Cloud native applications, infrastructure and patterns

Part 1: Origins and main characteristics

Renaud Lachaize & Thomas Ropars

Univ. Grenoble Alpes

M2 MoSIG

October 2025

Main references

- B. Scholl, T. Swanson, P. Jausovec. **Cloud native: Using containers, functions, and data to build next-generation applications**. O'Reilly, 2019.

- J. Garrisson, K. Nova. **Cloud-native infrastructure: Patterns for Scalable Infrastructure and Applications in a Dynamic Environment.** O'Reilly, 2017.

- Cloud Native Computing Foundation (CNCF) Web site: https://www.cncf.io

- Google Site Reliability Engineering (SRE) resources:
  - https://sre.google
  - Free eBooks: https://sre.google/books/

Introduction

- **During the first era of Cloud computing, most efforts were focused on facilitating the migration of existing applications** (and their legacy code bases) to Cloud platforms.
  - "Lift and shift" to IaaS (Infrastructure as a Service)
  - Migration of domain-specific applications (e.g., Web applications) to PaaS (Platform as a Service)

- Over the past ~decade, a number of principles have emerged for shaping the design of cloud-based applications and their underlying infrastructure.
- These **new applications have been designed from the ground up in order to take into account the specific characteristics (challenges and opportunities) of modern Cloud platforms and technologies**.

“Cloud native”: a definition

“**Cloud native technologies empower organizations to build and run scalable applications in modern, dynamic environments such as public, private, and hybrid clouds.** Containers, service meshes, microservices, immutable infrastructure, and declarative APIs exemplify this approach.

These techniques enable loosely coupled systems that are resilient, manageable, and observable. Combined with **robust automation**, they allow engineers to **make high-impact changes frequently and predictably** with minimal toil.

The Cloud Native Computing Foundation seeks to drive adoption of this paradigm by fostering and sustaining an ecosystem of open source, vendor-neutral projects. We democratize state-of-the-art patterns to make these innovations accessible for everyone.”

Source: CNCF Cloud Native Definition v1.0, 2018 [emphasis added]

(https://github.com/cncf/toc/blob/master/DEFINITION.md)

“Cloud native”: another definition

“A cloud native application is engineered to run on a platform and is designed for resiliency, agility, operability, and observability.

- **Resiliency** embraces failures instead of trying to prevent them; it takes advantage of the dynamic nature of running on a platform.
- **Agility** allows for fast deployments and quick iterations.
- **Operability** adds control of application life cycles from inside the application instead of relying on external processes and monitors.
- **Observability** provides information to answer questions about application state.”

“Cloud native applications acquire these traits through various methods. It can often depend on where your applications run and the processes and culture of the business.

The following are common ways to implement the desired characteristics of a cloud native application:

Microservices **Health reporting** **Telemetry data** **Resiliency** **Declarative, not reactive**"

(Source: J. Garrisson and K. Nova. Cloud native infrastructure. O’Reilly, 2017.)

Health reporting

- The developers of an application know best what it means for this application to be in a “healthy state”.
- Letting infrastructure providers figure out by themselves the current health of an application often results in fragile designs.
- **A cloud native application should expose its own “health check” interface.**
  - Such an interface can **typically be implemented as a Web endpoint** that returns a health status via an HTTP return code.
  - In addition to error codes, the absence of a prompt reply can also be interpreted as a symptom of a failed task or communication problem.
- **Application may have more than two states (“healthy” / “unhealthy”)**
  - For example, the application may be starting up or shutting down.
  - Giving previse feedback to the infrastructure can help it operate the application in a more robust and efficient way (e.g., distinction between states such as “ready”, “starting (not ready yet to receive traffic)” and “failed”.

Telemetry data

- Not the same thing as health reporting (although some data may overlap) – they serve different purposes.
- **Telemetry data provide information about business objectives**.
  - Sometimes named "**service level indicators**" (SLIs) or "key performance indicators" (KPIs)
  - These data can be **used to check if an application is meeting its "service level objective"** (SLO).
- Telemetry data is often stored in a time series database.
- Examples of questions and metrics ("RED"):
  - **Rate**: "How many requests per second does my application receive?"
  - **Errors**: "Are there any errors?"
  - **Duration**: "How long does it take to obtain a response?"
- Can be used to raise alerts about the global (application-level) behavior.
- Not the same thing as "logs" (logs are mostly used for debugging).

Service Levels: SLAs, SLOs & SLIs (1/3)

- These notions are complementary and all related to the level/quality of service provided to the users/clients using a service (internal or external “customers”).

- The quotes below are taken from Google’s SRE book: https://sre.google/sre-book/service-level-objectives/

- **SLI: Service Level Indicator**
  - **Definition:** “a carefully defined quantitative measure of some aspect of the level of service that is provided”.
  - **Examples:** request latency, system throughput, end-to-end latency, error rate, availability, data durability

Service Levels: SLAs, SLOs & SLIs (2/3)

- **SLO: Service Level Objective**
  - **Definition:** “a target value or range of values for a service level that is measured by an SLI”
  - **A natural structure for SLOs is thus SLI ≤ target, or lower bound ≤ SLI ≤ upper bound**.
  - For example: average search request latency should be less than 100 milliseconds.

  - **Why?** “Without an explicit SLO, users often develop their own beliefs about desired performance, which may be unrelated to the beliefs held by the people designing and operating the service. This dynamic can lead to both over-reliance on the service, when users incorrectly believe that a service will be more available than it actually is, and under-reliance, when prospective users believe a system is flakier and less reliable than it actually is.”

  - **Choosing a set of SLOs can be nontrivial.** For example, several SLIs might be connected behind the scenes:
    - Higher throughput often leads to higher latencies.
    - Many services have some performance cliff/drop beyond some input load threshold

Service Levels: SLAs, SLOs & SLIs (3/3)

- **SLA: Service Level Agreement**
  - **Definition:** “an explicit or implicit contract with your users that includes consequences of meeting (or missing) the SLOs they contain. The consequences are most easily recognized when they are financial—a rebate or a penalty—but they can take other forms.”
  - “An easy way to tell the difference between an SLO and an SLA is to ask "what happens if the SLOs aren’t met?": if there is no explicit consequence, then you are almost certainly looking at an SLO.”
  - “Whether or not a particular service has an SLA, it’s valuable to define SLIs and SLOs and use them to manage the service.”

- For more details, examples and advice:
  - https://cloud.google.com/blog/products/devops-sre/sre-fundamentals-sli-vs-slo-vs-sla
  - https://sre.google/sre-book/service-level-objectives/

Resiliency (1/3)

- **Resilience to failures:**
  - is generally the most important characteristic for an application.
  - is partially managed by the infrastructure but cloud native applications must also handle some part of that responsibility.
  - relies on two main aspects: **design for failure** and **graceful degradation**.

Resiliency (2/3)

- **Design for failure**
  - The SLO specifies the uptime guarantees for a given service.
  - Failure are almost unavoidable (over time) in any complex system.
  - A cloud native application is built with the assumption that failures will happen (although not all precise types/scenarios of failures can be anticipated).
  - Some (severe) kinds of failures cannot be addressed by the application (e.g., network partitions) and should be handled by the cloud platform.

Resiliency (3/3)

- **Graceful degradation**
  - Cloud native apps must be designed to handle excessive load (even though the cloud platform may also help).
  - Graceful service degradation consists in **servicing all/most requests** (i.e., providing available service) **yet with a lower quality of responses** (e.g., with less accuracy or less data – partial answers), in order to lower the request processing costs.

Declarative, not reactive (1/2)

- **Cloud native (distributed) applications should rely on the infrastructure to achieve some kinds properties instead of trying to manage them directly.**
  - For such properties, the application should **declare the desired outcome** and let the infrastructure reach this outcome. The steps to be taken are decided by the infrastructure.
  - This approach allows building simpler and more robust applications (and assemblies of applications/services).

- **Declarative communications**
  - **Applications trust that the network infrastructure will deliver the messages.**
  - The infrastructure can leverage a number of techniques to improve communication resiliency and efficiency, such as load balancing, load shedding, service discovery, retries and timeouts, and circuit breaking.
  - Such infrastructure-level features can be embedded within application using transparent proxies (e.g., implemented as “sidecar” containers).

Declarative, not reactive (2/2)

- **Declarative state**
  - Here, the “state” of an application refers to behavioral or structural properties.
    - For example : the number of replicas for a given service.
  - Old approaches are based on imperative configuration: a description of a sequence of steps to perform in order to reach a given state.
  - In contrast, **with a declarative approach, an administrator describes only the desired state of the distributed application** (not the steps to reach it).

  - Expected benefits:
    - **Fewer errors**. Since a declarative configuration describes an expected result, its impact can be immediately understood.
    - Allows leveraging usual software development tools for manipulating configurations: source versioning, unit testing.
    - **Provides simple support for configuration rollback** in case of problem (i.e., reversible operations).

“The twelve-factor app”

- One of the first historical milestones that lead to the design principles of Cloud-native applications.

- Context: **A (short) document / manifesto written in 2012 by the staff of Heroku, a Cloud provider for applications based on the Platform-as-a-Service (PaaS) paradigm**.
  - “The contributors to this document have been directly involved in the development and deployment of hundreds of apps, and indirectly witnessed the development, operation, and scaling of hundreds of thousands of apps.”

- Available online at: https://12factor.net

- We will summarize it in the next slides, using mostly verbatim quotes.

“The twelve-factor app” – Motivation (1/2)

A synthesis of experience and observations on a wide variety of SaaS apps in the wild.

A triangulation on ideal practices for app development, paying particular attention to :

- The **dynamics of the organic growth of an app** over time,
- the **dynamics of collaboration between developers** working on the app’s codebase,
- and **avoiding the cost of software erosion**.
  - Slow deterioration of software over time that will eventually lead to it becoming slow, faulty or unusable.
  - Typical cause: the software suffers from a lack of updates with respect to the changing environment in which it resides.

Source: Most of the contents of this slide are borrowed/adapted from https://12factor.net

"The twelve-factor app" – Motivation (2/2)

Goals:

- "**to raise awareness of some systemic problems** seen in modern applications development,
- **to provide a shared vocabulary** for discussing those problems,
- and **to offer a set of broad conceptual solutions** to those problems with accompanying terminology."

Source: Most of the contents of this slide are borrowed/adapted from https://12factor.net

The 12 factors – Overview

I) Codebase: One codebase tracked in revision control, many deploys.

II) Dependencies: Explicitly declare and isolate dependencies.

III) Config: Store config in the environment.

IV) Backing services: Treat backing services as attached resources.

V) Build, release, run: Strictly separate build and run stages.

VI) Processes: Execute the app as one or more stateless processes.

VII) Port binding: Export services via port binding.

VIII) Concurrency: Scale out via the process model.

IX) Disposability: Maximize robustness with fast startup and graceful shutdown.

X) Dev/prod parity: Keep development, staging, and production as similar as possible.

XI) Logs: Treat logs as event streams.

XII) Admin processes: Run admin/management tasks as one-off processes.

Source: Most of the contents of this slide are borrowed/adapted from https://12factor.net

“Beyond the 12-factor application”

A revised list of guidelines proposed by Kevin Hoffman (Pivotal) in 2016.

(Freely available booklet from Pivotal/O’Reilly: https://content.pivotal.io/ebooks/beyond-the-12-factor-app)

- Strongly inspired by the original 12-factor manifesto from Heroku.
  - Revisits some factors
  - Adds some new factors
  - Changes the order of some factor to highlight a sense of priority
- **Overall: 15 factors**

- **Warning: “**Rather than adopting an all-or-nothing approach, learning where and when to compromise on the guidelines […] is probably the single most important skill to have when planning and implementing cloud-native applications.”

“Beyond the 12-factor application” - List

1) One codebase, one application
2) API first
3) Dependency management
4) Design, build, release and run
5) Configuration, credentials, and code
6) Logs
7) Disposability
8) Backing services
9) Environment parity
10) Administrative processes
11) Port binding
12) Stateless processes
13) Concurrency
14) Telemetry
15) Authentication and authorization

Infrastructure as Code (IaC)

- **A set of tools and methodologies aimed at achieving automated and reproducible deployments of software configurations on (physical or virtual) machines.**
  - Promotes the usage of machine-readable configuration description files, rather than interactive configuration tools.
  - The configuration descriptions can be managed using a version control system.
  - Abstracts away and leverages low-level APIs (e.g., the machine provisioning API of a given cloud vendor, or the configuration setup of a given operating system).
  - Can address complex deployments requirements (e.g., multi-step coordination of dependencies between machines/services).
- **Typically based on a declarative approach:**
  - Goes beyond mere scripting of configuration steps.
  - Administrator defines the desired target configuration, and the tools perform the necessary actions to reach this target state, regardless of the initial state (and the potential intermediate failures).
- **Examples of tools**: Terraform (declarative), Ansible (imperative).

Infrastructure as Software

- An infrastructure management approach that goes further than “infrastructure as code”.

- **“Infrastructure as code” relies on a static description of the target infrastructure.**
  - There is limited support for managing the evolution of the infrastructure.
  - Risks of frequent configuration drifts.

- In contrast, **“infrastructure as software” is a continuously running service** that:
  - Builds and maintains a representation of the current state of the infrastructure
  - Monitors the infrastructure and the desired state specified by the administrator (declarative approach)
  - Mutates the infrastructure (and its representation), in order to reach (or maintain) the desired state.
  - Is based on a reconciler pattern (control loop) to converge towards the desired state.

- **Container orchestrator systems (like Kubernetes) are an incarnation of this approach.**

Immutable infrastructure

- A popular approach for the management of configuration modifications within a cloud infrastructure.
- **Main idea:**
  - One should **avoid mutating the existing/deployed infrastructure.**
  - Instead, **it is better to allocate/deploy (from scratch) a new version of the infrastructure (and then decommission the old one).**
  - Metaphor: “treat infrastructure resources (e.g., VMs) like cattle, not like pets!” (a.k.a. “phoenixes vs. snowflakes”)
- Based on some key observations:
  - Lower risks of configuration drifts/problems if we always start from a known, consistent state.
  - This approach is more heavyweight but nonetheless viable, because the creation/destruction of virtualized resources (e.g., VMs or containers) can be made efficient and fully automated.
- Expected benefits: improved configuration consistency, predictability and reliability.
- Limitations: **mainly suited to stateless components.**

Stateless vs. stateful components (1/2)

- One of the major traits of cloud native applications is related to how they handle application state.

- **Some definitions:**
  - By “state”, here, we mean the information that must be retained by a service/component after it has finished processing a request/job.
  - A **stateless component** does not retain any state after the completion.
  - A **stateful component** retains state, either temporarily (**“session state”**) or permanently (**“persistent state”**)

Stateless vs. stateful components (2/2)

- In order to handle scalability, autoscaling, and fault-tolerance in a flexible and efficient way, cloud native applications are based on design principles that strive to:
  - **Dissociate as much as possible the stateful parts and the stateless parts in an application** (“externalized state”)
  - If appropriate, **use distinct stateful components to store different types of state information** (see the “microservices” paradigm discussed later)
    - Different types of storage components and/or different instances of the same component
- In other words, compute and storage layers are decoupled, at a fine granularity.

CNCF: Cloud Native Computing Foundation (1/2)

- A non-profit foundation (part of the Linux Foundation), with many industrial members
- **The foundation's mission is to make cloud native computing ubiquitous through open-source projects.**
- Main responsibilities:
  - Stewardship of the projects
  - Fostering the growth and evolution of the ecosystem
  - Promotion of the underlying technologies, and approach to application definition and management, including: events and conferences, marketing (SEM, direct marketing), training courses and developer certification
  - Serve the community by making the technology accessible and reliable.
- **Useful links:**
  - Web site: https://www.cncf.io
  - **List of projects ("interactive landscape"): https://landscape.cncf.io**

CNCF: Cloud Native Computing Foundation (2/2)

“The CNCF will strive to adhere to the following principles:

- **Fast is better than slow.** The foundation enables projects to progress at high velocity to support aggressive adoption by users.
- **Open**. The foundation is open and accessible, and operates independently of specific partisan interests. […] the foundation’s technology must be available to all according to open-source values.
- **Fair.** The foundation will avoid undue influence, bad behavior or “pay-to-play” decision-making.
- **Strong technical identity.** The foundation will achieve and maintain a high degree of its own technical identify that is shared across the projects.
- **Clear boundaries.** The foundation shall establish clear goals, and in some cases, what the non-goals of the foundation are to allow projects to effectively co-exist, and to help the ecosystem understand where to focus for new innovation.
- **Scalable.** Ability to support all scales of deployment, from small developer centric environments to the scale of enterprises and service providers. This implies that in some deployments some optional components may not be deployed, but the overall design and architecture should still be applicable.
- **Platform agnostic.** The specifications developed will not be platform specific such that they can be implemented on a variety of architectures and operating systems.”

CNCF Cloud native trail map

- Available from: https://github.com/cncf/landscape/blob/master/README.md#trail-map
- “Provides an overview for enterprises starting their cloud native journey.”
- “[Describes] a recommended process for leveraging open-source, cloud-native technologies.”
- 10 steps:

1. Containerization	2. CI/CD
3. Orchestration & application definition	4. Observability & analysis
5. Service proxy, discovery & mesh	6. Networking & policy
7. Distributed database & storage	8. Streaming & messaging
9. Container registry & runtime	10. Software distribution