

**POLAR
STAR**
Consulting, LLC



Data Driven Resource Management

Maturity levels, capacity management, and application performance monitoring: raising the bar using data, and check lists

We all value maturity because it means a higher degree of ability and effectiveness. To develop an organization's management maturity, one can begin by following a common path, and use that path to guide the maturing process. As you grow in effectiveness, you can gauge the current organization's maturity, determine the right level to grow to, and use this path to structure project plans to develop organizational capabilities.

A good place to start is with a check list. This document provides a check list for documenting where you are, and where you need to be, plus how to find where to spend resources on improvement. This check list and the related questions are a simple place to start, but the real work happens next.

Authors: Jason Rupe, Ph.D., Director Operational Modeling; Steve Goering,
Director and Senior Network Engineer

Introduction

Data driven management of resources often follow common stages of maturity, which we have identified as five levels. The resources managed will have a wide range of attributes and behaviors, but as long as supply and demand vary over time, we need information with which to make decisions, and an organization can be in one of these five different levels of maturity.

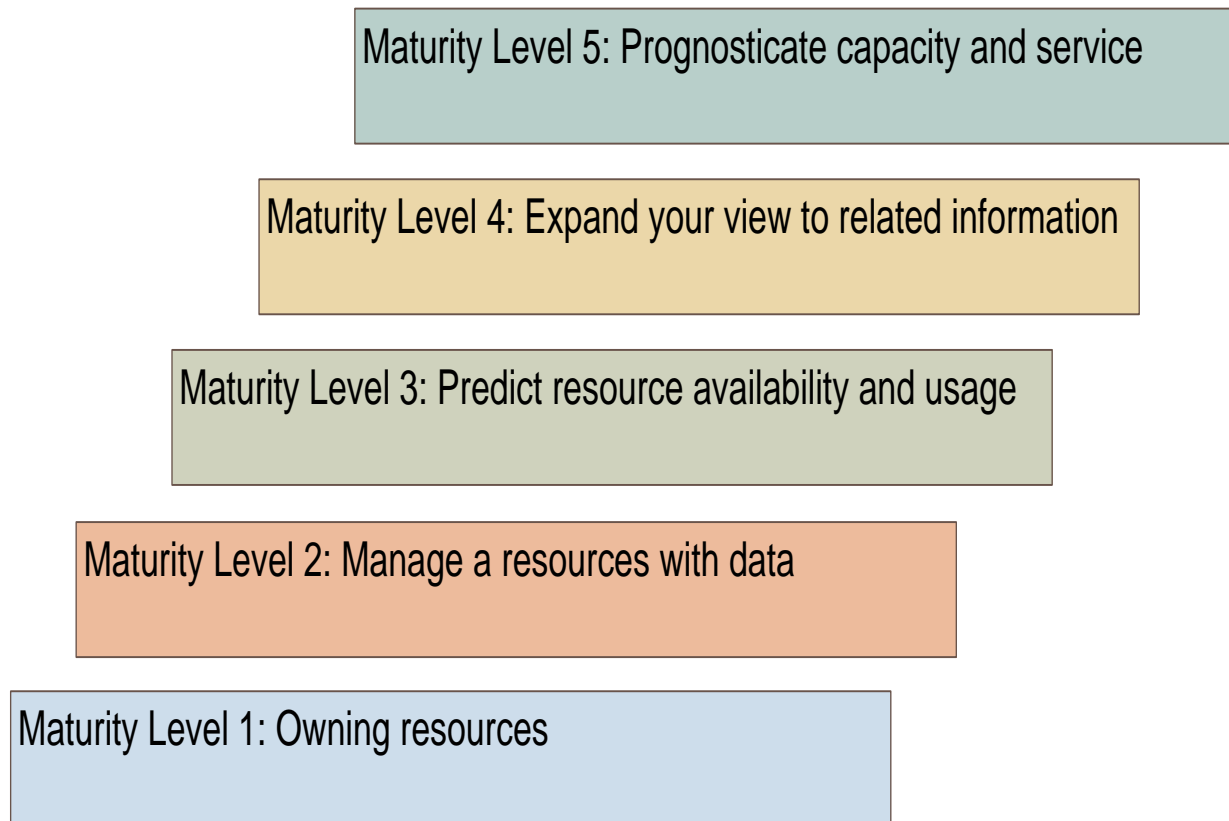


Figure 1 - Our recognized five levels of data driven management maturity

Levels of Maturity

Shown in Figure 1, we define the Maturity Levels as follows:

Maturity Level 1. Own a resource – the first maturity level is simply having a resource to manage, without perhaps any information aside from the fact that there is some amount of the resource, and that the resource is under some amount of demand that changes over time. If the resource is physical, the amount of it is usually well known as are its attributes and behaviors. But the amount of it being used may vary quickly, or be allocated so that it is essentially in use for longer or indefinite periods. The nature of its use is important because changing usage is what must be managed about the resource.

Maturity Level 2. Manage a resource with data – the second maturity level is to use data about the capacity and demand to make decisions about its adequacy. At this level, you are managing capacity to some extent by paying attention to the nature of demand (quantity, user behavior) for your resource, but you are still reacting to conditions.

Maturity Level 3. Predict resource availability and usage – the third maturity level is to begin to predict future conditions so you can manage to those expected conditions rather than simply react to shortfalls. Consequently, you are managing your risk of not having enough capacity.

Maturity Level 4. Expand your view to related information – the fourth maturity level is to use additional information for predicting and managing your capacity. Once you are predicting demand and availability of your resource, you want to improve on its success, so you investigate how the resources are being used, and what services are supported with the resource. Now you are looking at secondary information about the capacity you manage to better predict and manage that capacity, perhaps getting very close to managing a service.

Maturity Level 5. Prognosticate capacity and service delivery – the fifth maturity level is to use that new additional information for managing services with the capacity, thus using prognostics to decide when to manage the capacity of your resource as well. Now that you have secondary information, you can see how services can be directly supported, and capacity efficiently managed to enable those services, by prognostication of the service performance and availability.

Applications

This structure applies to a surprising number of cases. If there is something to manage, chances are we can consider it to be a resource with finite capacity, divided among tasks or users or functions in some way. And if there is uncertainty about how much of that resource is available for use, then there is a chance of not having enough of that resource. Further, if bad things happen when there is not enough of that resource, then we should pursue some level of maturity.

In the IT world, Capacity Management (CM) and Application Performance Management (APM) are two related functions which many IT organizations maintain for their own effectiveness as well as that of their customers. A small organization might be just fine with a part time IT capability, and no need for maintaining a database of any kind about the capacity of those IT resources. As a business begins to rely on those IT resources, likely because of some growth in the business, paying attention to the availability of those resources becomes important. Network management systems supply data for resource management purposes, which can be collected into databases and reported. As that IT function begins to manage its capacity through CM practices, their maturity grows. The organization begins to try to predict, formally or informally as needed, the need for more resources. If the importance continues to grow, eventually that IT organization may enter the business of managing their applications for their customers, thus becoming a service organization. The information required to manage the way those resources are used is more complicated, so has to be gathered and maintained in databases. If the lack of resources is sufficiently critical, then the service performance information can be used to manage applications directly, leading to APM practices, and perhaps even Prognostic Health Management (PHM) of the IT network and systems.

For commodity consumption, it is common for the consumer not to care much beyond whether a product is intended to function as it should, and its cost. However, as usage becomes more important, the nature of that functionality becomes more important. Consumers begin to care about reliability, and the degree of performance. When it becomes important enough, we try to use secondary indicators to avoid catastrophic failure. For example, we replace timing belts at a predetermined mileage to avoid breakage events that can make an engine fail. Realize that failure is just a special case of insufficient capacity to meet demand.

Even when the resource is a virtual one, there is capacity to manage, albeit more nebulous. Industry terms like virtual infrastructures and cloud architectures are no different. They only hide the true capacities to be managed, but highlight the more complicated layers of application and service. But as an architecture has a finite application and service capacity, we still have the task of managing capacity. But now, higher levels of maturity are required for sure.

What level of maturity is appropriate for a given situation? That depends on whether the cost of the resource, the impact of not having enough resource, the availability of information about the resource, and the nature of variability in the availability of that resource (including considering the demand for that resource). You can define all this mathematically and then find the right solution. But if you prefer to keep it simple, consider the cost tradeoffs. If a shortage event is likely and significant enough that its cost is higher than the cost of avoiding the shortage, then use that information and process to avoid the shortage. In other words, do your homework, and make an educated decision.

Start with a check list

Checklists are a very effective way to improve a manual process. When an enterprise realizes they have a resource that should be better managed, they consider CM or APM, in its general form if not specifically in the context of IT. The first step is to determine your best answers to the following questions, being sure to estimate the reliability of your answers as well.

- **What is the resource that should be managed?** Clearly identify the resource that needs to be managed, as crisply as possible. Define any boundaries to connected and related resources. If it is a pool of servers, how many, where are they located, which are included or excluded from the pool, and what resources on the servers are important enough to be managed?
- **Who owns that resource and its capacity?** For the sake of action and accountability, identify the owner or owning organization responsible for the capacity of the resources. It is key to identify the owner, and that owner must be able to control the resource's capacities. If this responsibility is somehow shared, then either divide the resource into two pools to avoid the conflict, or clearly define the nature of this shared responsibility so that there is no ambiguity for who must do what.
- **How do they know how much capacity they own?** If you can simply go and count those resources, maybe that is enough. But most resources in need of management will need some sort of inventory management capability.
- **How do they know how much capacity they need?** If change is rapid, then the information must be even more rapid. Rapid data collection usually requires automated data collection. And in any case the information must be reliable, but should otherwise be as inexpensive as can be while still meeting the requirements.
- **How do they know how much capacity they will need?** Either an analyst will need to do this work, or perhaps share the effort with an analytics engine. The reliability of the demand forecast is critical. Low reliability of the forecast leads to larger amounts of spare capacity, so we have a cost tradeoff to consider.
- **How does the organization add or subtract capacity?** Processes must be known, and effective, for changing capacity. Procurement processes play a key role usually, so supply chain management becomes important.

The first step of any resource management project is to answer those questions first, then to check how adequate the answers are. If the answers are not adequate, then it is time to consider increasing the enterprise's capability maturity.

- **Are the resource and responsibilities ill defined?** Then some process improvement work may be in order to clarify the roles and responsibilities.
- **Is the monitoring insufficient?** Then perhaps it is time to invest in tools, or improve the tools you have, or improve the use of the tools you have.
- **Is the information inadequate?** If there is information missing or incorrect, then work on filling those gaps. IT systems experts, systems engineers, or data analysts may need to complete the effort here.
- **Is the funding insufficient?** Balance budgets to be service oriented, or consider adding pressure for technology improvements to reduce costs.
- **Is the procurement process insufficient?** Increase the spare capacity in the supply chain to ease the pain, but look at supply chain optimization to rebalance the solution effectively and efficiently.
- **Do the teams lack coordination?** Much like the first problem in this list, perhaps we need more clarity of roles and responsibilities, but deeper in the organization. Clarify and get agreement on responsibilities, accountabilities, controls (establish or improve your RASIC or RACI), and Operational Level Agreements (OLA) so there is no lack of clarity.

There is a lot more to resource management than just getting a good tool to provide nice graphs. Much more. There is a responsibility that comes with ownership, and tools and techniques and agreements do not remove that responsibility. They can only help you be a more effective owner.