

algoseek Colocation Service Manifesto

AKA

A Paranoid's Rules for Thriving

Andy Grove, co-founder of Intel, once famously said:

“Only the Paranoid Survive.”

But his philosophies to turn a hightech business into a mass production machine is, in our view, fundamentally flawed — and Intel's rise and fall was doomed from the very beginning. At algoseek, our eyes aren't set on mere survival; we are here to thrive. We refuse to settle for Andy's serving “a three-minute soft-boiled egg” at “an acceptable quality level and the lowest possible cost” mindset, as mentioned in his other book ‘High Output Management’. We aim at the highest quality, unleashing our Paranoid Rulebook to crush downtime and data loss, ultimately delivering top-tier quality at the

lowest costs in the industry. After all, when our clients thrive, so do we. We've been building servers, writing in-house software, and collecting and computing market data for nearly two decades, and for clients for 10 years. Two decades are not very long. If you buy commercial Dell servers and replace them each time there's some sort of failure, you won't learn much during such time. But if you dig into every problem you encounter, from data congestion to CPU meltdowns and memory quirks, trying to figure out how to prevent it from happening at all, then you evolve from “just another vendor” to a true expert vendor.

So, what's in algoseek's Paranoid Rulebook?

We don't want to smother you in endless tech specs, but we are happy to share some of our approaches. Buckle up.

01 Server and Cabinet Building: OCD in Action

“Details matter, it's worth waiting to get it right.”

— Steve (Apple co-founder and former CEO)

Glue the Damn Fan!

Let's face it: like Intel, most manufacturers skip work whenever and wherever they can. So if you look inside a server, the fan is usually fastened with some cheap plastic pin that comes off easily. Vibration from the rattling of a loose fan will likely kill your server within a year. As a certified paranoid, what do we do? We glue the damn fan! So it stays where it belongs. And not just the fan. We glue cables so they

never wander into a fan's spinning blades, and if we ship a server, sometimes we glue the memory too. (Yes, we're serious.) That's not all — we also wipe everything clean, arrange cables neatly, snip zip ties to perfection.... Obsessive? You bet. But it's the difference between a rickety, failure-prone machine and one that runs for years without complaint.

Alu-MEH-nium? No Thanks

Pop quiz: What's the most critical component that server manufacturers cheap out on when they absolutely shouldn't?

If you answered the cooling system, congratulations — you're already thinking like a paranoid.

Take Dell, for example. Even in its high-end servers, it still ships standard-issue aluminum heat sinks. Why? Because aluminum is cheaper, lighter, and good enough for the average workload. But good enough doesn't cut it for us. We're obsessed with CPU performance, and these alu-MEH-nium coolers just don't make the grade.

So, we took matters into our own hands. We handmade our first batch of copper heat sinks—heavier, pricier, but infinitely better at dissipating heat—and sent them to a factory in Asia to produce custom copper coolers for our servers.

But we don't stop there.

- We meticulously apply high-quality thermal grease when installing vapor chambers. No sloppy paste jobs here—just perfect, even applications for optimal heat transfer.
- We validate every install because a poorly seated cooler is worse than no cooler at all.
- We never settle for off-the-shelf cooling solutions when we can build something better ourselves.

Because when it comes to protecting our CPUs, we don't skip, we over-engineer.

A Healthy Obsessiveness with Testing

Suzanne Massie taught Ronald Reagan to “trust, but verify”, and that's exactly our take on these soulless machines and parts. Not even a cable can be trusted without testing.

In fact, we do pre-installation testing not once, but **twice**. First, we test during initial setup at the office — confirming the OS sees all disks, running memory tests, etc., and, of course, testing every cable. Then, after bringing the server to the data center, we rerun the entire suite of tests again before inserting anything into a cabinet.

“Experience changes a person.”

(And apparently, a server too)

Servers don't emerge from life's trials unscathed. From office to data center, they're lifted up, put down, tossed on road bumps, and occasionally kicked by a careless foot. Bad things could happen during the journey — memories come off socket, hard drives wiggle loose, and cables lose their snug fit... That's why we test everything twice: first in the office, then again after the server survives its mini-odyssey. Because as stoic as a server may seem, it's our job to catch what's been shaken loose along the way.

The All or None Rule

So what if a server that passed with flying colors in the office fails even one test at the data center? We bring the entire thing back. Period. As a rule, we don't troubleshoot in data centers. With all the server fans and aircons, data centers are noisy, cold, and frustrating — perfect breeding grounds for sloppy “quick fixes” that come back to haunt you. Instead, we fix issues thoroughly in our own space, where patience and clarity reign.

Cabinet Planning: Neat Freak's Nightmare

We've all seen those data center horror stories: cabinets wrapped in messy cables like a ghost gear dumping ground, and an engineer was crawling on the floor, a flashlight in mouth, desperately trying to figure out which cable goes where. But never us. At algoseek, we'd rather walk barefoot over Legos than let our cabinets degenerate into such a nightmare. Here're some basic rules on how we keep it tidy:

- Combine high-power servers and low-power bulky servers (such as four-unit storage servers) to make full use of both space and power capacity.
- Put bulky servers on top, short and switches at bottom. So you have more room at the bottom.
- All the cables must be in the back. Don't use switches with front ports.
- Arrange servers of the same depth together. Mixing different depths leads to blocked ports and general chaos.
- Plan ahead with detailed cabinet planning diagrams all the way down to each cable's start and end points.
- Keep detailed logs that record every activity in the data center.

It may sound straightforward in theory, but in practice, it's like herding kittens. Still, it's worth it — our cabinets remain neat, efficient, and easy to manage.

Some Extra Rules

In the world of building reliable, cost-effective servers, we've already miles ahead of the competition. But we've got a few more tricks. For example, we found that motherboards and memory cards produced in the same country tend to have better compatibility and lower failure rate. So when building servers, we always check production year and location to match them. We also use network cards with their own processing chips for heavy-duty data-capture servers. It's not just about sparing the CPU some workload. If you rely on the machine's CPU, you're also using the PCI Express, which means latency, pollution of CPU cache, and all challenges of accessing the memory.

These may seem like small tweaks, but they're the kind of details that make a massive difference in long-term reliability and efficiency. **At Algoseek, we're committed to leaving no stone unturned — because excellence is all about the details.**


02 Data Center Maintenance: Anal-Retentiveness Takes You a Long Way

“Trifles make perfection, and perfection is no trifle.”

— Michelangelo (Renaissance artist)

If it wasn't obvious by now, we're massive believers in planning. “Fail to plan, plan to fail,” as the saying goes. In our world, it's more like, “Fail to plan, end up crawling through a jungle of cables while slowly losing your soul.”

Below is how we prevent our data centers from devolving into the modern-day equivalent of a dystopian labyrinth:



Plan Down to Every Step You Take

We draft a detailed script at least a week in advance of any data center visit — every single step is mapped out. We also print everything: cabinet diagrams, configurations and script, (and yes, on paper), because scrolling a laptop screen under harsh fluorescent lights and freezing AC is nobody's idea of fun.

How nitpicky are we? Let's just say we know precisely which cabinet, server, unit, slot, and port for the both ends of each cable — plus the ideal cable length. Yes, it's nerdy. But trust us, you'd prefer that over a cable scavenger hunt at midnight after market close.

Double/Triple Testing Any new gear goes through multiple test rounds. Pre-deployment, on-site, and again back at the office. Because if Murphy's Law taught us anything, it's to expect random failures at the worst possible moments.	Detailed Logs We document the who, what, when, where, and why of every action. Overkill? Maybe. But when something goes wrong months later, we know the exact footsteps to retrace.
--	---

Offsite Backups for Everything We're paranoid, remember? If the data center spontaneously combusts — or somehow gets swallowed by a cosmic wormhole — we can rebuild it in days, not weeks or months. That's why every script, every configuration file, and every log are also stored safely offsite.	Spares Are Not Optional We stock spare parts the way some people hoard coffee beans or toilet paper: excessively. Hardware loves to fail when you need it most, so we keep extras on hand — whether that's an entire server or just the cables and memory modules. Better paranoid than sorry.
--	--

Never Let the Little Stuff Slide

Happen to not having cable of the right length? We'll use a longer one for now, note it, and definitely replace it at our next visit. Because a single shortcut is the start of a slippery slope — one day, you'll wind up in a labyrinth of unexpected issues and ad-hoc fixes. Better to stamp out that chaos before it begins.

When all is said and done, our data centers remain neat, efficient, and — dare we say — boringly predictable. And trust us, when billions of your trading funds are at risk, “boringly predictable” is exactly the vibe you want.

03 A Paranoid's Take on Preventative Monitoring and Data Integrity

“When you monitor for failure, you already failed.”

— Art (Head of algoseek real-time arm)

We're paranoid, and proud of it. So naturally, we don't trust third-party monitoring tools. Their design approach—monitoring for failure—is flawed from the jump, and they don't know our business like we do. When you're juggling low-latency market data and mission-critical analytics, you want zero service downtime and absolute data accuracy.

Our solution? We developed our own monitoring and testing tools from scratch, so we can dig into the exact nooks and crannies that matter—and ignore the messy false positives that don't. After nearly two decades of reading logs, we've got a keen sense of where problems lurk and how to nip them in the bud before they blow up.

Early Indicators: Catch Problems Before They Become Failures

Our in-house monitoring dives deep — like really deep — into low-level operating system logs (think Windows event logs, /var/log/messages on Linux, and everything else in between). We're always on the lookout for early warning signs that most monitoring tools might wave off as “non-critical.” Turns out, “non-critical” can become “major outage” overnight if left unchecked.

Half-Speed SSD Any new gear goes through multiple test rounds. Pre-deployment, on-site, and again back at the office. Because if Murphy's Law taught us anything, it's to expect random failures at the worst possible moments.	Example 1
---	------------------

Disk Write Retry During an automated backup, our system detected a disk write retry. Was it a faulty network card? An iffy cable? A corrupted sector? Many possibilities, and all of them could cause serious trouble at a later time. Even though the process eventually succeeded, we investigated immediately, tested all suspects, and ended up replacing a memory card. Crisis averted.	Example 2
--	------------------

Congestion & Pause Frames We monitor for network congestion like our reputations depend on it—because they do. A switch issuing a “pause” command doesn't necessarily mean lost packets yet, but it's a giant waving flag that your network might be heading for a traffic jam. That's our cue to make some bandwidth upgrades before any real damage is done.	Example 3
--	------------------

Pro Tip: Knowing what not to worry about is just as crucial. Case in point: Linux and certain Samsung SSDs. They produce a flurry of dubious system errors that can safely be ignored — provided you've already confirmed no actual fault is lurking underneath.

Data Integrity: The Ultimate Goal

Spotting early warnings is one piece of the puzzle; ensuring your data packets don't get lost in transmission is another. We post “listener” servers on each data link, capturing every packet to ensure no invisible drops occur. Missing sequence numbers = dropped packets. We also compare exchange timestamp with our server timestamp to analyze latency.

We aim at **absolute accuracy** of our market data. Ask any hardcore quant or high-frequency trader: losing data mid-stream can ruin your day—and your profit margins. We'd rather be safe than sorry.

To sum up:	Early detection	=	Fewer nasty surprises
	Custom-built tools	=	Zero blind spots
	Dedicated link listeners	=	Bulletproof data
	Decades of experience	=	Knowing exactly which “chatter” in the logs matters

After all, the only thing worse than a failure is not knowing it's on the horizon. By the time most conventional tools say “Oops, something's wrong,” it's already cost you time, money, and endless frustration. We prefer to catch that “Oops” when it's just a faint whisper, well before it becomes an ear-splitting alarm. Because in our line of work, just a hint of paranoia can save a mountain of trouble.

04 Healthy Cynicism on Commercial Products & Services

“The prevalent organizational structure of relying on functional silos and short-sighted performance metrics often pits individual divisions' interests against those of the broader company and its customers. As a result, consumers bear inflated prices for offerings that are, at best, mediocre.”

— Chris (algoseek CEO)

Hang around long enough, and you see how big-name vendors spin run-of-the-mill products as “industry-leading” — and then charge a premium for them.

Dell Servers: Future Upgrades Never Not Apply

Dell's custom motherboards and CPUs lock you in. No standard slots, often no replacement parts—if anything fails or you need more capacity, guess what? New server. Even if you're merely short on RAM, your best bet might be to buy a whole new box.

Case in point: We once purchased Dell servers for ITCH data collection. They came with a poky Dell Perc RAID controller in a special non-standard slot. No direct upgrade existed, and Dell offered no alternative. In the end, we literally grafted in a simple HBA card, and our speed tripled immediately — something Dell's official gear couldn't provide.

Commercial Monitoring Software: The Half-Truth Hall of Mirrors

Commercial monitoring tools often claim to be comprehensive, but they rarely offer the deep visibility we need — like congestion stats on each server port, not just the switch port. Worse, they can be glitchy themselves, leaving you troubleshooting the monitor instead of your actual setup.

Why Most Colocation Providers Stick to Dell and Third-Party Monitoring

One word: inertia. It's easy, it's “safe,” and it doesn't require staff to learn or do much. Procurement teams don't get fired for buying from a big name, and support departments keep racking up budget by constantly swapping out failing gear. Everyone wins—except the customer footing the bill.

05 Why Algoseek is Your Best Option

“We are what we repeatedly do. Excellence, then, is not an act, but a habit.”

— Aristotle summarized by Will Durant

If you've come this far, you might be thinking, “These folks are really something else.” And you'd be right. At algoseek, we believe in:

- **Custom-Built Servers:** Higher performance, longer life span, and lower long-run total costs.
- **Preventative Monitoring:** Catch early signs of problems before anything fails.
- **Client-Centric, Perfectionistic Culture:** Market shifts and technologies change. Our longevity lies in striving to meet clients' need and evolving with our clients.
- **3-Dimensional, “Fuzzy”, Flat Structure:** Our productivity roots in each teammate. We adopted an organizational structure and incentive system that are both flexible and align our teammates' best interests with algoseek's best interests (and therefore our clients' best interests).

In short, we do the “paranoid” work so you can focus on what matters — **your trading, your analysis, and your growth.**