



# Corvil

# Review From A Customer

From IT Central Station, the leading review site for enterprise technology solutions.

---

## Review by a Real User

Verified by IT Central Station



Senior Network Engineer at a financial services firm with 501-1,000 employees

**SeniorNea968**

### WHAT IS OUR PRIMARY USE CASE?

Latency analysis is our primary use case.

### HOW HAS IT HELPED MY ORGANIZATION?

What we've done a lot of work on is tick flow. We generate prices on various instruments, for example DAX Futures. We need to understand, internally, how long does it take for the DAX Future tick to leave the exchange and to exit our pricing infrastructure, which generates the prices and feeds them into the apps that the clients use. We need to know at every step what that latency is. What we've built in Corvil is a dashboard that shows that. It shows the latency from the exchange to us. Then it shows the latency from us into what we call our MDS system, and from there into our calc servers, which actually do the grunt work of generating the prices, and from there into our price distribution system. At each point, we have a nice, stacked view that shows the latency of each component. We can look at this and say, "Oh, actually it's our calc servers that are causing the most latency." A good example is that recently our platform guys did some analysis on the kind of improvement we could expect if we put Solarflare network cards in our servers. The analysis showed we could get a 50 microsecond improvement. By using our Corvil data, we could say that, while Solarflare would give us 50 microseconds of improvement, our calc servers alone are generating something like 20 to 30 milliseconds of latency on a bad day. So in the grand scheme of things, spending all this extra money to save 50 microseconds isn't going to cut it when there is a lot more scope to save latency by just rewriting the code on our calc servers. That's a good example of how Corvil helps us. It gives us that kind of level of detail so we can pinpoint exactly where the latency is. Corvil helps us determine where to focus our performance improvement efforts. With the tick flow process, we've split it to cover four key parts of our infrastructure. We can see, straightaway, which part is generating the most latency, and that tells us where we should focus our efforts, where we should spend time and effort and money. You have to build that view. You definitely can't take it out-of-the-box and it will come up with that view. You have to understand how your traffic flows, and build that view. In addition, we do venue performance analysis. A good example is FX pricing. We take all the OTC pricing from various liquidity providers like the Tier 1 banks. Key metrics for us with FX are things like sending-time latencies. We look at that. We always knew anecdotally that one of our feeds was really poor when it came to latency. But without Corvil, we didn't have the numbers to prove it. We could just tell by looking at the quotes from the others to know how far out this particular feed was and from that deduce that the latency was really bad. Corvil helped us show that information in a nice, graphical manner and gave us some metrics to justify a scheme of works to improve matters. Corvil makes it really simple to extract the information required. For example, we are asked sometimes asked something in the line of "Could you supply us with quote ID where you observed these issues," maybe, in some respects, thinking it would take us a long time to get that information. But we can literally, in two clicks, export the spreadsheet and send it to them. Having latency information helps us improve order routing decisions. A lot of our trading is automated. It's not that the Corvil tool is used to directly feed the automation, but it has provided visibility to allow us to support the process. For example, we can determine precisely when a venue was down and what trades were impacted. That's definitely

helpful. We never had that kind of correlation in the past. It was a case of a trader coming up to us and saying, "Did you have a problem with XYZ at this particular time?" We'd have to dig out multiple logs from the network and other infrastructure components and try to figure out what was going on at that time. It allows us to narrow down on the problem a lot quicker. We've got a copy of all the messages, we know what went on at each point. Corvil has definitely helped reduce incident diagnosis time. Just the fact that it's so easy to pull out captures or the actual messages, whereas before that was probably the thing that would take us the longest, just getting the data. Before you can start looking at the data you actually have to get the data. Now, it's easy. We just say to the user, "Tell us the time XYZ happened." We can find it, we can zoom in on it, we can extract the messages. It happens a lot of when we're talking to third-parties about latency. In terms of Corvil reducing the time it takes to isolate root causes, quite a few times we've been able to look at a stream in Corvil and straightaway we can identify what the issue is. A good one is always batching. We can always tell, nowadays, when latency is being caused by batching. We can download the capture straightaway, take a look at it, and it's very easy to see that when a particular venue is sending multiple quotes together, queueing up one behind the other, rather than as they are generated. We're definitely able to diagnose things like that really quickly now, whereas before it would be a big struggle to get the data in the first place. If I had to put a number on how much time we save it's at least a good three or four hours. Finally, the dashboards have helped reduced the time it takes to answer business questions. We sit down with some of the application support teams and say to them, "What do you want to see?" So that's definitely there immediately for them so they don't have to be generating their own stuff. I don't know how much time it saves but the ability's there for them to log in, have a look at their dashboards for their particular products. The information is all there for them.

### **WHAT IS MOST VALUABLE?**

The performance metrics are pretty good. We've got everything from the network layer to the actual application layer. We can see what's going on with things like sending time and batching. Time-series graphs are very good for performance analysis. We can do comparisons. We can do minus one week, or we can say this is the latency in the last 24 hours, and this was the same 24-hour period a week ago and overlay the two time-series graphs on top of each other, so we can see the difference. That's a really powerful tool for us. We make improvements in the network all the time, so it's useful to be able to quantify what the effect of a change was. The performance analysis is pretty strong. The ability to dig into the messages is definitely a valuable feature. We don't have any other tool like that, at least in the network space where I work. It is very useful to be able to get such granular information. As network engineers, when we deal with packets, we can dig down into the TCP level and that's about it. We can't actually decipher the actual messages. For latency analysis, sometimes you're dealing with Exchange-driven time stamps and you actually do need to dig down to that level of detail. That's been the most valuable feature for us. The analytics features of Corvil are really good. It's the fact that you can build your own metrics. The fact that, as long as you know what the field is in the message, you can build your metrics based on that field and that is very good. It means you can do the analytics that you actually care for. You can customize it in a certain way, which is good.

### **WHAT NEEDS IMPROVEMENT?**

There is definitely room for improvement in the reporting. We've tried to use the reporting in Corvil but, to me, it feels like a bolt-on, like not a lot of thought has gone into it. The whole interface where you build reports and schedule them is very clunky and I find that, whereas on the GUI you can pull out all the metrics you want and it's very flexible and nice and easy to customize, the reporting is not very intuitive. And the metrics don't display nicely, as they do on the GUI. It's always a case of trial and error. You add a metric to a report, generate the report, and find, "Oh, that doesn't look quite right." Then you have to go back in, edit the report and fiddle around. There's no preview button, which would be useful. It's just very clunky and hard to use. We don't use it because it's not great. Alerting isn't great. It's very flexible in that you can put it in protection periods so it's not constantly giving you false positives. Sometimes if there's a latency spike, and it happens just once in five minutes, we might not particularly care about it. You can configure that kind of stuff with the alerting. But there's other stuff that's really not well thought through. For example, the email address that you use - this is a really simple thing that I've mentioned to Corvil a few times - I'm pretty sure you can only put one email address in. And that's for all kinds of alerting on the box. So we infrastructure guys might be interested if a power supply goes on the appliance, but for somebody in our application team, they don't care about that. They're more interested in: Did the latency go over X in a certain period of time? But because there's only one email address we can put in, it's very hard to manage that. The whole alerting and reporting side of things - more the reporting than alerting - definitely needs a bit of work. Corvil could do a bit more around system performance visibility (although I believe there is an action widget now which can help). It's quite difficult to see, sometimes, how hard your Corvil is working. When we had a very busy feed that chucked out a lot of data, it wasn't working very well on Corvil. We had to raise a case for it. It turned out to be that, in fact, we were overloading Corvil. It was very hard to determine that without raising the TAC case with Corvil. Their engineers have the CLI commands to dig this stuff out.

**FOR HOW LONG HAVE I USED THE SOLUTION?**

One to three years.

**WHAT DO I THINK ABOUT THE STABILITY OF THE SOLUTION?**

We've had no problems at all with its stability. We haven't had a failure, we haven't had to replace one, there's no disk failures or anything like that. You'd think a box like this, with disks constantly in use, that they would be the first things to go. That's been fine. We did have one issue with a decoder causing the whole capture daemon to crash and restart. That was fixed pretty quickly by Corvil. We said to them, "We keep getting these errors," and they took a look at it and said, "Yeah, it's the decoder." They were able to reproduce it in the lab. We gave them a sample of the data and by the next month - they do monthly Decoder Pack updates - they fixed it. You can tell when that happens because you see in the Decoder Pack release notes "such and such decoder: Improved the robustness of the decoder." That invariably means it was crashing the appliance. That was the only reliability issue we've had but luckily that was specific to one decoder.

**WHAT DO I THINK ABOUT THE SCALABILITY OF THE SOLUTION?**

I don't have anything to compare the scalability with. Our MDS spits out a lot of data, and it doesn't help that their messages have a lot of fields. This system was causing the Corvil to overload which we had to subsequently manage by being very prescriptive about what flows Corvil sees. You could argue, in some respects, that it's our fault because in the PoC we probably should have tested it against this particular MDS system. But Corvil doesn't publish the numbers. It's not obvious where to find numbers like "this is the number of flows an appliance can handle," "this is the number of sessions." So I don't know if scalability is a problem, but Corvil could do more around giving customers the information about how scalable a box would be.

**HOW ARE CUSTOMER SERVICE AND TECHNICAL SUPPORT?**

Their technical support is really good. I used to work for big organisations where I'm used to dealing with a lot bigger vendors where, when you raise a ticket it goes into this "first-line" type of queue and they don't really do a lot with it. Effectively, they just give you a ticket number, make some notes, and then pass it on to some engineering team. What's really nice about Corvil is, first of all, you don't have to fill in a web-based form or call some number. You can just email support. You get hold of someone straight away who knows what he or she is doing. There's no passing around. You're straight in touch with someone who understands and can support the product. They don't start asking for things like serial numbers. Straight away you're into, "This is my problem, here are my logs," and they'll help you fix it. And when they do fix it, they're very knowledgeable guys. Most of the stuff is fixed by the person I'm emailing. In a few circumstances they have to go back to "engineering," but most of the time they'll fix it there and then or tell me how to do it. Support is really good. Maybe as Corvil gets bigger that will change but I hope not!

**IF YOU PREVIOUSLY USED A DIFFERENT SOLUTION, WHICH ONE DID YOU USE AND WHY DID YOU SWITCH?**

We didn't have a previous solution. We were running custom scripts and parsing logs before Corvil. We had some latency-analysis tools but they were all things built in-house. When I joined the company two years ago, the Corvil PoC was going on, so the decision to move to something better was made before my time here. The way they were doing it before Corvil wasn't very scalable. It needed to be centralized. Each team was doing things their own way. It wasn't very joined up. You ended up with a lot of duplication or slightly different ways of doing it. These guys have got other things to be getting on with, rather than doing the analysis. That was the big gap. We had multiple teams using FIX, for example, and one person might build something that showed some kind of metric, how many market-data snapshots we get in X period. Some other team might build exactly the same thing, but in a slightly different way. So we just ended up with this myriad of tools and there was no compatibility between them and there was no easy way of comparing them. It was clear that we needed some centralized way. That was one problem. The other problem was that none of these tools - they were mostly scripts that people wrote - could work at that nanosecond precision that Corvil gives us. Market data has moved on and the trading has moved on where that kind of granularity matters. Whereas before we could probably get away with it - the millisecond range was okay for us - now, we need to know things like spending this

amount on a Solarflare card adds 50 microseconds. We need to be able to measure that 50 microseconds, which was something we couldn't do before.

## **HOW WAS THE INITIAL SETUP?**

Our initial setup was straightforward. 4 appliances with a central manager. Our SE from Corvil came in and helped us do the initial setup. It was very good because he let me do the work, just telling me what to do. His attitude was, if you do it, then you will remember how to do it, rather than having him do it for me. That was really good. It was very straightforward. We gave him a few flows that we were interested in and he showed us how to set them up. We had a few follow-up questions which he handled via email. It was very straightforward. Altogether, the setup took a couple of weeks to get it to where we wanted it to be. Some of that was scheduling because we had to schedule to get this guy in, but from initially unpacking the box and sticking it in our data center, cabling it up, to actually spitting out useful data, it took a couple of weeks. In terms of our implementation strategy, we did a PoC beforehand, which gave us an idea of how we wanted it set up. From there, we came up with a plan of having one per DC plus the central manager. Once we made the decision to buy Corvil, we put a little bit more effort into thinking about what places in the network we wanted to monitor, where were the key points we wanted to put SPANs and TAPs in. We also thought about how we were going to use our aggregation switches. In terms of Corvil itself, apart from the PoC, we didn't spend a lot of time with it. In some respects, we didn't know all it could do at that point. We knew the basics and we had to make a decision on what we saw in the PoC. So some of that came out during setup. We knew, for example, that we wanted to monitor our FX feeds. But when we got to the day where the SE came, he showed us that we could display this kind of view and have these kinds of metrics displayed. It was a case of, "Oh, you know that's actually good." And then we fine-tuned it from there. For the deployment, it was pretty much just me involved from our side. Since it was deployed and configured, it has been just me maintaining it. That is a bit of a sore point because it makes me something of a single point of failure. We have to sort that out because it's not sustainable in the long-term. We have about 15 people using it. They're mostly application development guys. Obviously, we network guys use it for troubleshooting and the like. But in terms of the non-infrastructure teams, they are mostly application developers or those who run those teams.

## **WHAT ABOUT THE IMPLEMENTATION TEAM?**

Implemented via a mixture of Corvil themselves and in-house. Corvil team was excellent.

## **WHAT WAS OUR ROI?**

Our ROI is probably more cost-avoidance than anything. Corvil allows us to see where our issues are and then we don't waste money on areas that aren't going to give us the biggest gains. The whole Solarflare example I mentioned elsewhere is a good one.

## **WHAT'S MY EXPERIENCE WITH PRICING, SETUP COST, AND LICENSING?**

My perspective can be summarised as Corvil is a bit expensive but you get what you pay for. I like the way they've decoupled the hardware now. That makes sense because hardware's a commodity now. It's natural that they had to do that. Everything's based on the licensing side now. The way they do the packs is fair. It is very flexible in that we are not charged per decoder; we are charged for a certain pack. Whether we use 1 decoder or 20 decoders, as long as they're in the same pack, there's no extra charge. It's expensive but it's no more than I would expect for this kind of product. They throw in some things for free, which is nice. We've just started doing the UTC clock sync, where you can use the Corvil to analyse your time signals and generate a report. They don't charge for that kind of stuff, which is nice. I remember I was pleasantly surprised when I found that out. Expensive but fair is how I'd summarise it.

[Read 7 reviews of Corvil](#)

### **WHICH OTHER SOLUTIONS DID I EVALUATE?**

We evaluated the product against ExtraHop which you could argue was not the correct thing to do as the products are targeted for similar but different use cases. We, quite late in the day shifted the PoC towards comparing Corvil with Velocimetrics but as we were running out of time and already liked what we saw with Corvil, we decided to proceed with them.

### **WHAT OTHER ADVICE DO I HAVE?**

Definitely understand how your traffic is flowing and remember that Corvil won't magically fix your latency issues, rather help you identify where they are and what impact they are having on your business. Corvil is only as good as the data you put into it, so if you're monitoring in the wrong places, for example, you're not going to pick up the whole story. The advice I would give is to educate your users. You don't just spend money and, miraculously, all your issues are fixed. Rather it will help you understand where your issues are. But to do that you have to know the various application flows, how they work. Make sure you're monitoring the right places and that Corvil has all the right data, to a high level of detail. Ideally we would like to use Corvil for other things, apart from just pricing-related stuff. For example VOIP and PTP. I'd like to use it more on the enterprise side. We've got a file-sharing issue at the moment between two locations and if we could get that traffic into Corvil it might be useful to help show us what's going on. There are some powerful analytics on Corvil, especially for network engineers, like the TCP side. We want to use it more for that. Maybe that means we buy another Corvil that is dedicated to those uses. That's where we'd like to go. Regarding Corvil and productivity, in a small company like this we're not as siloed as they would be, say, in a big bank. Someone in my role is expected to know a bit about the trading side and a bit about protocols like FIX, etc. Corvil has changed the way we work in that it means that we need to know more about how the whole business works. In terms of what it does, I'd give Corvil a ten out of ten. I've never seen a tool like it.

[Read 7 reviews of Corvil](#)