

We hang together or we hang alone

Joe Skorupa is a research vice president at Gartner Inc. and an expert on the increasingly vital interplay of networks and applications in the blossoming world of SOA. He spoke with F5 World's editorial director, Bill Laberis. A single-page version of this discussion is available in F5 World Magazine, Winter 2007 edition.

Q: How is the role of the network changing?

A: There have been some significant changes in what the network needs to do to deliver appropriate application performance and availability. As more business processes [and] business-critical applications moved over to an Ethernet IP network, what we saw emerge was the concept of *application awareness*—meaning that the network could recognize whether the traffic was Oracle traffic or file transfer or SAP or email. But it didn't understand the content or the context of the traffic. All Oracle traffic looked the same.

Now virtually all business-critical processes are network based. At this point, we see the emergence of a concept that we call *application-fluent networking*. By way of example, two people could be talking in Spanish and I could be aware that they're talking in Spanish, but I have no idea what they're saying and I can't contribute to that conversation. If I'm fluent in Spanish, I can become actively involved and hopefully can contribute.

The network now needs to be an active participant in the delivery of the application. It needs to understand the underlying protocols; it needs to understand the messages; and it needs to understand what's going on with a particular transaction in order to understand how to best treat it. So all Oracle traffic isn't the same. For example, order entry traffic might be much more important at the beginning of the month and shipping traffic may be much more important at the end of the quarter. You have to be able to interact in a much more integrated manner with the application.

Q: So this means that the networking and application groups need to work together more closely?

A: If they don't, the application won't work. And by won't work, I mean you may need a response time of five seconds a transaction to get the number of transactions per day that a particular group has to perform. If it's five minutes per transaction, functionally you get the answer back, but it doesn't matter. You're going out of business—you can't serve your customers. So if they don't work together, you won't have the performance, availability, and cost structure that you need.

Q: Are there roadblocks to networking and applications working together?

A: The roadblocks are cultural and organizational. They speak different languages. Application developers don't think in terms of roundtrips and latency and messages.

They think about a business transaction they're trying to automate. The people don't know one another. They've never had to work with one another. It's always been, "We write the code, we throw it over the fence, and it works. If it's a problem, we just tell the networking guys to buy more bandwidth."

What they [the networking team] usually wind up doing is deploying a combination of application delivery controllers and WAN optimization controllers to compensate for inefficiencies in the protocols. For example, inefficiencies in the Web servers—network-induced problems in terms of latency and bandwidth constraints. But quite often it's after the application has been deployed. It is an unbudgeted expense. Now I'm the network guy and I've got a million dollars unbudgeted. I'm wiped out for the rest of the year.

Q: Why don't these problems show up earlier, say in prototyping?

A: Because it's typically prototyped [on] one terminal, one workstation, on a dedicated gigabit Ethernet, dedicated high-end server. It has nothing to do with the real world.

Q: How can this disconnect be better addressed during development?

A: We recommend something we call the *lifeboat development process*. When you kick off the project, bring application architecture, application development, the operations folks, the network architecture and operations folks, storage, security together in a room, with a representative of the business stakeholder. The instructions are, "You are all in a lifeboat lost at sea. There will be no individual survivors, meaning no individual heroes. You will either all reach shore successfully, or you will all die a terrible death."

Q: We hang together or we hang along.

A: Bingo. So I don't want to hear, "Hey, it works fine on my workstation, it must be the network's fault," or, "If only those storage architects knew what they were doing, they could build an infrastructure that has the I/O requirements I need." You're in it together. And by the way, the person who will ultimately judge whether you live or die is the business stakeholder."

This gets an iterative process going very early, where all of the involved people—all the guys in the lifeboat—begin to sit down with that business stakeholder and ask questions: What are the business processes we're trying to automate? How many transactions a day do you expect to run? Where will the people be that are running the transactions? What's the peak load? What time of month does that occur? What's the average load? Now folks are beginning to scope the impact. How many messages are going around? How much data has to be moved around? Now the application and storage guys can talk about how big is the data store? What are the I/O requirements? What database is being used? How should it be structured?

The networking guys can begin to think: What's the impact in terms of bandwidth on the network, and what's the impact of the network on the application in terms of latency?

They can sit down with the development guys and say, “By the way, if you’re structuring a browser-based application, here are things you can do.”

So you begin to have this iterative process of, “What are the assumptions you’re making? Let’s look at it.” The networking guys can supply a device that we call a WAN emulator. You can put it between a developer and the server and program them to provide the network characteristics that end users will actually see. You can use this in the very beginning when the first prototype is tested and see the impact. You can determine whether performance is going to look OK from a latency standpoint. You can begin to see how much bandwidth you were using.

It gives the developers much earlier visibility. Remarkable as it is, most developers, if you tell them 250 milliseconds latency and a megabit of bandwidth, that’s a very abstract concept to them. But when it takes them 30 seconds to run a transaction they’re used to taking a tenth of a second, you don’t need to leave [the WAN emulator] in place very long for folks to become sensitive.

Now the tradeoffs get made. It may be that the development tools the company has selected don’t give the developer the choice of whether or not to use persistent connections, whether or not to cache locally, a whole series of things. Or you may not be able to make changes in the appropriate time frame because you’re stepping into this one particular development midstream. Whenever you turn the process on, something is already going to be in development. Or the third-party service you’re working with has all the capabilities you need, except it’s really poorly structured.

So the right answer is, we will put in app delivery controls. We will deploy WAN optimization controls. We will use specialized XML appliances to do security and transformation. That’s OK. That’s a perfectly acceptable answer. It’s just you want to know up front so you can budget appropriately for equipment [and] staff. You can have the appropriate equipment ready so that you can have it tested and qualified before you go live. The right answer is to embed intelligence in the network. It’s just doing it up front. The goal is for the application to succeed on day one.

Q: This allows for better budgeting up front, and potentially cost savings?

A: Absolutely. Quite often, by the time it gets to the network, the app guys have already added memory to the servers, they’ve put more processors in the servers, they’ve maybe even bought some additional application licenses. We’ve had cases where people are a quarter million dollars into it, a half million dollars into it, before they ever call the network guys. And frankly, the money is wasted. They could have spent less up front by letting the network intelligence offload processes that have no business running on a Web server, or on an application server. The other advantages are you reduce the number of servers in the data center. That means you reduce power at the thermal footprint. You reduce rack space. You reduce complexity in terms of management configuration. You simplify the cabling. If you can reduce power and thermal in a data center, that’s a big win.

And so by embedding intelligence in specialized networking components, you can see a 10-to-1 to 100-to-1 improvement in performance. And that may allow you to reduce the number of servers by maybe a factor of 50%. If you can take 10 or 20 servers out of a data center, if you've got a lot of applications, it adds up across the applications. It might be 30 or 40 or 50 servers. That winds up being a meaningful amount of power and verbal footprint.

Q: By embedding in this intelligence, can the network provide a strategic advantage?

A: If you take the approach that we've spoken about, you can now begin to use your network as a strategic advantage. I've had some folks tell me the instruction from management has been, "Make sure we're no worse than any of our competitors. We should be doing the same thing [as] everybody else at [a] similar size company." You shouldn't. This is an opportunity to use the network as a significant advantage to lower your costs, reduce development cycle time, improve the effectiveness of applications, get to market sooner with better customer service. By embedding intelligence in the network, by building an application-fluent network, you can actually make your company more competitive.

Q: As you point out, it is an iterative process, so you would simply get better at it, the more you do it?

A: You get better at it the more you do it, it becomes engrained in the culture of the company, and that's why it's really important for someone at the level of the CIO to be the one who says, "This is how we now do business."

The earlier the network and applications and end-user population and the rest of that team get together, the smoother the process goes. And the good news is the application vendors—SAP, Microsoft, Oracle, go down the list—are having a change of attitude. They've come to realize that it's cheaper in the long run to tell people, "By the way, you've budgeted \$30 million for this deployment; you're going to need some networking gear. It might cost you another million, a million and a half, but the app will work." The app vendors realize that they're better off if the application works well right out of the gate. What people don't like is that unhappy surprise six months later.

Q: What can the network folks do to be proactive?

A: Even if your company doesn't have the right process, go out on your own and start meeting on a regular basis with the applications stakeholders. Find out what their key business initiatives are over the course of the next 18 months. What new business processes, what new applications do they expect to deploy? Figure out who your cohorts are in the other organizations—on the server team, on the database team, on the application team. Begin to build those personal and professional relationships. Have lunch with them occasionally to find out what they're working on. Begin to get this process going bottom-up.

Q: So in terms of the relationship between networking and applications, good communication is just as important as technology?

A: There are only three reasons to install a network. The first is to run business applications, the second is to run business applications, and the third, you can probably guess. So if the network folks get the right mindset, they begin to realize that interaction with the business stakeholders and the app developers and the rest of that group is an important part of their business.

They've [networking professionals] been able to live in isolation for a long time. The downside is the world has changed and they can no longer do that. The upside is, they have the opportunity now to become much more involved in the business, and that bodes well for their career opportunities long term. If you're seen as someone who understands the business and is supporting the critical business processes, you're more visible, you've got a greater chance for personal reward. And your budget is less likely to be the first one whacked because no one knows what value you bring. You become seen as something more than just a plumber.