



Transformation Leader | Vol. 2

TSO Logic



An Interview with William Giard Principal Engineer Intel Data Center Group

How to use analytics to de-risk and fast track enterprise IT transformation.

Given the marketing power and innovative services of major public cloud service providers (CSPs), it is no surprise that C-level executives are increasingly calling on IT to ask about the benefits of public cloud migration. While there are real advantages to public cloud consumption — lower capital costs, improved resource utilization and availability, and faster time-to-market – it’s also important to recognize that enterprises can also benefit from using cloud technologies via private clouds (on-premises or as a managed service) or hybrid environments that mix both private and public.

To determine whether cloud is worth the investment, organizations need to clearly understand how their data center environments work today, how they might work in the future, and the steps that will be involved—technical, operational, cultural—to reach their goal state.

To help work through it, let’s review some unique algorithmic analysis from TSO Logic on optimizing future-state options for on-premises and cloud compute. Then, we’ll hear from Intel leadership about how they are working with enterprises to plan and mitigate barriers to transformation.

New Server Research from TSO Logic Reveals Economics of Cloud Migration and Server Consolidation


TSO Logic performed a statistical analysis of operating system (OS) instances across an anonymized sampling of its North American customer base, representing multiple industries. The analysis encompassed hundreds of millions of data points from over 10,000 physical servers (including both hypervisors and non-hypervisors) and 25,000 virtual OS instances.



26% of current OS instances are over-provisioned based on historical utilization.



25% of servers were more than three years old.



54% fewer servers needed if older servers were updated and resources right-sized.

Based on historical utilization patterns, from a pool of 25,000 OS instances currently running on-premise:

- 26 percent of current OS instances are over-provisioned based on historical utilization.
- Right-sizing those instances—migrating them to cloud instances sized for actual utilization—would generate a 36 percent cost reduction compared to provisioning them in the cloud exactly as they’re provisioned on-premises.

Even in traditional data center environments, the research revealed ample opportunity for optimization. According to an anonymized statistical analysis of 10,000 physical servers⁽³⁾ across the TSO Logic North American customer base:

- 25 percent of servers were more than three years old.
- Accounting solely for the processing power gains of newer hardware, the same workloads running today on Generation-5 servers—provisioned “like-for-like,” using the same resources per OS instance provisioned right now—could run on 30 percent fewer Generation-9 servers.
- By updating older servers to modern hardware and right-sizing resources for each workload based on historical utilization, organizations could support the same workloads with 54 percent fewer servers.

Planning and Mitigating Barriers to Transformation: An Interview with William Giard, Intel Leadership

Q1. How closely does the TSO Logic research map to what you see when working with customers?

A: We've seen two things that align to those research findings. First, if you're running straight bare-metal instances, or even virtual machines (VMs), moving to an automated orchestration environment certainly yields value in and of itself from operational efficiency perspective. Certain workloads can take advantage of public cloud, and other workloads can take advantage of private cloud. We see that just moving your traditional application server deployments to a cloud environment—whether public or private—can yield significant results in terms of efficiency, agility, time-to-market, and cost.

Second, in terms of OS instances, we also find that there are significant efficiencies that can happen by upgrading and consolidating hardware, and addressing things like over-provisioning of memory and resources. A lot of customers can benefit from just the native advantages of upgrading and consolidating the server environment.

At Intel, we have a pretty significant run rate of improvements within our own internal IT environment. We have a large server environment and have moved over 90 percent of our office and enterprise environments from traditional infrastructures to virtualized OSs. Even without being cloud native, we've seen some pretty significant improvements. (Click [here](#) to read more about Intel's Data Center Strategy)

“ In terms of OS instances, we also find that there are significant efficiencies that can happen by upgrading and consolidating hardware, and addressing things like over-provisioning of memory and resources.”

Q2. What do you see as the biggest misconception among organizations contemplating cloud migrations?

A: Probably the notion that if you want to take advantage of cloud, you have to commit to one approach for your entire environment. There is definitely a lot of marketing out there to drive a specific strategy, and a lot of push to go all in one direction. If you read some of the news stories or talk to some of the big cloud providers, a lot of the message is around how to go public cloud. That makes a lot of sense in many instances – but not all. And then others want to keep everything on-premises, but that's not the right answer either.

The biggest thing we educate our customers around is that cloud is an architecture that can be deployed in both locations. It's not all public and it's not all private. It's absolutely a mix—that's truly where the world is going.

Q3. When you're working with customers, who do you see taking the lead role in driving transformation? What are their motivators?

A: Normally, the CIO and executive staff drive a lot of these projects and it often starts with a conversation about cost reduction. And security, of course is always a top consideration that we hear. But where we see more pull from the lines of business is really in the agility space. We see more and more of the business coming in and saying, "Hey, we need to deploy new solutions or bring in new capabilities faster." The mix of who

is driving transformation is starting to change, with the business taking a much larger role. In many cases now, the application line of business is becoming a primary decision-maker in a lot of cloud decisions, where it used to be traditionally just the infrastructure engineer.

The good news is that cloud, and specifically the public cloud efforts that you see from the big CSPs, are really changing expectations for the users, and driving more agility and self-service into mix. That's causing disruption - in a good way.

Q4. How important is it to have real data about historical utilization and the cost of various future-state options?

A: This is an area where having more in-depth insights, such as what TSO Logic is doing, can give you a jumpstart to make smarter cloud-based decisions. Just having the visibility to understand what's running on the server, what it's interacting with and how it's performing is extremely valuable. And that holds true whether you're landing and upgrading a system internally, or landing and launching it to a public cloud environment. As we've seen, both can be viable. But having the data is critical to making informed cloud decisions and helping businesses get out of what we at Intel call the "cloud stall" – when confusion and uncertainty about the right strategy stalls all progress. The power real data has in getting enterprises moving in the right direction again is simply unmatched.

“ Having the data is critical to making informed cloud decisions and helping businesses get out of what we at Intel call the “cloud stall” – when confusion and uncertainty about the right strategy stalls all progress. The power real data has in getting enterprises moving in the right direction again is simply unmatched.”

Q5. What kinds of risks do you take if you don't have accurate data?

A: We have seen many cases where organizations begin to pursue cloud, but in lacking the right visibility into what was already running on their current system, mistakenly over- or under-provision applications. Sometimes complete app migrations have to be re-worked because the migration was much more expensive than anticipated. Changing strategy mid-way like this ultimately causes delayed deployment and value realization, which has obvious business and financial consequences.

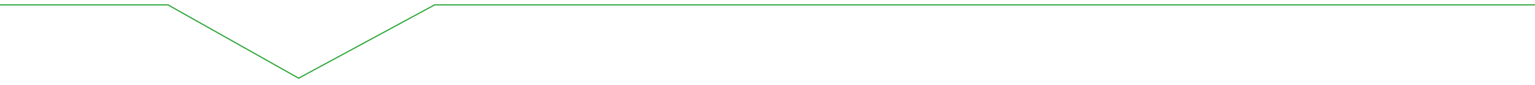
At Intel, we want our customers to be able to focus their dollars on growing their business. Helping them work through those complicated decisions upfront so that they don't have to move things back and forth, or experience business impact, is where we get very interested. That's where really precise insights into how the solution is running and what it needs become important.

Q6. When does it make sense to keep applications on-premises versus moving them to a public cloud?

A: Deciding where to land a given workload really is something you have to evaluate on a case-by-case basis; there is no “one size fits all” answer. And there are many cases where keeping applications on-premises makes sound business sense. At Intel, we have a mix. Some workloads we host externally, and others we keep internal because we can get better cost advantages.

In general, if it's an application that doesn't require global scale, we often find we can manage it more effectively on our internal cloud infrastructure. In other cases, we do go externally, especially in our SaaS solutions. So, it's a balance. We have both a public and a private cloud strategy, and we apply them to different things.

The high-level guidance I'd give organizations is to look at where your business is focused and where



your users are located. For example, if something has high interaction and your users are local, then keeping it internal makes sense. In other areas, moving externally and taking advantage of public cloud capabilities can help the business focus more narrowly on what's most important to them.

From another perspective, if you're a smaller business and/or don't already have a major data center infrastructure investment, the public cloud makes a lot of sense for most of your capabilities. If you're a large business, you may have a larger mix on-premises. But I think it's definitely a mix between the two.

For more information on our point of view, please see our Optimal Workload Placement for Public, Hybrid, and Private Cloud white paper on www.intel.com.

Q7. In your experience, what are some of the biggest reasons that transformation projects stall?

A: I believe the biggest one is underestimating the level of integration that you have with other systems and infrastructure. So you move something, and then you realize, hey, I've got these other system integrations, or these other user interactions, that rely on that data. From a technical perspective, the number one barrier is really integration and user interactions.

Now, the other big way that organizations stall is really in the culture—the engagement between the business and application teams, and the IT organization. We see a lot of transformational projects hit a wall because they don't engage both the consumers and the providers together. You have to have that cultural discussion and ask: "Are we bringing the capabilities to enable both teams?" That cultural engagement, making sure that you're setting a strategy that serves both, is really important.

Q8. When do data and analytics become most important in a transformation project?

A: There are many phases to a transformation journey, and having a good understanding of your environment is pretty critical. Whether you realize that at the beginning or you wait until halfway through the project to go get the detail, you do eventually have to get there. It's always better to have the data earlier and not have to revisit assumptions.

If you look at some of the insights you get from TSO Logic into data size and memory consumption, for example, you gain a more granular level of detail to understand what something should be running on, what it's consuming. So it's important. At Intel, we did a lot of the work to get that information manually, and we added probably several months to our project assessments on some of our strategy. If we'd had automated capabilities to get those answers, it probably would have been a quicker journey for us.

Q9. How important is it to have that data on an ongoing basis, even after you move something to the cloud?

A: After you've deployed, data is still important. You want to be able to manage capacity and costs. You want to be able to understand availability thresholds. Whether you need it on a daily flow or via passive collection becomes part of standard operations. We have seen a lot of cases, especially as you start running older hardware, where the data provides better insight into how to improve the business. If you're considering specific hardware-level refreshes to be able to take advantage of the latest memory throughput and capacity, for example, then memory analytics becomes important. Over time, as you're trying to justify upgrades or figure out the best way to consolidate or refresh hardware, having that data does become

an effective way to continue to manage and scale your business.

Q10. As customers weigh their transformation options, they're getting lots of messages from vendors with a stake in a specific solution. How important is it to have impartial advice on how to proceed?

A: It's not hard for our customers to get advice, but are you getting impartial advice? It depends on who you talk to. I think it's important to have a balanced view. The good news about Intel is that we're in most of the solutions in both public and private cloud. We really, truly, just want to find that right mix for our customers that allows them to make the best decisions for their business.

The two most important things we've found in talking to different organizations is, first, making sure you have that balanced view, and second, using data to help shape the decision. Data becomes the normalizing factor. So as you have that discussion, you can answer, "What's the data behind that recommendation?" That's where data-driven recommendations become really important.

Conclusion: Making Data-Driven Decisions

Any organization considering a cloud migration, or even a hardware refresh, faces a long list of questions:

- Which deployment model makes the most sense?
- What impact will it have on my applications and users?
- How much will I actually save by doing this?

Getting at those answers can be an intensive process, but the one commonality across every decision point is the value of accurate data. The more you understand about your environment, the more confident you can be that you're making the right decision.

The TSO Logic platform was developed specifically for this reason—to eliminate the frustration of wanting to drive innovation, but not having the clarity and visibility to know you're making the right decisions. Its backbone is a set of proprietary algorithms that quickly creates a fine-grained statistical model of compute resources and historical utilization patterns to determine the most cost-effective place to run each workload.

Successful transformation initiatives encompass a broad range of factors and a deep understanding of security, integration, cultural, and other considerations. No organization can answer every question overnight. But with total visibility into compute resources and application delivery, you can eliminate much of the heavy lifting involved in analyzing the best deployment option for a given workload. You can drive out the guesswork that's holding back many organizations from capitalizing on transformation. And you can make smarter decisions about your cloud strategy and data center investments.

Learn More

IT transformation requires visibility and the tools to help you navigate a path to your ideal future states. Together, TSO Logic and Intel are the easy first step to bring that clarity to your fingertips.

To find out more, please contact:

TSO Logic

T: 1.866.379.9688

E: sales@tsologic.com

W: tsologic.com

About William Giard

William Giard is Principal Engineer for Intel's Data Center Group (DCG), working with Intel customers to help them map out the best strategy to grow their business. William has over 20 years of experience designing enterprise architectures and developing software solutions to support mission-critical systems across supply-chain, product development, and enterprise infrastructure segments. Prior to joining DCG, he led the effort within Intel IT to modernize the application and computing environment, delivering secure and usable solutions across multiple client computing platforms while hosted in the cloud to enable new business models.

Footnotes

1. Calculations for on-premises costs include server amortization, licensing, maintenance and power. Labor is not included in this analysis. Virtual instances reported on come out to an average guest: host ratio of 10:1.
2. Cloud catalog comparison and pricing data is based on Amazon Web Service (AWS) and Microsoft Azure. AWS T-Series instances represent a large portion of the cloud savings.
3. On-premises OS instance cost is based on the amount of resources provisioned to a virtual machine, licensing, maintenance, hardware amortization and power. Physical servers reported contain both hypervisors and logical servers.