AVAILABILITY & RECOVERY · THE QUESTION EVERYONE ASKS

"What if the server *goes down?*" This is our answer.

The most natural worry about owning your own server is this: a single point that can fail. In a big cloud platform — the intuition goes — nothing ever fails. The intuition is understandable. It just isn't true.

Centralized cloud platforms have outages too — sometimes regional, sometimes global, sometimes for hours. That's no secret; it just rarely makes the marketing slide. The structural difference isn't "never fails" vs. "fails." It's a different question entirely: **Who sees the outage? Who can intervene? How traceable is the recovery?**

> *The honest question isn't "fragile single server or indestructible cloud?" It's: a controllable risk with a documented plan – or an uncontrollable risk with no influence at all?*

WHAT HAPPENS IN AN OUTAGE

01 How fast do we know something broke?

1

MONITORING & ALERTING

The alert reaches us before your team notices.

Your NexaStack is monitored around the clock, automatically: system health, running services, disk space, certificate validity, unusual load patterns, failed login attempts. Outages or irregularities trigger an alert **on our side immediately** — usually before a single one of your employees notices anything is wrong.

With a centralized cloud platform, you experience an outage the other way around: it reaches you through colleagues who can no longer work. What you have then is a status page and the hope that the vendor responds in some reasonable window. You can't intervene. You can't set priority. You are one of millions of customers in the same queue.

02 How completely can we recover?

2

BACKUP ARCHITECTURE & ABILITY TO RETURN

Two locations, different companies, daily – and actually tested.

Every NexaStack is backed up daily and fully automatically — and the backup is stored in a geographically separate, second data center. Concretely: if your production server runs with one host, your backup sits with a second, independent host — a different network, a different company, a different location. An outage on the production side never touches your backup. An outage on the backup side never touches your production.

The recoverability of those backups is tested regularly — a backup that has never been restored, even as a drill, is neither legally nor practically a backup. And because the entire NexaStack is built from globally used open-source software, the system can be rebuilt on any Linux server — with another host, in another region, if you wish. You're never dependent on one particular machine, one location, or one company.

03 How long does that realistically take?

3

HONEST TIMEFRAMES, NOT MARKETING PROMISES

Hours, not minutes — on a clear, documented path.

Here is the honest answer you'll never get from a big cloud vendor: a serious hardware failure means recovery in **hours, not minutes**. We are not a hyperscaler platform with a live hot-standby — and we don't claim to be. Instead, we work from a clear, documented recovery path out of the backup.

A single service glitch (say, an office module that won't start) is usually fixed within the same working hours. A serious server failure at the host is bridged by provisioning a new server and restoring the latest backup — realistically a matter of a few hours. A simultaneous total loss of both data centers — a scenario that essentially never happens — could still be recovered from the longer-retained backup history.

With a centralized cloud platform, you know none of these timings in advance. You don't know them during the outage either. Recovery takes as long as it takes — and you have no influence over order, priority, or pace.

> **THE HONEST LINE**
>
> What we promise: a documented plan, a human to call, geographically separate backups, recovery within hours, and an open-source stack that can be rebuilt anywhere — even without us.

STEP BY STEP

04 What a serious outage actually looks like

The sequence during a serious hardware failure — minute by minute.

Minute 0	**Automated monitoring detects the outage.** The alert goes to us, not to you. In parallel: first automated diagnosis, log analysis, narrowing down the cause.
Minute 0–15	**First assessment and communication.** We reach out to you actively: what happened, what's affected, how long we expect it to take. You don't call us — we call you.
Hour 1	**Depending on the fault: repair or restore.** Smaller issues are fixed directly. For a serious hardware failure: provision a new server at the host, restore the most recent backup.
Hour 2–6	**Recovery & function check.** Data is restored from backup, services brought up systematically, every module checked, and you're notified as soon as your team can access it again.
Next day	**Follow-up and documentation.** A written incident report: cause, timeline, measures applied, recommended next steps. Transparent, traceable, on your desk.

Figures for the typical serious-hardware-failure case on a running NexaStack. Pure software issues are usually resolved much faster. These are real-world operational guide values — not a contractual availability guarantee (SLA).

THE STRUCTURAL DIFFERENCE

05 Two ways to experience an outage

WITH CENTRALIZED CLOUD PLATFORMS

You notice the outage when your team complains. You have a status page and a hotline. The vendor — not you — decides the window, priority, and order of recovery. Your company is one of millions in the same queue.

WITH YOUR NEXASTACK

The alert reaches the provider before your team notices. You get a call with the situation and a time window. The recovery path is documented. Afterward, a written incident report. A human on the other end — not a ticket system.

What we honestly don't promise

No "99.99% guaranteed uptime." No hot-standby redundancy in standard operation. No zero-downtime on hardware failure. Hyperscaler platforms with billion-dollar infrastructure can only rarely hold to that exact wording either — the difference is just that there, it's printed on the slide.

WHAT WE DO PROMISE

A documented plan, a human as your point of contact, geographically separate backups, recovery within hours, and an open-source stack that can be rebuilt anywhere — even without us.

WHAT IT ADDS UP TO

06 No system is unbreakable. The only interesting question is what happens when it breaks.

With a big cloud platform, the answer is structurally the same as everything else in that architecture: you wait. No insight, no influence, no point of contact, no time window. The outage is invisible until it's over.

With your NexaStack, the answer is the opposite: the outage is visible, assessed, and time-bounded. You get a call, a plan, a report. Your data sits in two separate locations. Your infrastructure can be rebuilt on any Linux server in the world.

> **THE BOTTOM LINE**
>
> An outage at NexaStack is an event with a point of contact, a time window, and a plan. That is worth more to your business than any uptime percentage you can't verify.

STILL HAVE A QUESTION?

Ask us on a strategy call.

No obligation, no sales script. — nexastack.co/en/strategy-call