



Five Stages to Streaming Platform Adoption

Lyndon Hedderly
September 2018

Table of Contents

Recommendations for business executives	3
Prior to streaming	4
Stage 1. Streaming awareness and pilot	5
Stage 2. First steps into production streaming	6
Stage 3. Mission-critical integrated streaming	7
Stage 4. Global streaming	8
Stage 5. Central nervous system	9
Summarizing the stages	10
In conclusion	11

Recommendations for business executives

At Confluent, we've worked with a number of organizations that have architected themselves around streaming platforms, to the extent the platform has become fundamental to how the business operates.

These organizations range from digital natives to more traditional companies across Financial Services, Retail, Automotive, Telecom, Healthcare and Government. In fact, we've seen adoption in any industry in which data is critical, which of course is every industry. Indeed, this year's [Apache Kafka® report](#) found that 94% of organizations plan to deploy new applications or systems using Kafka as their streaming platform. And two-thirds [67%] plan to deploy between 1 and 10 new applications or systems.

Use cases for streaming platforms vary from improving the customer experience, to facilitating new business models, to driving increased efficiency and/or mitigating risk. Regardless of the use case, through working with many of these organizations, we have synthesized some common themes of streaming maturity and have identified five stages of adoption, as shown in Figure 1.

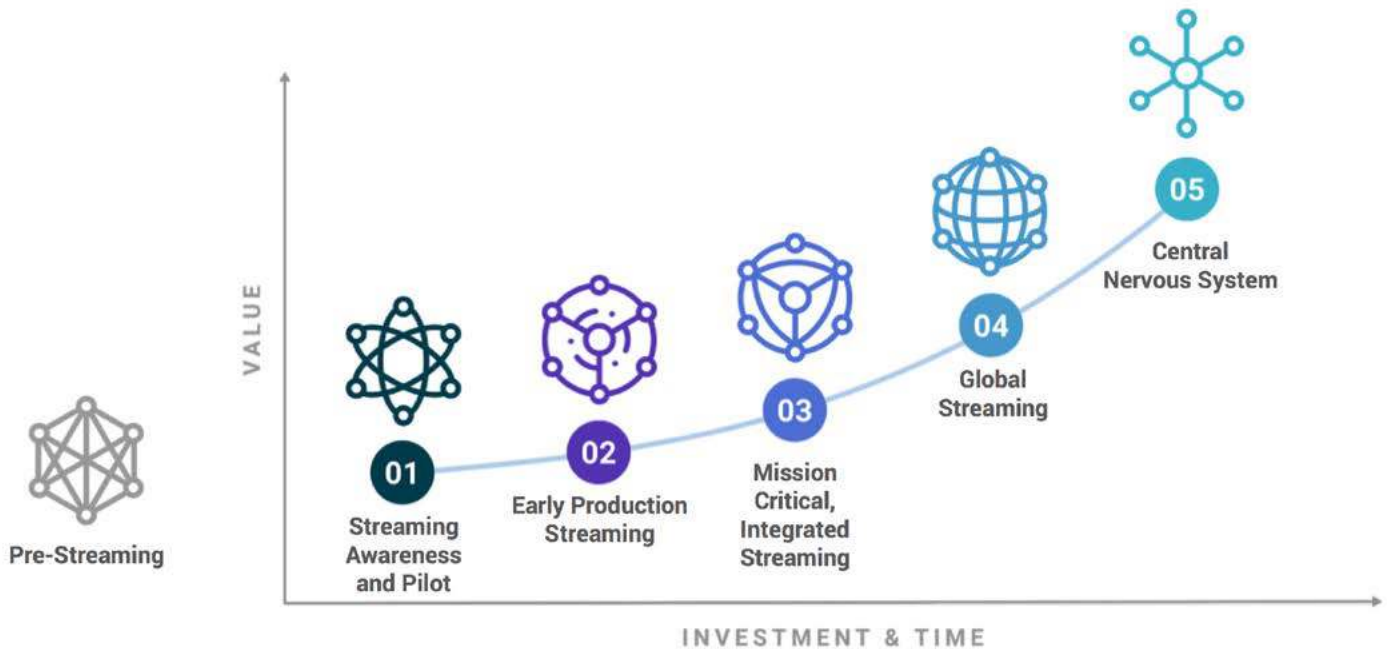


Figure 1. Migrating to a streaming platform is a multi-phase adoption journey. We see five stages to streaming maturity.

The journey starts with awareness and moves up a curve, ending with a streaming platform acting as the central nervous system of the enterprise. This whitepaper communicates these common stages and provides guidance for organizations in different streaming maturity stages. First, we'll start with pre-streaming.

Prior to streaming

Before an organization implements any form of streaming, they sit in the pre-streaming stage. Here, we often see a group, or groups, of data engineers, or developers, struggling with moving data around an organization to deliver required business functionality.

This struggle is usually a result of 'evolved complexity'. That is, legacy core systems that have grown around entrenched, key business processes, three-tier architecture and simple point-to-point connections, with a growing number of integrations. In most organizations, legacy architecture has become a tightly-coupled, complex and highly rigid ecosystem, as depicted in the diagram below.

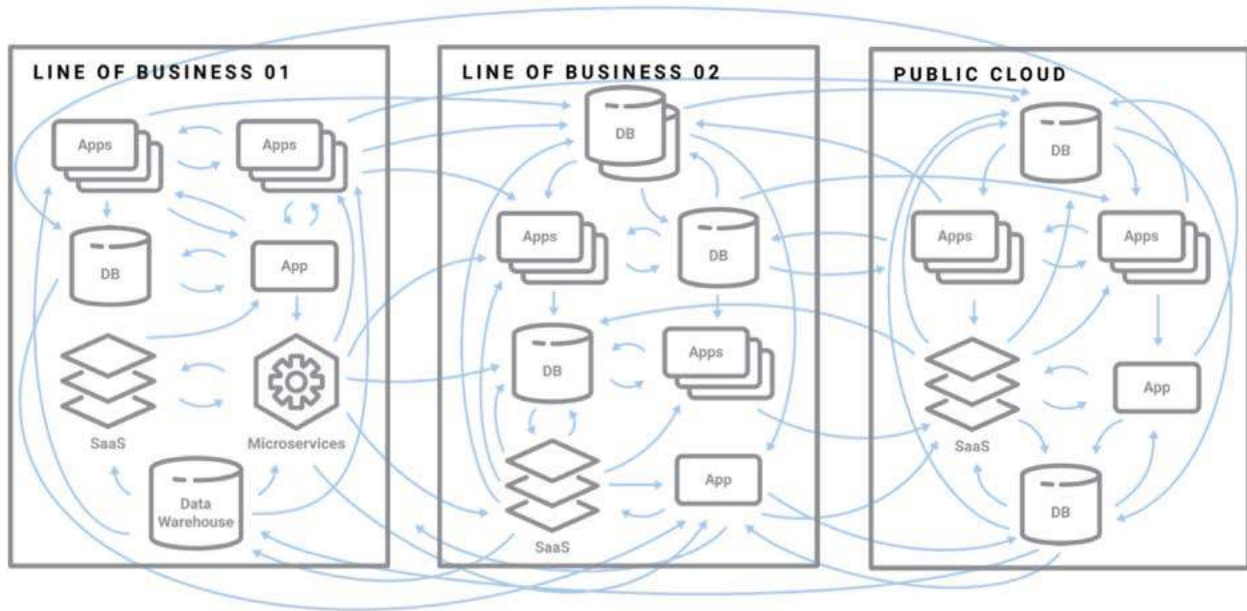


Figure 2. Pre-streaming: Data moves around in bespoke batches and hence processing is often scheduled, rather than real time.

Data becomes 'stuck' within the various legacy systems and/or in business silos.

Time and focus that should be spent building responsive, event-driven applications is instead mostly spent on landing data in the right places, communicating what that data and fixing the misinterpretations and idiosyncrasies that inevitably creep into this kind of piecemeal approach to data movement.

This complexity is exacerbated in organizations with a history of mergers and acquisitions, where different systems have been patched together.

We want to help organizations move away from this legacy complexity, not only to realize value from real-time events, but to also realize the many benefits of simplification, which includes developer velocity and lower costs.

Stage 0 Recommendations:

- Increase awareness of streaming. [Watch this short video](#) by Confluent CEO, Jay Kreps.
- Undertake a Confluent Streaming Discovery.
 - We will work with you to assess, pilot, or develop proof of concept (POC) use cases, business processes, or data challenges that could benefit from streaming.
 - We typically see two options here:
 - Connecting to existing data stuck in external data stores or legacy systems
 - Producing new streams of data (new business models)
- Assess value potential—saving money, making money or protecting money—in moving to a streaming platform. Value could include:
 - Cost savings from architectural simplification.
 - New revenue streams or enhanced revenue through better customer experience.
 - Managing risk, for example implementing fraud detection.

Stage 1. Streaming awareness and pilot

Whereas enterprise technology transformations are traditionally driven from the 'top-down,' we often see streaming platforms entering the enterprise from the bottom-up. The journey mostly starts when the data engineers or developers (mentioned in the previous section) hit a problem, or pain-point with their current legacy infrastructure and architecture.

The tech team may have already worked with products like Kafka, or they come across streaming at a community event or through their own research. This awareness initiates a pilot, or POC. And because Kafka is open source, the barrier to entry is minimal. Teams typically start by tackling one or two edge cases and are modest in scale and ambition. As we explained in stage 0 (pre-streaming), there are two key drivers for a pilot or POC:

1. Connecting to data stuck in external data stores or legacy systems.

- a. Examples include a bank offloading data from the mainframe, into a streaming platform, in order to make these locked-up datasets available to new initiatives, process the 'events' more rapidly and meet regulatory compliance requirements.

2. Producing new streams of data (new business models).

- a. The tech team may play with streaming, producing and processing far larger and faster streams of data than were previously possible. An example in the automotive industry includes a connected car initiative; moving data from car sensors to a central hub and processing the data in real-time for predictive maintenance.

It is also important to note that the evolution of the pre-streaming world into an event-driven one doesn't always mean completely replacing existing or legacy infrastructure. The old paradigms can coexist with the new event-driven one. This stage is typically about supplementing the architecture, where the change to event-centric thinking makes sense. Traditional transactions might still need to be processed using synchronous communication, running in parallel.

Stage 1 Recommendations:

- Move to early production use cases: Start small, think big and act fast.
- The tech team will often be most interested in the increased agility and flexibility of a streaming platform. Their developer velocity increases significantly.
- The business team will obviously see benefits here, but will also see a wider business use case, such as improved customer experience and reduced costs. They get a technology infrastructure that actually describes what their customers do, mapping these onto their various existing business processes. So, a customer or business person can reason about how it all fits together.
- Overall, this transition isn't just about adopting a different technology, it's about changing how you think about data, or events, in your business. It's about making a fundamental shift to event-centric thinking where you are continuously capturing, assessing and responding to streams of events that matter to the business.
 - There may be minimal business value in the first POC step—and this will mostly be in the data engineering area—but the business value will come as you move up the maturity curve.
 - Look for early opportunities to expand.

Stage 2. First steps into production streaming

We saw the transition from stage 0 to 1 is mostly driven from the bottom up the tech team (developers, data engineers, architects) champion the need for streaming platforms because they are seeking efficiency, flexibility and developer velocity.

Organizations typically run small-scale pilots and progress these into production, for non-mission critical (or experimental) use cases. It is stage 1 to 2 where tech and business minds meet. It is here the business application owners see clear business benefit in real-time data, offering improved customer experience (CX), new business models and/or ways to mitigate risk.

Of the two broad categories of 'business-application' adoption we see slightly different value drivers:

1. Connecting to data stuck in external data stores or legacy systems.

- This year's [Apache Kafka report](#) found the top benefit of Kafka is increased agility. Developers and data engineers appreciate being 'loosely coupled,' with increased speed, flexibility and extensibility. As such, they often champion the change, as it simply enables them to get stuff done. Developers and data engineers often realize this benefit before the business 'gets it.' This is mostly around savings: doing more for less.

2. Producing new streams of data (new business models).

- The tech team may play with streaming, producing new streams of data that were previously not possible to process with legacy systems. As with the example of the connected car initiative, this is often driving top-line business value via new business models. Another is capturing, and analyzing, fine grained customer interactions, for example understanding how a customer navigates a website or mobile application to measure engagement or optimize sales conversion ratios.

Using a streaming platform, pilot teams—both tech and business—are able to show how events can be handled in real time, with the corresponding technical and business benefits, proving the value of the platform.

Whilst the tech and business teams work closely in stage 2, the project teams are typically self-contained, or 'bounded.' Only a small number of highly-skilled resources really appreciate the difference between the streaming platform and, say, legacy messaging technology. The wider business may remain unaware of the power of the platform and how streaming can be applied across multiple projects. Projects remain siloed.

Stage 2 Recommendations:

- Tech: Educate the business around the early production use cases, freeing data from legacy systems or silos.
- Business: Look for business opportunities of real-time data flows.
- Consider a Confluent Training and Professional Services offerings.
- For pilots in Stage 1: Demonstrate the benefits of the pilot, either:
 - Freeing data that was stuck in legacy systems or external data stores, or
 - Producing new streams of data that were previously not possible to process with legacy systems.
- Look for opportunities to expand early production applications to mission-critical apps.
- Think about the team in terms of skills and experience. Think: Operating model (see next stage).

Stage 3. Mission-critical integrated streaming

As organizations onboard more apps onto the streaming platform, they move into the third stage, which has two elements: **mission-critical** and **integrated**.

Mission critical: Rather than managing edge cases, this stage is characterized by the streaming platform being tied more closely to the business's overall vision and strategic objectives. Capabilities of the streaming platform that matter in this stage include security, durability, exactly-once guarantees, together with the ability to monitor the event flows across multiple applications and maintain data completeness service-level agreements (SLAs).

Integrated streaming: This means, different business processes or business units working together where before they might have worked in silos. We see a number of organizations breaking silos with Kafka, by joining events that originate in separate parts of the business. a streaming platform can be used to join streams of business facts like orders, trades or payments to tables of reference data like customers, accounts or foreign exchange rates, leveraging the stream table duality to provide a unified view of data in real time.

In this stage, the organization has to think about their operating model. This includes how the business and data functions are structured, organized and managed (the governance). For example, a new mission-critical tech stack most likely requires new skills within the organization and potentially a new delivery model.

Often, at this point, the C-suite and executives get to hear about a streaming platform for the first time. Whilst they see the benefits, some education around the technology is required. Even tech-savvy executives can fail to appreciate the differences in existing or legacy messaging technology, and why streaming is fundamental to business success.

Stage 3 Recommendations:

- As with any major business initiative, streaming should have its own strategic direction. To create a streaming strategy beyond a few use cases, look for opportunities to join-up teams, for efficiencies and economies of scale.
- Leaders should look to create enriched datasets that abstract downstream systems from the complexities of upstream business architectures.
- Focus on the mission-critical use cases, and the ‘ilities:’ reliability, scalability, operability, extensibility etc. as well as security.
- Integration: Think about the business applications and how they are integrated across the organization.
- Think about the Operating model, including delivery model and team (skills etc.). The C-suite, together with whoever is tasked with leading the company’s streaming initiatives, should set up a series of workshops for the executive team to coach its members in the key tenets of advanced streaming and educate the broader management audience. Three crucial questions for the C-suite and company’s leaders:
 - How can streaming help the company make more money?
 - How can it help save operational costs?
 - How can it mitigate risk?

Stage 4. Global streaming

As organizations move beyond the Integrated Streaming stage, they enter the Global Streaming stage. In this stage, the streaming platform has grown, within the business, to the point where it must service customers internationally, or where systems are migrating from on premise infrastructure to cloud.

There’s power in global streaming, but there are also several big challenges a business needs to address in this stage, such as:

- How do you make data available across different data centers or regions?
- How do you serve data efficiently from closer geos?
- How do you implement data sovereignty rules, like GDPR?

The old (pre-streaming) ways of solving these problems are obviously operationally challenging. This includes running hot standbys and/or doing manual failovers for disaster recovery that take several hours or breaking up a global service into bespoke regional services that requires all the complexity of geo partitioning into the application logic.

In global streaming, applications become both portable and federated. For example, a mainframe running in an on-premises data center uses global streaming to make account information available directly in different geographical regions where it is close to users and sovereignty rule compliant, while also performing real-time fraud prevention using machine learning services deployed on the cloud. In this way, businesses can leverage global streaming both as a mechanism for consolidating business processes, aligning with customer’s needs, as well and evolving their architectures to make use of the benefits provided by cloud: putting data close to customer or close to the applications that need it.

Kafka, as a global streaming platform, is evolving and at Confluent, we are working with global organizations to ensure Kafka plays a role to really enable Global Streaming. Please [contact us](#) if you want a Kafka global roadmap presentation.

Stage 4 Recommendation:

- Look to qualify and quantify the business value of global streaming and how will it impact the global business?
- Design the Global Operating Model (Target Operating Model, or TOM).
 - Understand and align with how the business is organized.
 - Streaming capabilities need to be embedded in the business, resulting in an effective streaming organization structure.
- We have observed that organizations with successful streaming initiatives embed streaming capabilities into their core businesses.
- Define roles and responsibilities around streaming.

Stage 5. Central nervous system

The final stage in our maturity model is stage 5, in which the streaming platform effectively becomes the central nervous system of the entire enterprise.

This is the stage that is associated with the digital natives (businesses that were born digital). These organizations have often architected themselves around a streaming platform from the beginning, without the burden of any legacy complexity.

Everything in a digitally-native business is an event and all the data in an organization is managed through a streaming platform. These organizations tend to progress along the streaming journey quickly. They tend to handle large amounts of data in real time, operate applications and services 24/7 with high nines uptime and look to continuously foster future innovation.

Netflix is a prime example of a digital native using Kafka at massive scale. Netflix runs approximately 50+ Kafka clusters, with 4,000+ brokers, processing an astonishing 2+ trillion messages every single day. This is a pretty powerful state for a business to be in and the possibilities are immense.

In this stage, everything happening in the business is available instantly to all applications in the company through the streaming platform. The technical team is happy as the architecture is greatly simplified and they can work efficiently. The business team is happy because they can get real-time insights and actions on data, or events, as they happen vs. when it is too late.

Stage 5 Recommendation:

- Look for opportunities of continuous improvement and benefits realization.

Summarizing the stages

Along the stages of adoption, we have identified different areas of focus. In the earlier stages (0,1,2), the technical team tends to drive streaming and the focus is more bottom-up on infrastructure and architecture that enables data flow. In the later stages (3,4,5) the focus is more top-down, including the strategic objective, the operating model and the business applications.

We have summarized the areas of focus in our Streaming Adoption Framework:

Top-down & bottom-up perspective (incorporating; people, process & technology)

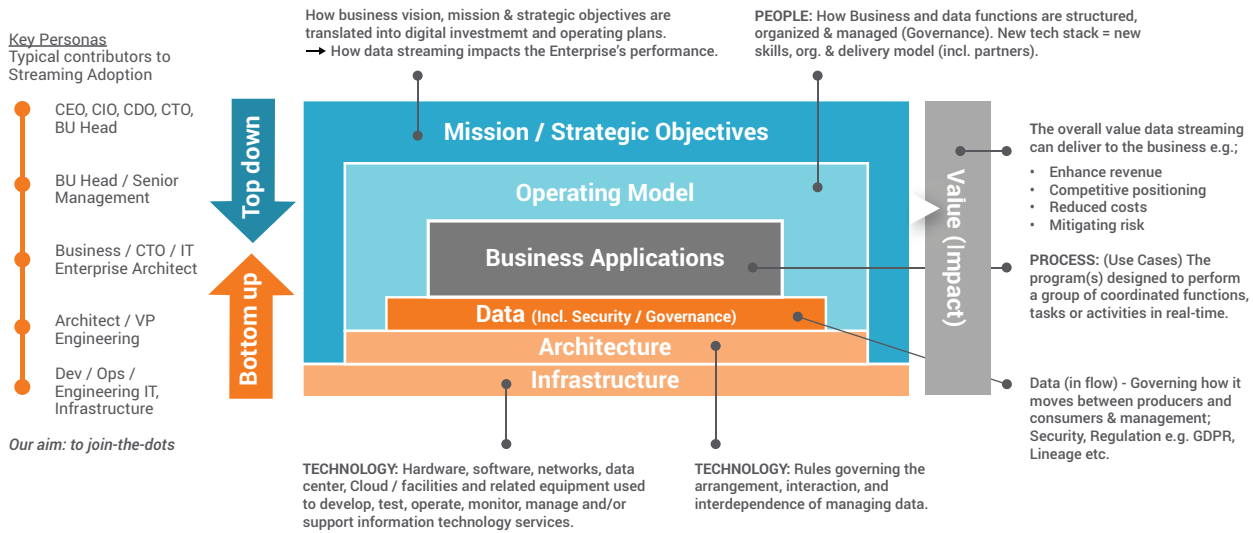


Figure 3. The Confluent Streaming Adoption Framework

A summary of the stages is shown here in Figure 4 including typical characteristics and recommendations.






 1. Streaming awareness & pilot	 2. Early production streaming	 3. Mission critical, integrated streaming	 4. Global streaming	 5. Central nervous system
Characteristics Developer downloads Kafka and experiments, pilot(s).	LOB(s); Small teams experimenting; pub/sub/integration but fragmented.	Multiple mission-critical use cases in production with; scale, disaster recovery and SLAs.	Streaming platform managing majority of mission critical data processes, globally, with multi-datacenter replication.	All data in the organization is managed through a single streaming platform.
Recommendations Educate the business on streaming.	Demonstrate value and review opportunities to expand.	Focus on the mission-critical capabilities: (reliability, scalability etc.) And on overall integration and operating model.	Organize around a Global Operating Model (Target Operating Model, or TOM) for streaming.	Continuous improvement and benefits realization.

Figure 4. Streaming maturity stages, characteristics and recommendations

In conclusion

We hope this level of guidance proves useful, as organizations think about their use of Apache Kafka, as a distributed streaming platform. This document can be used to provide a common language of adoption of streaming with fellow Kafka users, or to help pitch Apache Kafka internally within an organization.

If you have additional information to add, your own story to tell, or any questions or comments on this article, please [get in touch](#). We would love to hear how your organization is moving along the phases of the streaming journey, to arrive at the point where it is truly event-driven.