

How to Do Capacity Planning



IT is playing an increasingly important role in business, with IT innovation expected to drive business transformation. In fact, 85 percent of IT leaders surveyed by Foundry in 2023 said they see the CIO role becoming more digital- and innovation-focused.

Given the growing dependence on IT to drive digital transformation and business success, ensuring consistent service delivery is extremely important. In today's business world, you can't afford a single bottleneck or a single minute of downtime.

Capacity planning plays a critical role in making sure everything stays up and running at maximum efficiency. Whether you open another office, launch a new product or service, or simply want to be sure that your infrastructure will be reliable and cost-effective, capacity planning will help you keep service levels up and costs down.

A Changing Landscape

While capacity planning has long been important for minimizing the risk of downtime and maximizing uptime, reliability, and efficiency, the best ways to optimize service delivery have changed over time.

Not long ago, capacity planning was strictly focused on infrastructure in the on-premises data center. The top priority was keeping an eye on servers, storage, and network bandwidth to ensure resources were properly utilized.

Adoption of virtualization eliminated the fixed tie between an application and the infrastructure. The capacity planning process now had to take things like workload placement and resource pools into consideration.



More recently, organizations are moving workloads into the cloud. The elasticity of IaaS means you can worry less about capital expenditures and provisioning lead times, but you still need to plan ahead and optimize your allocation strategy. Capacity in the cloud may seem infinite, but the cost for allocations needs to fit your budget.

Regardless of technology, you need to know about future business demands and how they will impact capacity requirements over time. Only then can you optimize your sourcing strategy and cost.

It's also important to remember that paradigm shifts are rarely binary nor do they happen overnight. You will most likely have services running on older technologies that still need to be managed. Your capacity planning process has to cover all technologies used by your organization.

How to Do Capacity Planning

Start here and discover what you need to know in order to do capacity planning today. If you want to do capacity planning right, you need to follow these six steps:

Steps in Capacity Planning

1. Identify a Capacity Planning Champion
2. Get Buy-in from the Business
3. Determine Your Requirements
4. Analyze Your Resources
5. Plan for Future Needs
6. Influence Your Audience

1. Identify a Capacity Planning Champion

While having a dedicated capacity planner isn't always possible, you do need to put someone in charge to outline the strategy and drive the capacity planning initiative. Capacity planning requires involvement and assistance by other parts of the organization, not just IT. So, one key responsibility of the capacity planner or champion is to define and coordinate the work of the cross-functional capacity planning team.

2. Get Buy-In from the Business

Getting the business involved early is critical to your success. To get buy-in, you need to highlight the benefits of capacity management:

- Better planning will improve service quality
- Improved efficiency will make services more affordable
- Better business and IT alignment will improve agility

To accomplish these objectives, you need to make it clear that the business needs to provide you information to base your planning on. There are many business decisions that will impact the demand for a service, like:

- Marketing campaigns
- Acquisitions
- Expanding service offerings

By being involved in planning discussions and getting this type of information early on, you can do a proper job planning.

Additionally, when the business understands the value of capacity planning you will be empowered to make necessary changes. Capacity planning is a strategic discipline that touches many parts of IT service delivery. You rely on contributions from other teams to be successful. If you're not getting the information you need, having the backing of your customer, the business, makes behavior and process discussions easier.

3. Determine Your Requirements

Once you have buy-in from the business, it's time to establish formal performance and availability requirements.

The requirements should be dictated by service level agreements (SLAs). Meeting your SLAs is critical to your business success and your reputation. Examples of common SLAs are:

- Uptime requirements for a business service
- Response time for an online service
- Hourly throughput rates for back-office component

Be sure to evaluate your SLAs with a critical eye. Too many organizations have SLAs focused on incident resolution and service restoration. Or, their SLAs are simply too generic. To determine the quality of service and promote accountability, SLAs need to include specific availability and performance metrics. You will also have to map high-level SLAs to the various components involved, determining each individual requirement.

It's also important to understand the business cycles and how they correlate with IT demand, for example, seasonality in retail businesses. Any type of capacity analysis or requirements-gathering needs to take cycles like these into consideration.

4. Analyze Your Resources

Once you've established your performance requirements, it's time to analyze your infrastructure and current capacity situation.

First, you need to make sure you have access to performance data for your relevant infrastructure. This probably means combining multiple sources. Without

data, you can't analyze. If data is missing for part of your infrastructure, you need to find a way to collect it.

Once you have access to the relevant data, you can start the analysis to look for things like:

- Seasonality in the utilization of the infrastructure
- Correlation of utilization across different components
- Highly utilized or saturated systems in need of additional resources
- Components with low utilization that can be consolidated or reclaimed
- Upward trends in utilization and growth rate

This analysis will help you establish a baseline of your current situation. And it will probably help you find some looming problems before they happen.

5. Plan for Future Needs

You know your performance requirements and you've analyzed your current state. Now it's time to think about your future capacity requirements.

Here's what you'll need to develop a capacity plan for the future.

Forecast Business Demands

Figuring out the right scenarios to plan for is crucial.

You need to engage with the business to capture their expectations and validate those against historical demand patterns.

Choose a Capacity Planning Method

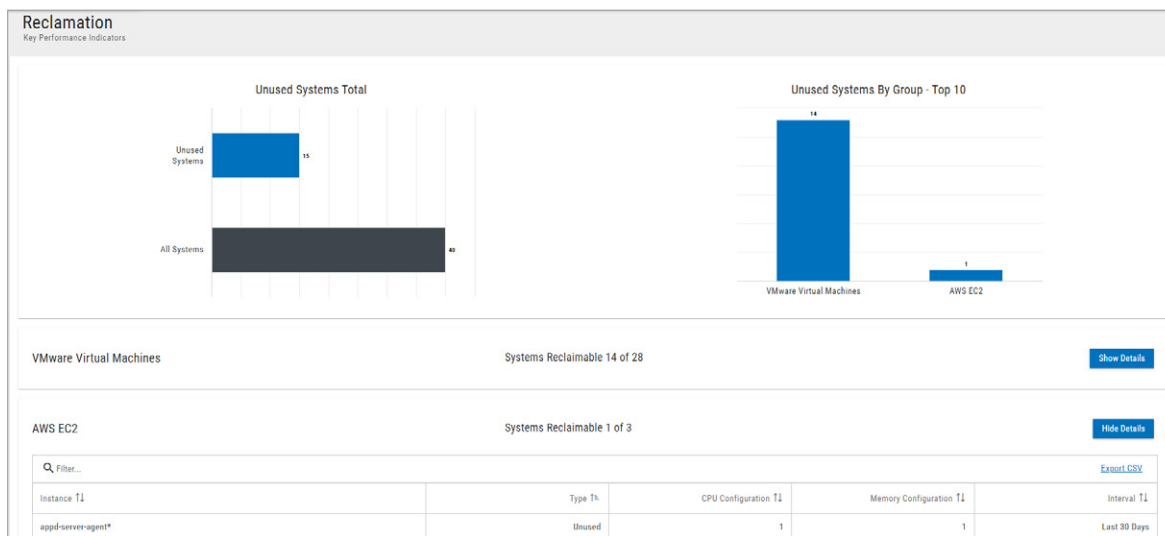
You need to decide which capacity planning method is best for your business. This is how you'll be able to determine how the forecasted demands will impact the performance and availability of your services. You'll learn more about the various capacity planning methods later on in this guide.

6. Influence Your Audience

The final step of capacity planning is just as important. You need to engage and influence your audience—both within IT and within the business at large.

IT and the business need to work together and get the right groups involved to carry out tactical activities.

The true value of capacity planning is created the moment the business follows your recommendations and an impact is made. Tailoring your message to each audience is critical to making sure that impact can happen.



Above is an automated Reclamation Report from Fortra's Vityl Capacity Management solution that shows systems that have not been used continuously for a specified amount of time. In this example, the time is the last 30 days and there are unused VMware Virtual Machines as well as 1 AWS EC2 cloud instance.

If you want your operations team to follow your recommendation about reclaiming resources, you should provide them with sufficient technical details. They'll need just the right amount of background information to be convinced.

If your audience is on the business side and you're talking about how business continuity can be improved, avoid talking about technical details. The business audience will be swayed by the value that capacity planning adds.

Capacity Planning Methods

Now that you know the steps in capacity planning, how do you actually do capacity planning?

You will want a solution that provides reliable answers to questions like:

- What will happen to the response time if the transaction rate for this application doubles?
- Can I migrate this workload from a legacy platform to our new standard platform without hurting performance?
- Do we need to be concerned about available capacity if I add a new tenant to a shared service next month?

There are three different methods for doing it, each with different characteristics.

Capacity Thresholds

A simple way to do capacity planning is to set performance or capacity thresholds. Once those thresholds are hit, actions need to be taken.

An example of this would be setting a threshold that all systems with a CPU utilization above 50 percent should be upgraded.

Using capacity thresholds is a good method to use if you only need basic capacity planning. The thresholds will help you mitigate issues as they arise.

But you can't take proactive measures, like provisioning new resources, based on capacity thresholds alone. That's because provisioning new resources usually involves a lead time. And you would need to account for potential continued growth during that lead time when defining your thresholds.

Plus, using capacity thresholds is often prone to errors. And capacity thresholds often lack the precision of other capacity planning methods. Setting accurate thresholds requires insight about individual workloads. If you're uncertain, you'll add safety margins—which leads to inefficiency and a slack capacity.

Linear Trending

Linear trending involves looking at historical data over time and creating a trend line to predict future needs. It's a commonly used method that is simple to implement and use.

Linear trending is a useful capacity planning method for workloads that increase at a steady rate. Your historical data will help you create a trend line and make sure you have enough capacity down the line.

The major limitation with linear trending is it assumes that workloads increase at a steady rate, something that is not always the case. Many times you may want to consolidate workloads or plan for events or new conditions that are not represented in the trail of historical data. In such cases, trending provides very little guidance and the whole process is reduced to guesswork.

Linear trending also falsely assumes that system performance is linear. The reality is that once you hit a bottleneck, the performance of the application drops exponentially. A capacity planning tool needs to assess more than just past performance to make accurate predictions about the future. Having the ability to determine when bottlenecks will occur is crucial.

Capacity Modeling

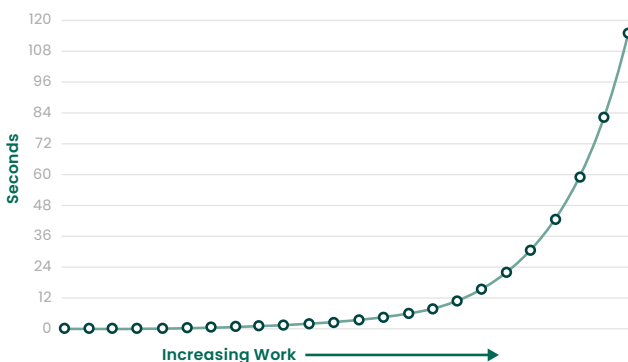
Capacity modeling helps you understand system behavior and predict if the current resources will be sufficient in different scenarios.

This capacity planning method requires analytical tools and data from monitoring solutions. Algorithms like queuing theory are typically built into capacity modeling tools. These formulas help you calculate processing times and delays, so you can predict the behavior of a system under varying loads.

The best part about capacity modeling is that you can do it without physically testing it on your infrastructure.

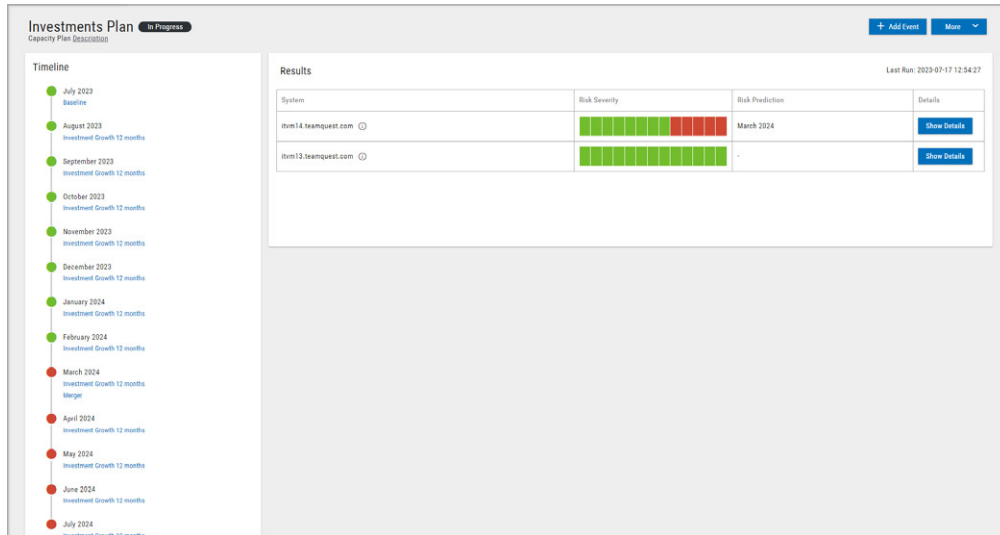
Here are some questions you can use capacity modeling to answer:

Response Time Does Not Scale Linearly



- How many VMs running this workload can safely run on each physical server?
- Which of my applications are in danger of failing to meet service levels within the next six months?
- Where will my future bottlenecks be?

A key differentiator of capacity modeling is that it can identify risk and prescribe solutions to the problem. Capacity modeling also has some distinct advantages over linear trending and performance thresholds.



Above is a 12-month capacity plan created by Vityl Capacity Management capacity plans which uses capacity modeling technology. The timeline on the left shows the plan month by month with growth events in blue. These events include a month-by-month growth, and in March 2024 a significant growth event was caused by a merger with another company. The right-hand side shows the two servers involved and when problems (if any) would be expected to start based on growth events. Clicking on the blue 'show details' button would reveal a drilldown report showing what the problems are (CPU, I/O etc.) and advice on how to fix them. These fixes could be added to the timeline and the plan re-run.

Capacity Modeling vs. Linear Trending

It gives you the ability to define your own what-if scenarios. You're not bound to what is reflected in the historical data.

Capacity Modeling vs. Capacity Thresholds

Capacity modeling is also workload-agnostic. It automatically adapts to changing circumstances. That means you don't have to define custom thresholds based on the specific characteristics of each application.

Capacity Planning Tools

Understanding the nuances of capacity planning and choosing the right method for you can be difficult. But getting started with capacity planning can be a lot easier if you enlist the right toolset.

It is important for the capacity planner to use technology and tools that simplify and automate data access, analysis, and reporting information. Without proper technology, you risk being caught up in tactical work, losing track of the end goal: making sure you have the right infrastructure resources to meet business needs without overspending.

Ready to learn How to Choose a Capacity Planning Tool?

Ready to Begin Capacity Planning?

Our solution makes it easy for organizations to identify wasted spend across hybrid IT, optimize performance, and make adaptable plans. Try Vityl Capacity Management today for free.

START TRIAL



Fortra.com

About Fortra

Fortra is a cybersecurity company like no other. We're creating a simpler, stronger future for our customers. Our trusted experts and portfolio of integrated, scalable solutions bring balance and control to organizations around the world. We're the positive changemakers and your relentless ally to provide peace of mind through every step of your cybersecurity journey. Learn more at fortra.com.