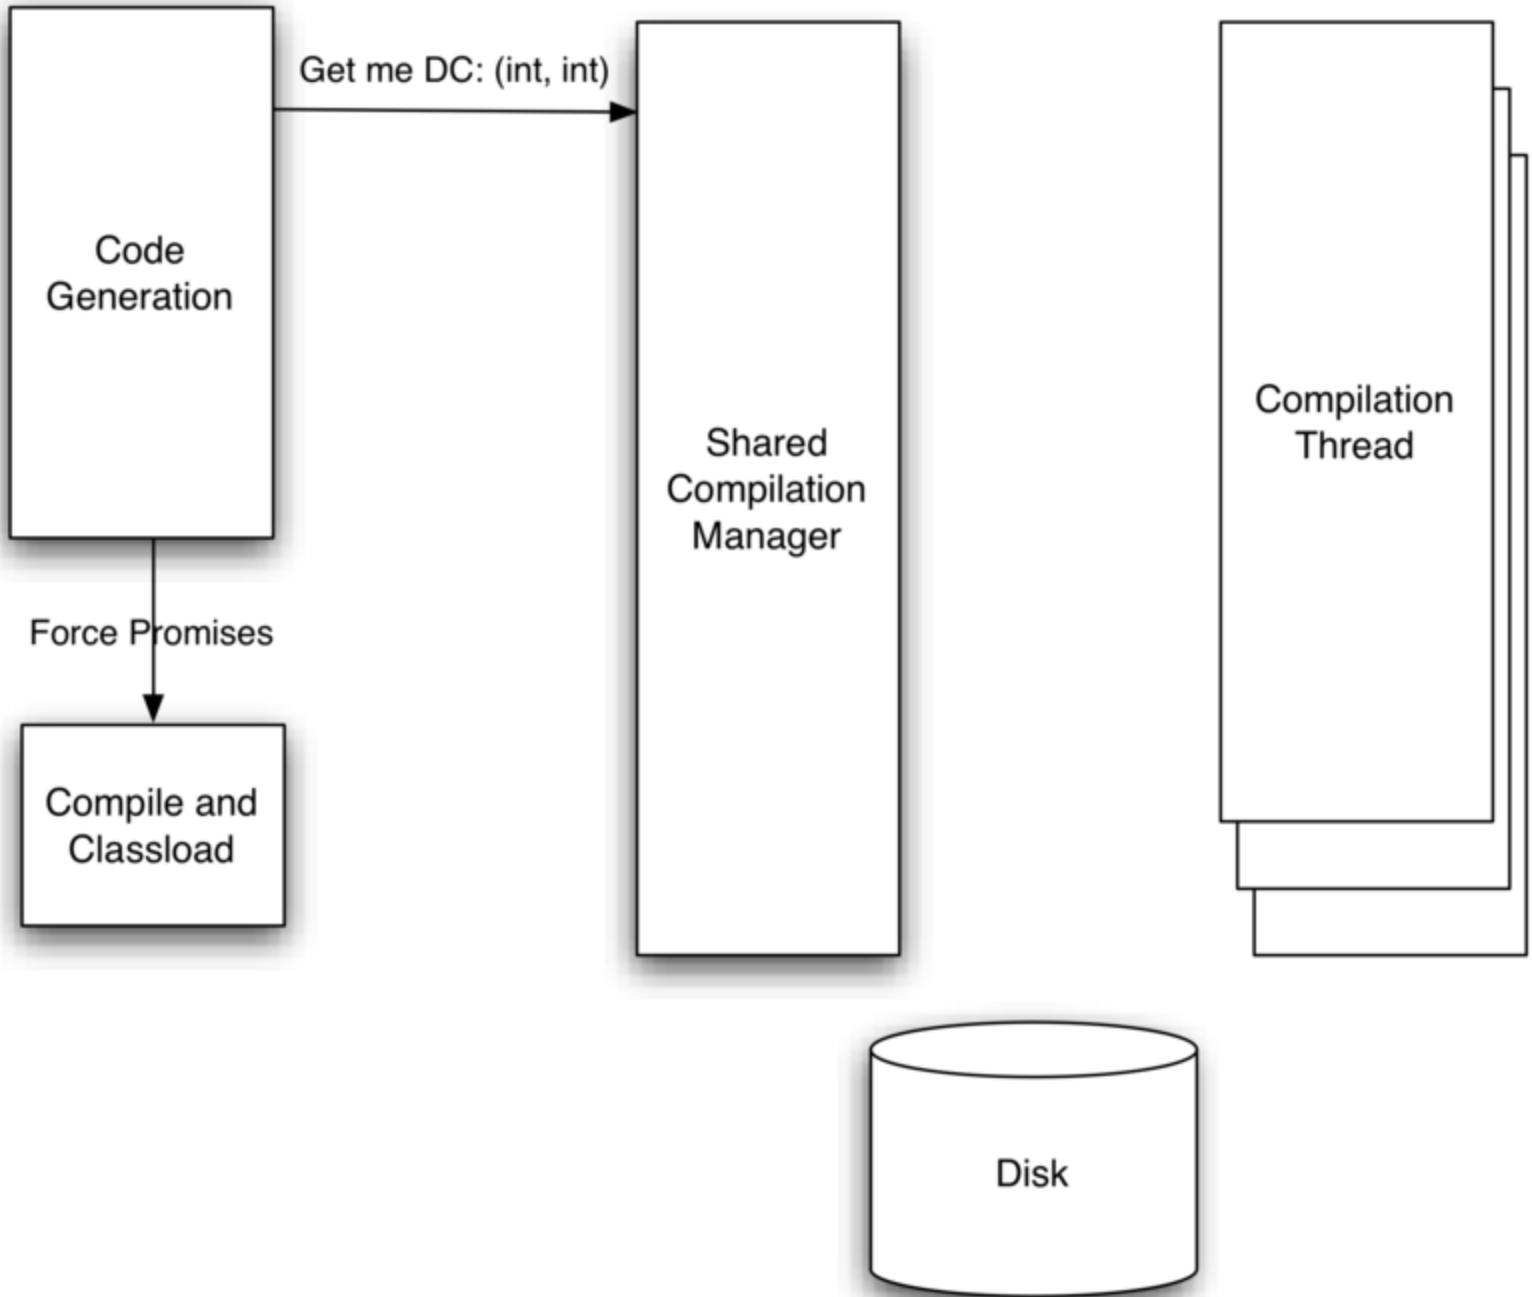
- Spaghetti code is bad; spaghetti code generation is worse
- Choose between generated code vs runtime library
- Parallel APIs for generating code and for runtime code
- **Type specialization makes code fast**
- Separate compilation requires planning
- Classloaders hurt the unwary
- Improve tooling

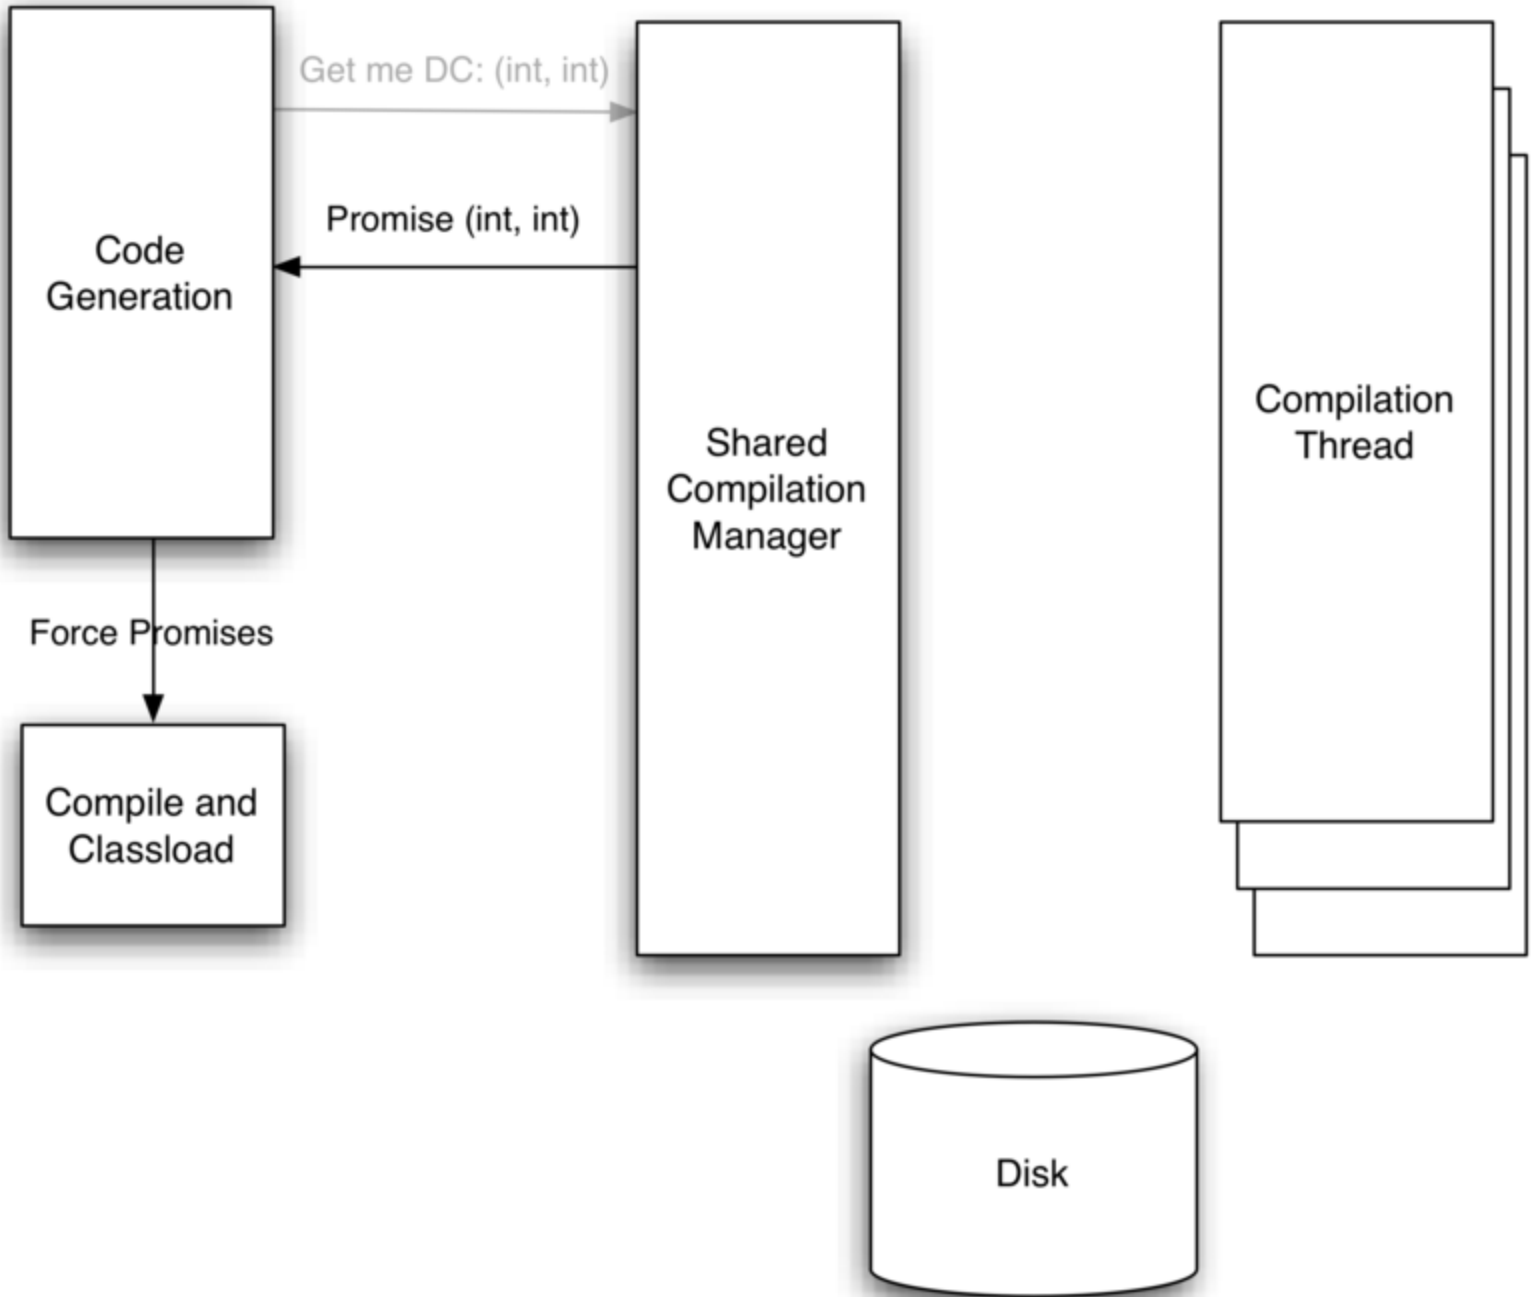
Type Specialization

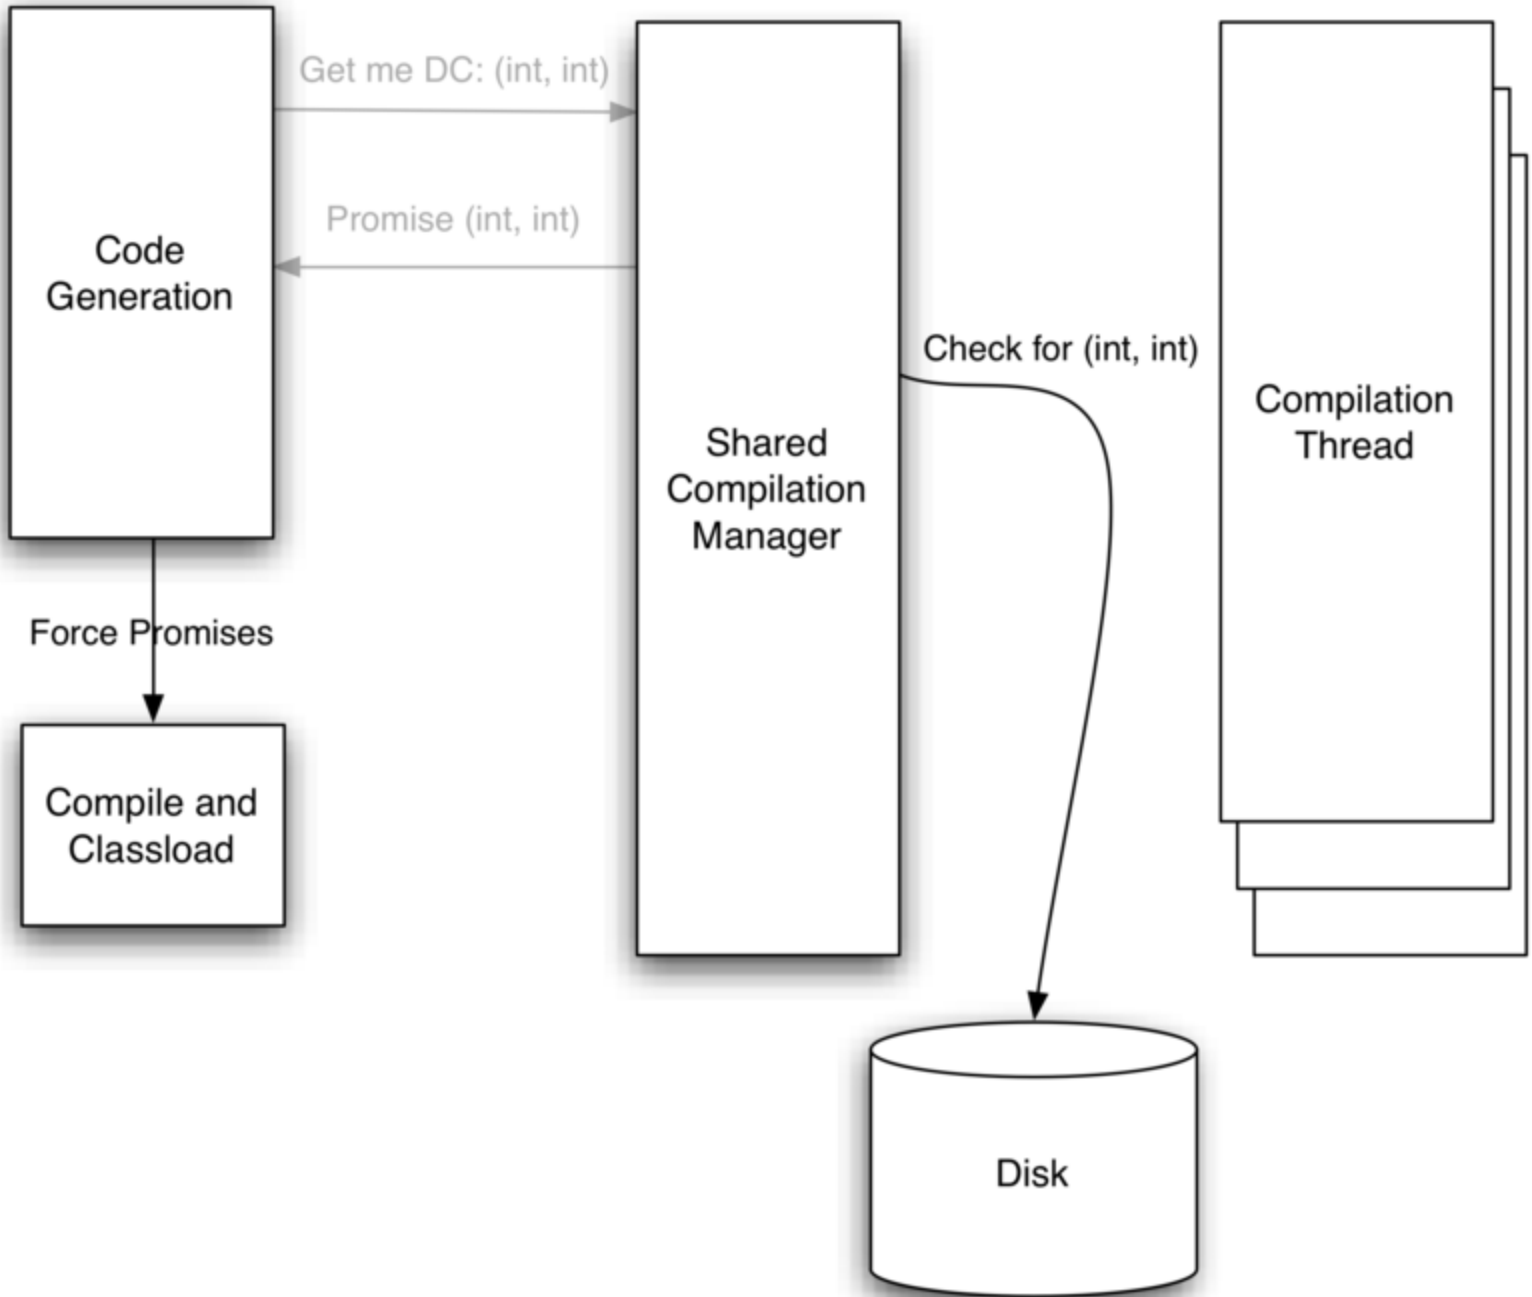
- Monomorphic call sites
- Unrolled loops
- Avoids garbage
- Lengthens compile time
- Can lead to massive methods

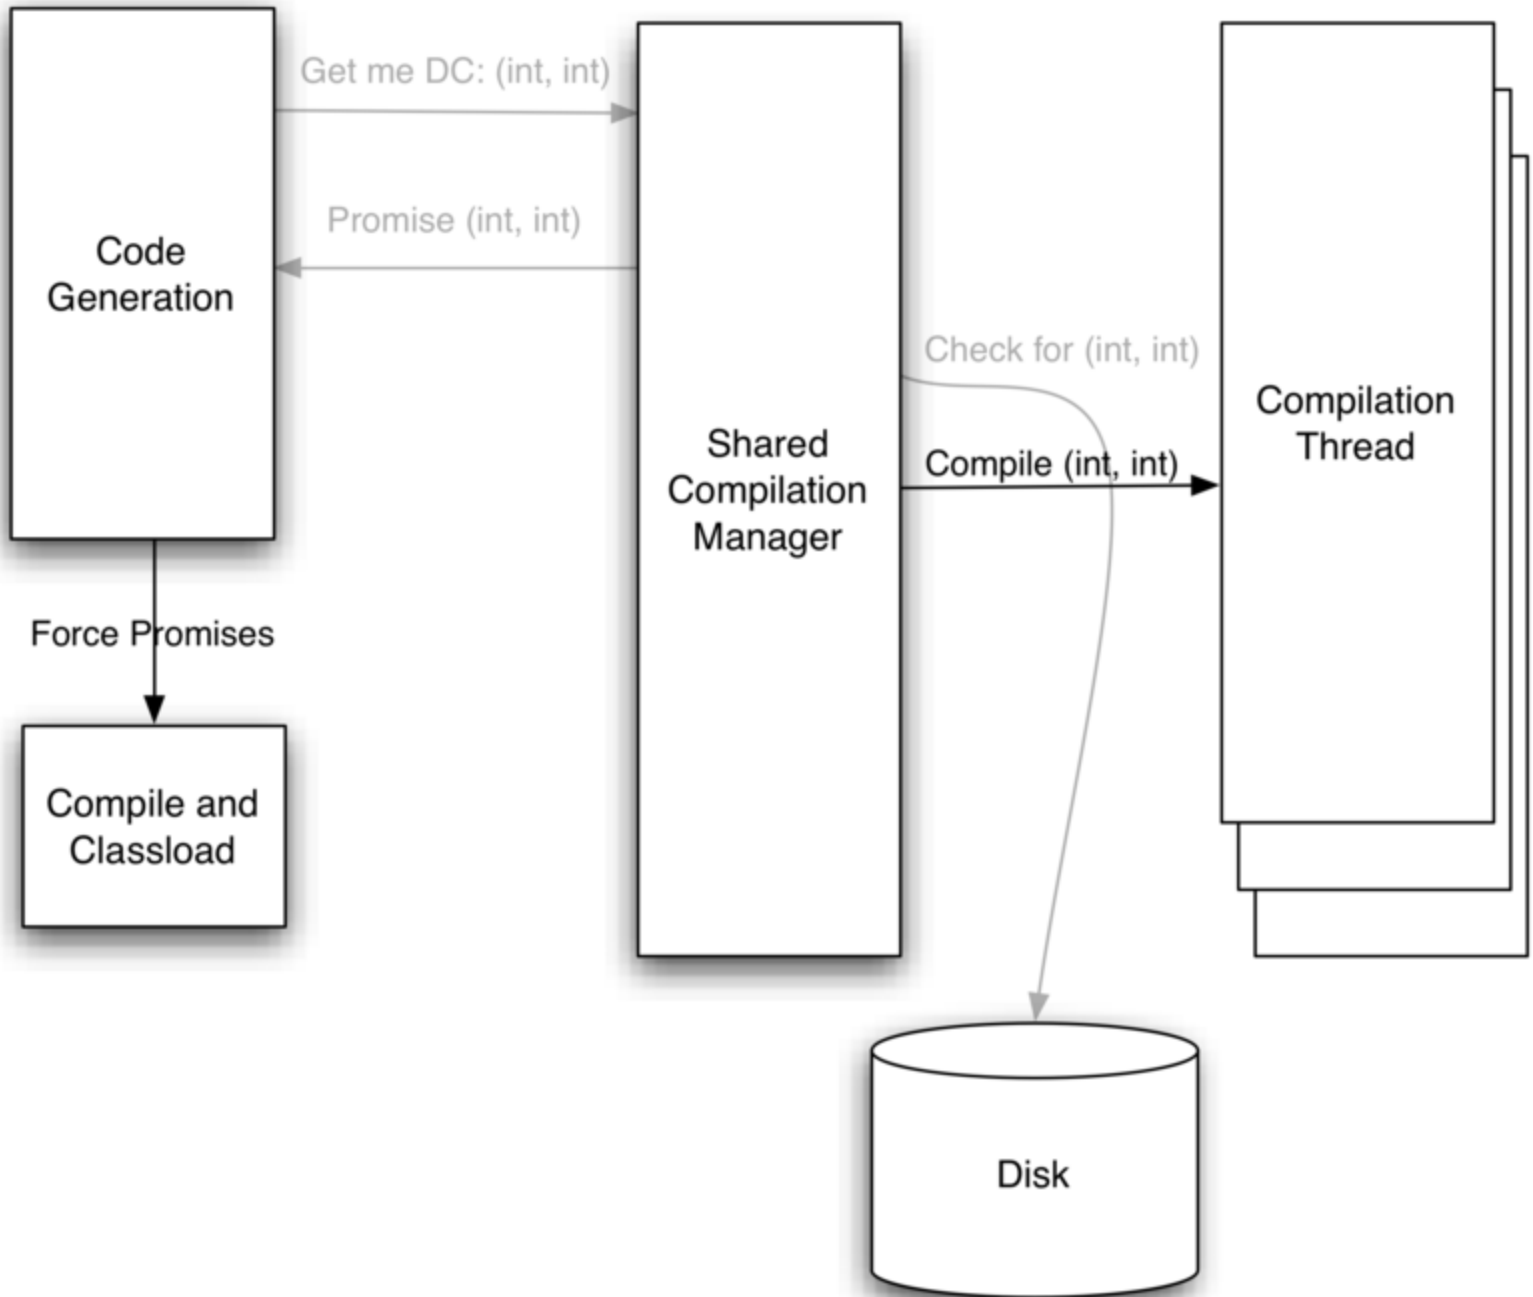
Lessons Learned

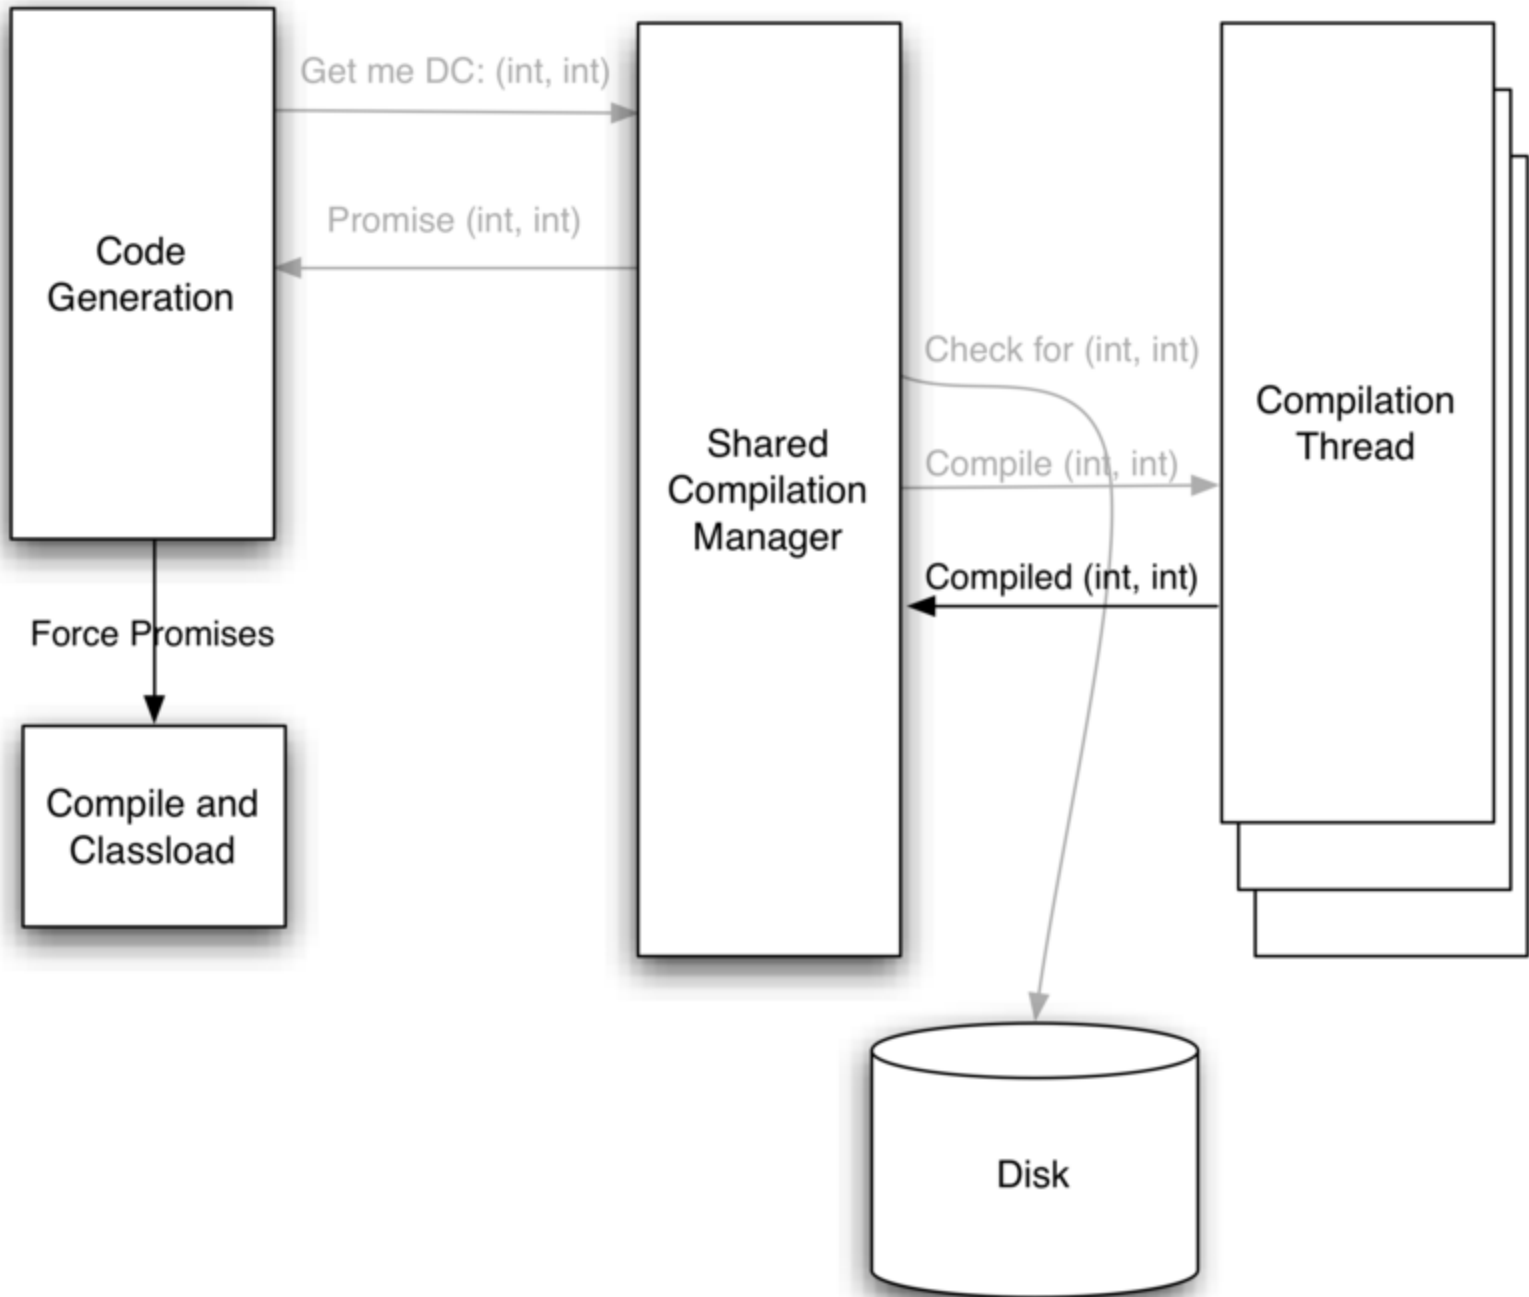
- Spaghetti code is bad; spaghetti code generation is worse
- Choose between generated code vs runtime library
- Parallel APIs for generating code and for runtime code
- Type specialization makes code fast
- **Separate compilation requires planning**
- Classloaders hurt the unwary
- Improve tooling

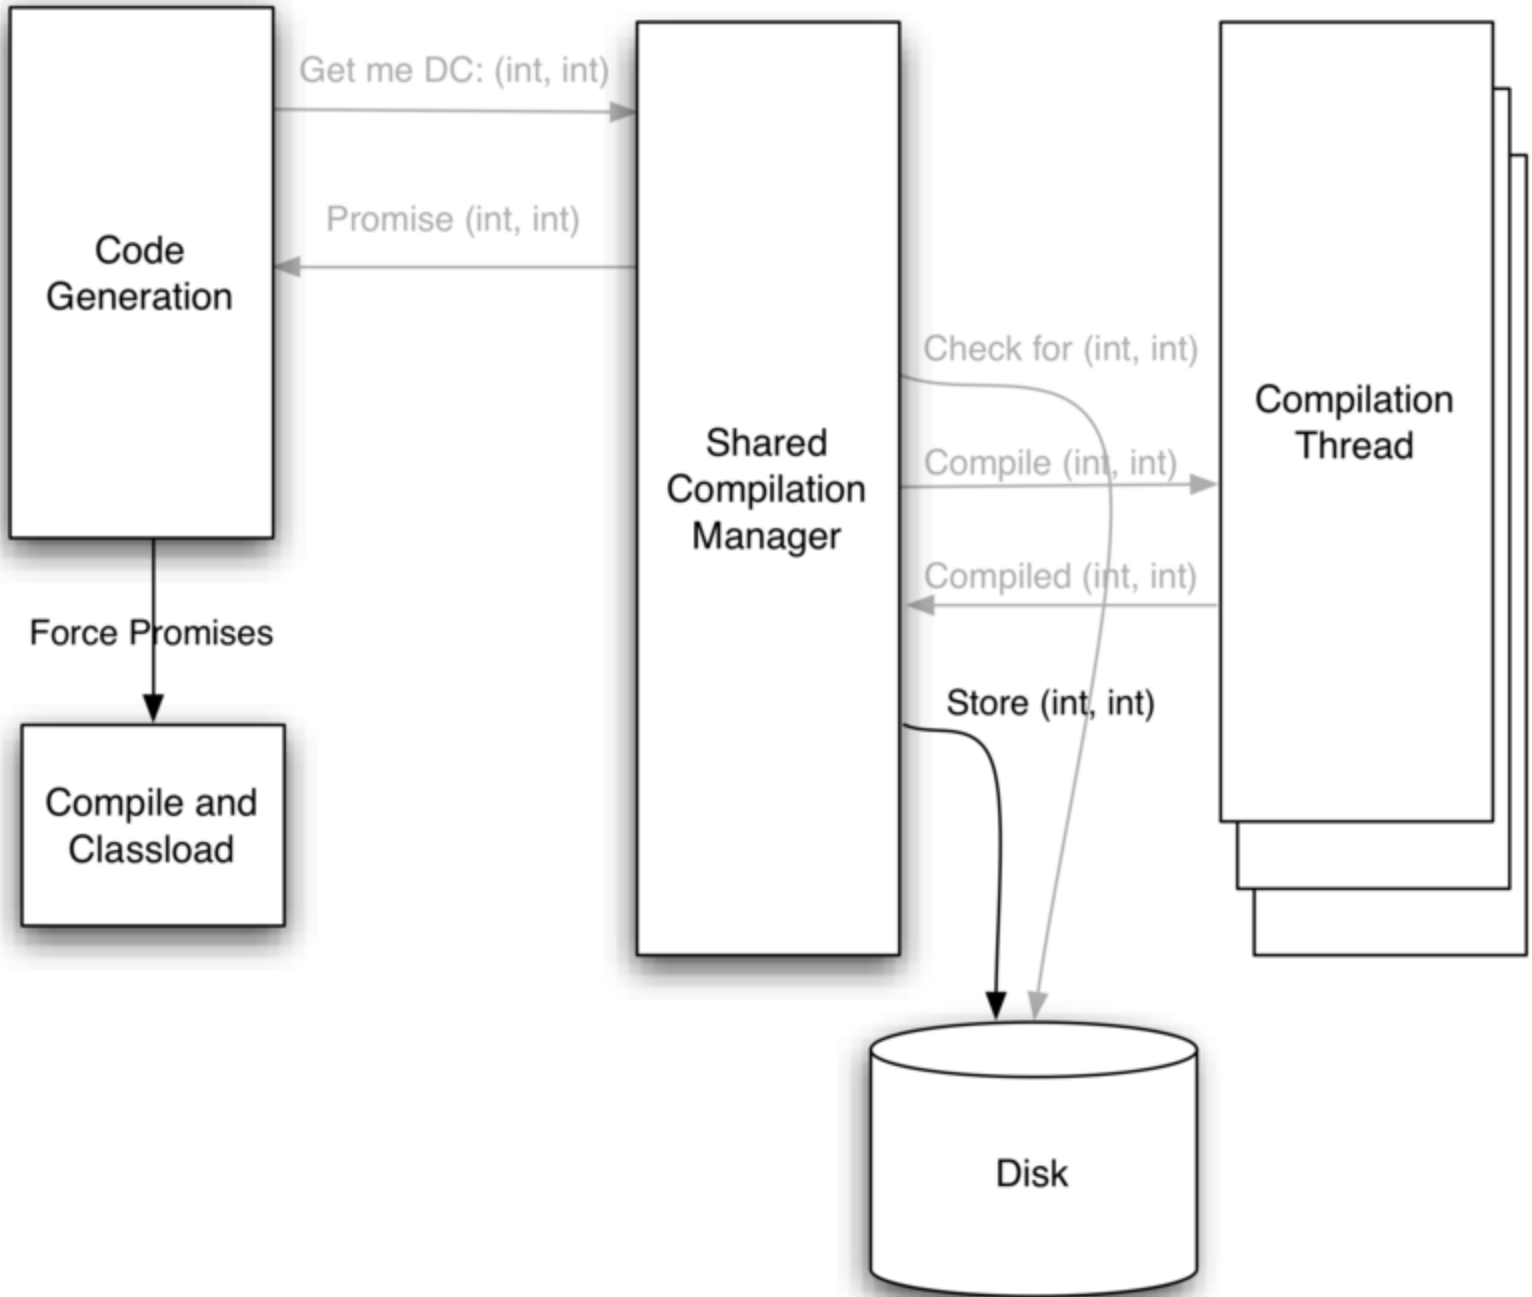


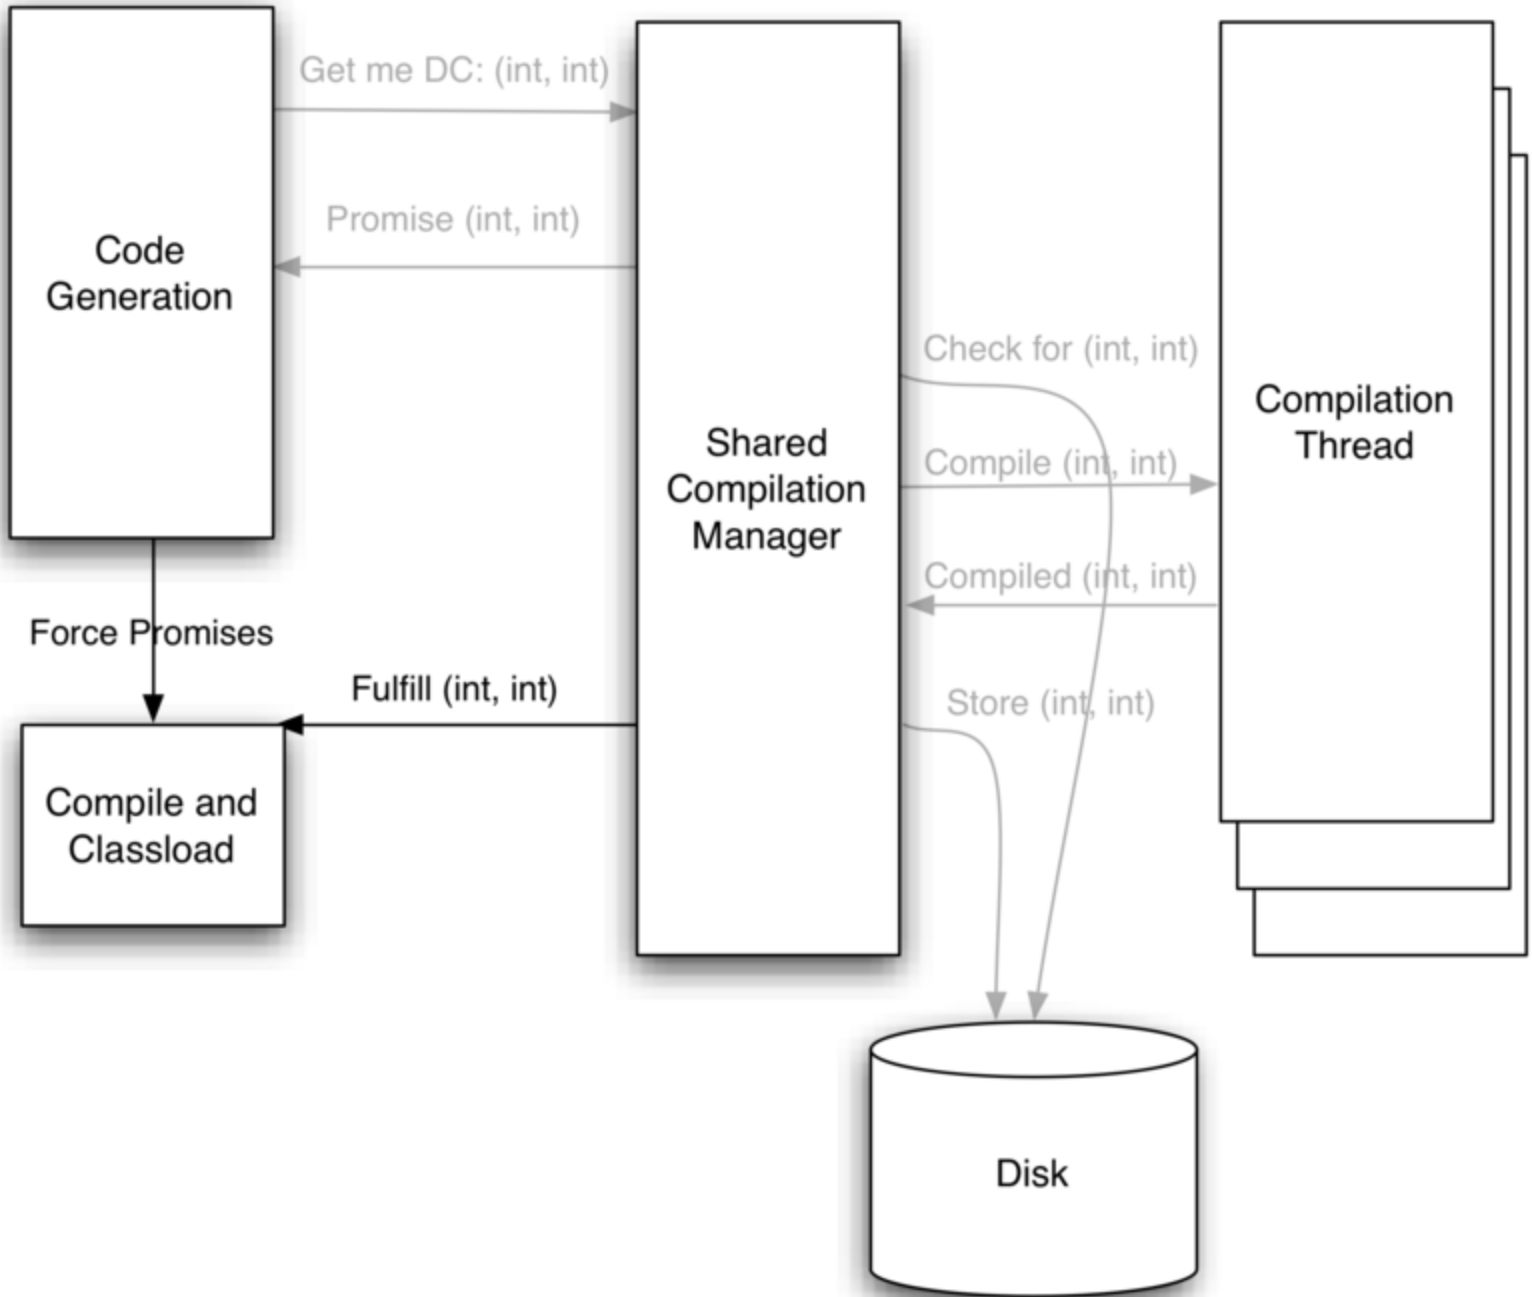













Lessons Learned

- Spaghetti code is bad; spaghetti code generation is worse
- Choose between generated code vs runtime library
- Parallel APIs for generating code and for runtime code
- Type specialization makes code fast
- Separate compilation requires planning
- **Classloaders hurt the unwary**
- Improve tooling

Classloaders

- Sharing
- Lifecycle
- MS Windows



Lessons Learned

- Spaghetti code is bad; spaghetti code generation is worse
- Choose between generated code vs runtime library
- Parallel APIs for generating code and for runtime code
- Type specialization makes code fast
- Separate compilation requires planning
- Classloaders hurt the unwary
- **Improve tooling**

Shameless Plugs

- **StreamBase**
 - You could build one of these yourself, or use ours
 - Download, Buy, OEM
 - We're hiring
- **DEBS – Distributed Event Based Systems**
 - Academic (ACM) Conference: Arlington TX in July 2013
- **EPTS – Event Processing Technology Society**
 - <http://ep-ts.org>, industry consortium
- **IDDS – International Development Design Summit**
 - <http://iddsummit.org/>

The background of the slide is a dark blue grid with various financial data visualizations. In the upper left, there's a horizontal bar with the text '99.68 -1.98' and 'Average' below it. To the right, there's a candlestick chart with several bars. Below that, a line graph with multiple peaks and troughs is visible. In the bottom right, there's a bar chart with many vertical bars of varying heights. A black pen nib is positioned horizontally across the middle of the grid, pointing towards the right. The word 'Questions?' is written in a large, white, sans-serif font in the center of the image.

Questions?

Download StreamBase and More Information

<http://www.streambase.com>