

# ByteByteGo — The Complete Channel Summary

Every video summarized, in order, for study

168 videos · 30 Shorts · 867 min total runtime

2022-06 → 2026-08 · Generated 2026-08-15

*For personal study only. All content (c) ByteByteGo.*

# Table of Contents

---

## **I Fundamentals: Networking, HTTP, and How the Web Works** (35 videos)

- 1** What happens when you type a URL into your browser?
- 9** HTTP/1 to HTTP/2 to HTTP/3
- 21** What Is GraphQL? REST vs. GraphQL
- 22** What Is Single Sign-on (SSO)? How It Works
- 24** What is RPC? gRPC Introduction.
- 25** SSL, TLS, HTTPS Explained
- 27** What is OSI Model | Real World Examples
- 38** The Most Beloved Burger for Developers
- 47** Top 6 Most Popular API Architecture Styles
- 51** OAuth 2 Explained In Simple Terms
- 55** Top 7 Ways to 10x Your API Performance
- 56** Why Google and Meta Put Billion Lines of Code In 1 Repository?
- 65** HTTP Status Codes Explained In 5 Minutes
- 66** Python Vs C++ Vs Java!
- 68** How Big Tech Ships Code to Production
- 69** Top 8 Most Popular Network Protocols Explained
- 71** Why is JWT popular?
- 72** How Does Linux Boot Process Work?
- 80** Top 9 Most Popular Types of API Testing
- 87** Linux File System Explained!
- 91** Good APIs Vs Bad APIs: 7 Tips for API Design
- 92** Top 6 Tools to Turn Code into Beautiful Diagrams
- 94** Top 3 Things You Should Know About Webhooks!
- 104** Top 12 Tips For API Security
- 108** Top 9 Most Popular API Protocols
- 109** Do You Know How Mobile Apps Are Released?
- 115** Linux Crash Course - Understanding File Permissions
- 116** Session Vs JWT: The Differences You May Not Know!
- 119** How the Internet Works in 9 Minutes
- 120** HTTP 1 Vs HTTP 2 Vs HTTP 3!
- 127** Everything You NEED to KNOW About Web Applications
- 128** How SSH Really Works
- 130** API Pagination: Making Billions of Products Scrolling Possible
- 133** API Vs SDK! What's the difference?
- 168** HTTP vs HTTPS Explained

## **II Databases, Storage, and Data Systems** (25 videos)

- 2** System Design: Why is Kafka fast?
- 3** System Design: How to store passwords in the database?
- 11** The Secret Sauce Behind NoSQL: LSM Tree
- 13** Back-Of-The-Envelope Estimation / Capacity Planning
- 14** How To Choose The Right Database?
- 28** CAP Theorem Simplified
- 32** 8 Key Data Structures That Power Modern Databases
- 34** Top 5 Redis Use Cases
- 40** Cache Systems Every Developer Should Know
- 45** Secret To Optimizing SQL Queries - Understand The SQL Execution Order
- 49** How Discord Stores TRILLIONS of Messages
- 95** Caching Pitfalls Every Developer Should Know
- 100** ACID Properties in Databases With Examples
- 102** Roadmap for Learning SQL
- 110** KISS, SOLID, CAP, BASE: Important Terms You Might Not Know!
- 111** What is Data Pipeline? | Why Is It So Popular?
- 112** Kafka vs. RabbitMQ vs. Messaging Middleware vs. Pulsar
- 113** 7 Must-know Strategies to Scale Your Database
- 124** Top Kafka Use Cases You Should Know
- 131** Apache Kafka Fundamentals You Should Know
- 149** Trillions of Web Pages: Where Does Google Store Them?
- 154** How Key value Stores Work (Redis, DynamoDB, Memcached)?
- 159** System Design: Why is Kafka Popular?
- 161** What Is Redis Really About? Why Is It So Popular?
- 162** What is a Data Lakehouse?

# 1. What happens when you type a URL into your browser?

Runtime: 5:20 · Published: Jun 2022 · Type: Video

## What it covers


- URL anatomy: scheme (`http://` vs encrypted `https://`), domain, path, and resource — the last two act like directory/file on a server.
- DNS lookup resolves `example.com` to an IP; caching happens at every layer (browser cache → OS cache → DNS resolver → upstream chain), each holding answers briefly.
- A TCP connection is established via a multi-round-trip handshake; modern browsers reuse connections via keep-alive to avoid paying it per request.
- HTTPS adds an expensive SSL/TLS handshake; browsers mitigate with SSL session resumption.
- Over the established connection the browser sends a simple HTTP request; the server responds, the browser renders the HTML, then repeats DNS/TCP/HTTP for extra resources (JS bundles, images).

## Key takeaways


- DNS answers are cached at each hop — that caching is why repeat lookups are near-instant.
- The TCP handshake's round trips are the real cost; keep-alive reuse is the browser's big optimization.
- "Type a URL and hit enter" is a cascade: DNS → TCP → HTTP → render → repeat for subresources.

## Beyond the video

- Modern browsers fetch subresources over HTTP/2 or HTTP/3 (QUIC over UDP), which multiplexes streams over one connection and cuts the per-request round trips the video implies. Per RFC 9113 / the QUIC transport RFC 9000.
- TLS 1.3 cut the handshake to one round trip (0-RTT resumption), making session resumption much cheaper than when this video aired. Per RFC 8446.
- Since 2022, DNS-over-HTTPS (DoH) is common, so "browser asks the OS" often becomes an encrypted query to a DoH resolver. Per RFC 8484.



URL components breakdown



full flow: DNS → TCP → HTTP → render

# 9. HTTP/1 to HTTP/2 to HTTP/3

Runtime: 4:06 · Published: Aug 2022 · Type: Video

## What it covers

- HTTP/1 (1996): built on TCP, but every request to the same server needs a separate TCP connection — expensive because each new connection repeats the TCP 3-way handshake.
- HTTP/1.1 (1997): adds keep-alive so one connection is reused across requests, cutting latency by avoiding the handshake per request; also added HTTP pipelining, which lets a client send multiple requests before waiting for responses — but responses must come back in request order, it was hard to implement correctly, many intermediary proxies mishandled it, and browser support was eventually removed.
- HTTP/1.1's flaw: head-of-line (HOL) blocking — later requests on a connection wait for earlier ones; a single stalled request (e.g., packet loss) blocks everything behind it. Browsers compensate by opening multiple parallel TCP connections to the same server.
- HTTP/2 (2015): introduces streams — many independent request streams multiplexed over one TCP connection. Unlike pipelining, streams don't need in-order send/receive, so HOL blocking disappears at the application layer. HTTP/2 also adds server push: the server sends updates without the client pulling.
- But HTTP/2 still runs on TCP, so transport-layer HOL blocking remains: one lost packet can stall the whole TCP connection and every stream on it.
- HTTP/3 (draft 2020, published June 2022): swaps TCP for QUIC, which is built on UDP. QUIC makes streams a first-class transport-layer concept: streams share one QUIC connection (no extra handshakes to open new ones) and are delivered independently, so packet loss on one stream doesn't affect others — eliminating HOL blocking at the transport layer.
- QUIC is designed for mobile-heavy usage: people switching between networks (Wi-Fi → cellular) get sluggish handoffs with TCP; QUIC's connection ID lets a connection move between IP addresses and network interfaces quickly and reliably.
- Adoption at the time: HTTP/3 already used by 25% of websites and supported by many browsers.

## Key takeaways

- The story is one problem per generation: 1.0's per-request connections → 1.1's keep-alive + pipelining → HOL blocking → 2's independent streams → TCP-level HOL blocking → 3's QUIC/UDP streams.
- Pipelining failed not because the idea was wrong but because ordered responses and broken proxies made it impractical; HTTP/2 streams fix exactly that by removing ordering.
- QUIC's connection ID is the key mobile win: connections survive IP changes, unlike TCP sessions tied to the 4-tuple.

## 2. System Design: Why is Kafka fast?

Runtime: 5:02 · Published: Jun 2022 · Type: Video

### What it covers


- "Fast" is ambiguous: Kafka is optimized for throughput (records moved per second), not latency — "a big pipe moving liquid."
- Sequential I/O: an append-only log as the primary data structure means reads/writes hit the disk arm in order. On modern HDD arrays sequential writes reach hundreds of MB/s vs hundreds of KB/s random — several orders of magnitude apart.
- HDD economics: ~1/3 the price, ~3× capacity of SSD, so Kafka can retain messages long-term cheaply — unusual for messaging systems at the time.
- Zero-copy: without it, sending a disk page to a consumer takes 4 copies (disk→OS cache→Kafka app→socket buffer→NIC buffer) and 2 syscalls. With the `sendfile` syscall, data goes OS cache → NIC buffer directly; with DMA the CPU isn't involved at all.

### Key takeaways


- Sequential access on spinning disks is orders of magnitude faster than random access — the append-only log exploits this.
- `sendfile` + DMA collapses a 4-copy path to one copy, which is what makes Kafka's high-throughput reads possible.
- The HDD cost/ trade-off is deliberate: cheap retention is a Kafka feature, not a compromise.

### Beyond the video

- Kafka has evolved beyond the log: KRaft (KIP-500) replaced ZooKeeper for metadata in Kafka 3.3+, and tiered storage (KIP-405) pushes old segments to object stores. Per [apache.org Kafka docs](https://kafka.apache.org/docs).
- Modern deployments run on SSDs/NVMe where sequential vs random still matters but the gap is smaller; the zero-copy principle remains central, and `sendfile` is still the documented path for log reads.
- The video's "large pipe" framing matches the producer batching model — clients batch records (`linger.ms`, `batch.size`) to keep the log append pattern sequential. Per Kafka producer docs.



zero-copy vs 4-copy data path



producer → OS buffer → disk persistence flow

See also: [8](#)

## 3. System Design: How to store passwords in the database?

Runtime: 3:44 · Published: Jul 2022 · Type: Video

### What it covers

- Never store plain text — anyone with DB access (or an attacker who compromised it) gets everything.
- Hash first: hashing is one-way; a stolen hash can't be replayed to log in. Use modern, deliberately slow hash functions designed for passwords; fast legacy hashes (MD5, SHA-1) are insecure for this.
- Salt second: a unique random string per password, combined before hashing. Without it, precomputation attacks (rainbow tables, DB lookups) crack one-way hashes in seconds; a unique salt makes each hash unique and those attacks unattractive.
- Salt is not secret — store it in plain text next to the hash.
- Login flow: fetch the user's salt → append to entered password → hash → compare with stored hash; equal means valid.

### Key takeaways

- Rules: never plain text, use a slow modern hash (per OWASP), always use a per-password random salt, store salt + hash together.
- Salting defeats rainbow tables even though the salt is public — uniqueness is what matters.
- Validation is just `hash(password + salt) == stored hash`.

### Beyond the video

- OWASP's current Password Storage Cheat Sheet still recommends Argon2id as the first choice, then `scrypt`, `bcrypt`, `PBKDF2` — the "slow functions" the video describes; it also sets per-hash parameters (e.g., `bcrypt cost ≥ 10`). Per [owasp.org](https://owasp.org).
- Since 2022, NIST SP 800-63B-4 guidance shifted toward checking passwords against breach lists and using pepper/encryption as defense-in-depth; Argon2id has been the mainstream recommendation. Per [nist.gov](https://nist.gov).
- The video's "MD5/SHA-1 are fast, don't use them" holds: they're designed for speed, which is the opposite of what password hashing needs.

## 4 4. Big Misconceptions about Bare Metal, Virtual Machines, and Containers

Runtime: 7:02 · Published: Jul 2022 · Type: Video

### What it covers


- Bare metal: single-tenant physical server, complete control, no noisy neighbor, immune to side-channel attacks; downsides are cost, manual management, slow hardware acquisition.
- VMs: hypervisor (VMM) on a host OS virtualizes hardware so many guest OSes run alongside each other; a bare-metal hypervisor skips the host OS for more control and performance. Cheap to run many, easy to scale (even live-migrate), but subject to noisy neighbor and side-channel attacks (Meltdown, Spectre).
- Containers: package app + libraries + runtime; the container engine virtualizes the OS, not the hardware. Faster provisioning, lighter than VMs, far more containers per host, near-instant startup as native processes — but weaker isolation since they share the host OS kernel.
- Containers-in-VMs is a real pattern: it trades flexibility for a reduced attack surface.
- Serverless/edge are the next rung, with their own trade-offs.

### Key takeaways

- The decision is a trade-off spectrum: bare metal = max performance/isolation, containers = max density/portability, VMs in between.
- Noisy neighbor and side-channel attacks are VM/container problems that bare metal's physical isolation eliminates.
- "Bare metal hypervisor" ≠ "bare metal server" — the term refers to the hypervisor running directly on hardware.

### Beyond the video

- Kernel-based isolation for containers has hardened since 2022: user namespaces, seccomp, and gVisor/Kata-style sandboxed runtimes narrow the shared-kernel attack surface the video flags. Per [linuxcontainers.org](https://linuxcontainers.org) / [kata-containers.io](https://kata-containers.io).
- The cloud vendors' shift: on-demand bare metal (AWS EC2 Bare Metal, Azure BareMetal) makes the "slow to acquire" downside less severe than when the video aired. Per [aws.amazon.com EC2 docs](https://aws.amazon.com/ec2/docs).
- Meltdown/Spectre mitigations (e.g., L1TF, MDS) added measurable VM overhead, which reinforces the video's isolation-vs-cost framing rather than overturning it.



three-stack comparison: Bare Metal / Virtualized / Containerized



zoomed hypervisor + guest OS layers

## 17 17. What Are Microservices Really All About? (And When Not To Use It)

Runtime: 4:45 · Published: Oct 2022 · Type: Video

### What it covers

- Microservices are loosely coupled services, each handling a dedicated function (e.g., shopping cart, billing, user profile, push notifications); these functional areas are sometimes called domains. They communicate via well-defined interfaces with small service areas.
- Small service areas limit the blast radius of failures/defects and keep each service easy to reason about. Communication mixes RPC (e.g., gRPC — faster response but larger blast radius when a service goes down) with event streaming/message brokers (better isolation, but slower processing).
- Independent deployment: each service is smaller and safer to change, giving operators confidence to deploy often. Services scale up independently — operational flexibility that is "invaluable."
- Strong information hiding: the monolith's database is broken into logical components (a separate schema in one cluster, or a separate physical database) hidden inside each microservice. The big drawback: the DB can no longer enforce foreign keys/referential integrity across units, so data integrity moves to the application layer.
- Supporting components: an API gateway routes requests to services, relying on an identity provider for authentication/authorization; a service registry + discovery service lets the gateway locate services and lets services find each other; plus monitoring/alerting and DevOps tooling.
- When to use it: microservices cost real money to build and operate, so they make sense only for large teams — independent teams per domain, fast movement, contained blast radius. Not a fit for small startups; the advice is to design each function behind a well-defined interface so a later migration is manageable if the business and team grow.

### Key takeaways

- Microservices are about organizational independence and blast-radius control, not just technology.
- Breaking the monolith DB means losing cross-unit referential integrity — integrity becomes the application's job.
- Startups should design clean interfaces now; migrate to microservices only when team size and business growth justify the overhead.

## 5 5. FAANG System Design Interview: Design A Location Based Service (Yelp, Google Places)

Runtime: 24:40 · Published: Jul 2022 · Type: Video

### What it covers



- Requirements: given location + search radius, return businesses within it; business owners can CRUD businesses (changes may lag a day); users view business details.
- Non-functional + estimates: 100M DAU, 200M businesses, ~5 searches/DAU → ~5,000 QPS. Business table ~1KB-10KB/row → low TB range; geospatial index table 200M × 24B (id + lat + long, 8B each) ≈ 5GB — tiny enough for in-memory options.
- API: GET search (lat, long, optional radius; simplified without pagination) + standard CRUD for businesses. RESTful.
- High-level design: load balancer routes to a stateless Location-Based Service (read-heavy, 5K QPS) and a Business Service (low write QPS); primary-replica DB cluster since reads dominate and consistency lag is acceptable.
- Geospatial indexing: circle-scan with bounding-box SQL needs a table scan — slow at 200M rows; per-axis indexes don't help because 2D intersection is still huge. Fix: map 2D → 1D.
- Hash-based: even grid has uneven business density (NYC vs rural); geohash subdivides the world into quadrants (2 bits each), base-32 encoded, precision 4-6 covers ~0.5-20km radii; prefix sharing ≈ proximity, with boundary edge cases solved by querying the 8 neighbors.
- Tree-based: quadtree (recursively split until ≤100 businesses/cell) and Google S2 — in-memory structures built by app code, not a DB.
- Chosen design: geospatial\_index table (geohash + business\_id compound key, precision-6 geohash column, LIKE prefix queries) ≈ 6GB — fits one server; scale reads with replicas rather than sharding. Business table gets a cache in front; omit cache for the video's design.
- Request trace: 500m radius → geohash length 6 → 9 geohashes (cell + 8 neighbors) → fetch IDs + lat/long → compute distances, rank, return.

### Key takeaways

- Numbers first: 5K QPS search, 5GB index table, 1TB business table — these drive every choice (replicas over sharding, cache placement).
- The core trick: reduce 2D location to 1D (geohash string) so a normal relational index works.
- Geohash edge cases are real (no shared prefix across equator/meridian; shared prefix across grid boundaries) — the 8-neighbor query is the standard fix.
- Don't over-engineer: when the whole dataset fits one machine, read replicas beat sharding for complexity/ops cost.

### Beyond the video

- Redis 3.2+ ships GEO commands and Postgres has PostGIS (both named in the video); Postgres 14+ also added built-in earthdistance-free ST\_DWithin optimizations via GiST indexes. Per [redis.io](https://redis.io) / [postgis.net](https://postgis.net).
- S2 has largely displaced raw quadtrees in production (Google S2 is what powers Google Maps-style indexing); Uber's H3 (hexagonal) is a modern alternative with more uniform cell shapes. Per [s2geometry.io](https://s2geometry.io) / [h3geo.org](https://h3geo.org).
- Since 2022, geohash remains a fine simple choice; the "LIKE '9q8zn%'" prefix pattern is still the standard way to do it in MySQL/Postgres. Per [mysql.com](https://mysql.com) [geospatial docs](https://mysql.com/docs/8.0/geospatial/).



See also: [7](#)

## 8 8. System Design: Why is single-threaded Redis so fast?

Runtime: 3:38 · Published: Aug 2022 · Type: Video

### What it covers

- In-memory: memory access is orders of magnitude faster than random disk I/O → high throughput, low latency; trade-off is dataset ≤ RAM, and simpler code (no on-disk structure management) contributes to stability.
- Single-threaded core: counterintuitive, but multithreading needs locks/synchronization that add complexity and bugs; a single-threaded code path is easy to reason about.
- I/O multiplexing is how one thread serves thousands of connections: the OS lets a thread wait on many sockets at once — select/poll scale poorly, but Linux epoll handles many thousands of connections in constant time.
- Drawback: doesn't use all CPU cores; common fix is running several Redis instances per server.
- Bonus: in-memory design lets Redis use efficient low-level structures (linked lists, skip lists, hash tables) without disk-persistence constraints.

## 7 7. Consistent Hashing | Algorithms You Should Know #1

Runtime: 8:04 · Published: Aug 2022 · Type: Video

### What it covers


- Problem: horizontal scaling spreads data across servers; naive hashing ( $\text{hash}(\text{key}) \% N$ ) distributes evenly while  $N$  is fixed, but adding/removing a server changes  $N$  and remaps most keys — a storm of misses and data movement.
- Core idea: hash server names too, onto the same hash space; connect the ends into a ring. Objects are placed by their hash; each object is served by the first server found going clockwise.
- Adding a server (e.g.,  $S_4$  next to  $S_0$ ) moves only the keys whose clockwise path now hits  $S_4$  — in the example, only  $K_0$  moves.
- Removing a server ( $S_1$ ) remaps only its keys ( $K_1 \rightarrow S_2$ ); the rest stay put.
- Problem: random ring placement gives uneven partitions (one server can end up holding most data), worse when servers churn. Fix: virtual nodes — each server appears at many ring positions (e.g.,  $S_{00}/S_{01}/S_{02}$  for server 0). More vnodes = more balanced, but more metadata; tune per system.
- Real uses: DynamoDB, Cassandra (data partitioning, minimal rebalancing movement), Akamai CDN (edge distribution), Google load balancers (sticky connections with minimal re-establishment on failure).

### Key takeaways


- The invariant: only keys between the removed/added server and its neighbor on the ring are affected.
- Virtual nodes fix skew at the cost of metadata; vnode count is a tunable trade-off.
- Same hash function + same hash space for servers and objects is the whole trick.

### Beyond the video

- The video's ring model matches how DynamoDB/Cassandra actually hash tokens; Cassandra's default Murmur3Partitioner uses 256 vnodes per node (configurable). Per [cassandra.apache.org docs](https://cassandra.apache.org/doc/).
- Google S2 and CRUSH (Ceph) are alternatives for placement with different trade-offs — CRUSH uses a hierarchical pseudo-random algorithm rather than a ring. Per [docs.ceph.com](https://docs.ceph.com/).
- Consistent hashing also underpins cache sharding (e.g., in memcached-tier designs) exactly because it minimizes cache misses when nodes change — the "storm of misses" the video warns about.



hash ring with servers s0-s3 and clockwise lookup



virtual nodes: multiple ring positions per server

See also: [5](#)

## 12 12. Bloom Filters | Algorithms You Should Know #2 | Real-world Examples

Runtime: 5:39 · Published: Sep 2022 · Type: Video

### What it covers

- Semantics: false positives are possible, false negatives are impossible — the "probably yes" is what makes it probabilistic; it trades perfect answers for much less memory than a hash table.
- You cannot remove an item — the filter never forgets.
- Real uses: NoSQL LSM databases skip costly multi-file lookups for missing keys; Akamai avoids caching one-hit wonders (75% of pages are requested once; caching only on the second request cuts workload and raises hit rate); Chrome pre-screened malicious URLs before an expensive full check (retired as the list grew to millions); password validators reject weak passwords.
- Mechanics:  $k$  fast hash functions with even, random output; a bit array starting at 0. Insert sets  $k$  buckets; membership re-hashes and checks all  $k$  buckets are 1 — constant time.
- Walkthrough: 10 buckets, 3 hash functions — "apple" sets 1, 3, 4; "potato" sets 0, 4, 8; "pork chop" hits bucket 5 = 0 → confident no; "lemon" hashes to 1, 4, 8, all set by other keys → false positive.
- Size the filter against the expected number of entries to control the false-positive rate — the space-vs-accuracy trade-off.

### Key takeaways

- Zero false negatives; false positives are tunable via size vs expected entries.
- No deletions — once a bit is set it stays set.
- Ideal when "definitely absent" is the common case (DB/cache misses).

### Beyond the video

- Invented by Burton Howard Bloom (1970), "Space/time trade-offs in hash coding with allowable errors", CACM 13(7). Per [acm.org](https://acm.org/).
- Chrome's current Safe Browsing uses hash prefixes with a full-hash server query on match, not a Bloom filter. Per [security.googleblog.com](https://security.googleblog.com/) (Aug 2022).

## 35 35. Debugging Like A Pro

Runtime: 5:47 · Published: Feb 2023 · Type: Video

### What it covers


- Premise: debugging is rarely taught in school — most developers learn on the job, and it demands discipline; the video's steps are a toolkit to use "in any order — there's no right or wrong way to debug."
- Mindset: computers are logical, so there is always a logical explanation even when it seems impossible to find; being stuck is only temporary and persistence eventually resolves the issue; know your limits and when to seek help from more expert colleagues; and prioritize — some bugs aren't worth significant effort, so weigh each bug's potential impact and severity.
- When a bug report arrives, gather as much as possible: request a screenshot or screen recording from the customer, collect detailed steps to reproduce, and pull all relevant logs including error messages; then isolate the environment/context and try to recreate the problem on a staging server — "having a reproducible environment is half the battle."
- Investigation strategies once reproducible: use print statements liberally to construct a timeline of what happened and check it against the events you expected; some language ecosystems have excellent debuggers worth reaching for first — the video's example is Erlang/OTP's introspection capabilities.
- When you can't reproduce: understand why — some bugs only appear under production load, some only in specific race conditions under that load, and some need specific context/environment on the customer's device. Then: trace the code from the error line up the call stack, comb through logs for clues, build a timeline by following one failed request through its entire lifecycle, form a theory, add logging to prove it, deploy to production or ship to the customer, and repeat — a cycle that is "lengthy and frustrating, but persistence is key."
- When completely stuck: take a break and come back with a fresh mind (sleep on it, exercise, do a different activity); get a rubber duck — talking the problem out loud can trigger a light-bulb moment, and writing the problem out or emailing an imaginary mentor can generate new perspectives; get help — collaborate, since a fresh set of eyes and a different perspective often surface new solutions.

### Key takeaways


- Reproducibility is half the battle: screenshot/recording, repro steps, and all logs — then recreate on staging before touching code.
- For non-reproducible bugs, go forensic: trace the call stack, follow one failed request end-to-end, form a theory, and add logging to prove it in production.
- Prioritize by impact and severity — debugging discipline includes deciding a bug is not worth the effort.
- When stuck: break, rubber-duck it, write it out, and ask for a second pair of eyes.

### Beyond the video

- Rubber duck debugging is named after an anecdote in Andy Hunt and Dave Thomas's *The Pragmatic Programmer* (1999), where a programmer explains code line by line to a rubber duck before asking a human. Per [en.wikipedia.org](https://en.wikipedia.org).
- Erlang/OTP remains a strong fit for the video's "excellent introspection" claim: OTP 26+ ships a built-in, observable debugger and tracing via the observer tool. Per [erlang.org](https://erlang.org).
- Modern observability (distributed tracing with OpenTelemetry, structured logging) formalizes the video's "follow one failed request" advice for microservices; the practice of a formal postmortem/root-cause analysis is also standard in SRE culture. Per [opentelemetry.io](https://opentelemetry.io).



opening: why debugging isn't taught in school and the systematic checklist approach



gathering information from a bug report — screenshot, repro steps, logs — and reproducing on staging

See also: [30](#)

## 54 54. DevOps vs SRE vs Platform Engineering | Clear Big Misconceptions

Runtime: 4:44 · Published: Jul 2023 · Type: Video

### What it covers

- DevOps was coined by Patrick Debois and Andrew Shafer in 2009 as a culture, not a role: it bridges traditionally separate development and operations teams so they work together through the entire software development life cycle — from design through development to production support.
- Amazon's "You Build It, You Run It" is DevOps in action: the team that builds the software also deploys and maintains it, accelerating feature delivery to end users.
- Roles fluidly blend in practice: a cloud engineer strapping virtual instances with Ansible and running Kubernetes on top can look like a DevOps specialist; a sysadmin who builds tools and uses monitoring systems could be considered an SRE.