

Realtime  
publishers

*The Definitive Guide™ To*

# Application Performance Management

*sponsored by*



*Greg Shields*

Chapter 10: Your APM Cheat Sheet .....	181
Part 1—What Is APM? .....	181
Defining APM, More than “On vs. Off” .....	183
Part 2—How APM Aligns with the Business .....	185
What Changes with APM? .....	186
Part 3—Understanding APM Monitoring.....	187
Part 4—Integrating APM into your Infrastructure .....	188
Part 5—Understanding the End User’s Perspective .....	191
Monitoring from the End User’s Perspective .....	192
Where Does EUE Fit?.....	193
Part 6—APM’s Service-Centric Monitoring Approach .....	195
Flow Up, Drill Down.....	198
Part 7—Developing & Building APM Visualizations .....	199
Part 8—Seeing APM in Action.....	200
Visibility.....	200
Prioritization .....	201
Problem & Fault Domain Isolation.....	201
Troubleshooting, Root Cause Identification, & Resolution .....	201
Communication with the Business.....	201
Improvement.....	201
Part 9—APM Enables Business Service Management.....	202
Linking BSM to APM .....	203
APM Is Required Monitoring for Business Services .....	204

## Copyright Statement

© 2010 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers or its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at [info@realtimepublishers.com](mailto:info@realtimepublishers.com).

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library for IT Professionals. All leading technology eBooks and guides from Realtime Publishers can be found at <http://nexus.realtimepublishers.com>.]

## Chapter 10: Your APM Cheat Sheet

---

APM is obviously a big and complex topic. With monitoring integrations that span networks, servers, and end users along with all the transactions in between, a fully-realized APM solution quickly finds itself wrapped around every part of your network. Truly understanding how an APM solution adds value to your business services requires every part of a multi-chapter book. This book's definitive approach to explaining APM strategy and tactics gives you the information you need to make smart purchasing decisions. Once decided, its chapters then help guide you towards the best ways to lay it into place.


However, not everyone has the time or the interest in poring through a 200-page tome. Digesting this guide's 200-odd pages will consume more than an afternoon making the topic hard to approach for the busy executive or IT director. To remedy this situation, this final chapter is published as a sort of "Cheat Sheet" for the other nine. Using excerpts from each of the previous chapters, this "shortcut" guide summarizes the information you need to know into a more bite-size format.

So, how is this chapter best utilized? Hand it out to your business leaders as a walkthrough for APM's business value. Pass it around your IT department to give them an idea of APM's technical underpinnings. Show it to your Service Desk employees as an example of the future you want to implement. Then, for those who show particular interest, clue them in on the other chapters for the full story. In the end, you'll find that APM benefits everyone. You just have to show them how.

### Part 1—What Is APM?

If you're an IT professional reading this guide, you've heard these stories many times before. You know about the host of potential problems that an IT infrastructure can and does experience on any particular day. You've experienced the nightmare situation where a critical service goes down and no one can track down exactly why. You've sat in the "war room" where highly-skilled individuals from every IT domain—network engineers, systems administrators, application analysts—sit around the conference table for hours attempting to prove that the problem isn't theirs. Whether you're an IT professional, or someone who directs teams of them, you know that any downed service immediately signals the beginning of a bad day.

The problem is that the idea of a “service that is down” is often so much more than a simple binary answer; on versus off, working versus not working. As you can see in Figure 10.1, IT services are made up of many components that must work in concert. Servers require the network for communication. Web servers get their information from application servers and databases. Data and workflow integrations from legacy systems such as mainframes must occur. These days, even data storage must be accessible over that same network.



**Figure 10.1: An IT service is comprised of numerous components that rely on each other for successful operation.**

If any of those pieces experiences an unacceptable condition—an outage, a reduction in performance, an inappropriate or untimely response, and so on—the functionality of the entire service is affected. This can happen in any number of ways:

- The service or hardware hosting the service is non-functional
- A server or service that is relied on is non-functional
- One or more servers or services that make up the service are not performing at an acceptable level
- An individual component or function of the service is non-functional or is not performing at an acceptable level

All of these are situations that can and will impact the ability of your critical IT services to complete their stated mission. No matter whether the actual service itself is down or the cause is some component that feeds into the functionality of that service, the ultimate result to the end customer is a degradation in service. The ultimate result to your business is a loss of revenue, a loss of productivity, and the inability to fulfill the regular needs of business.

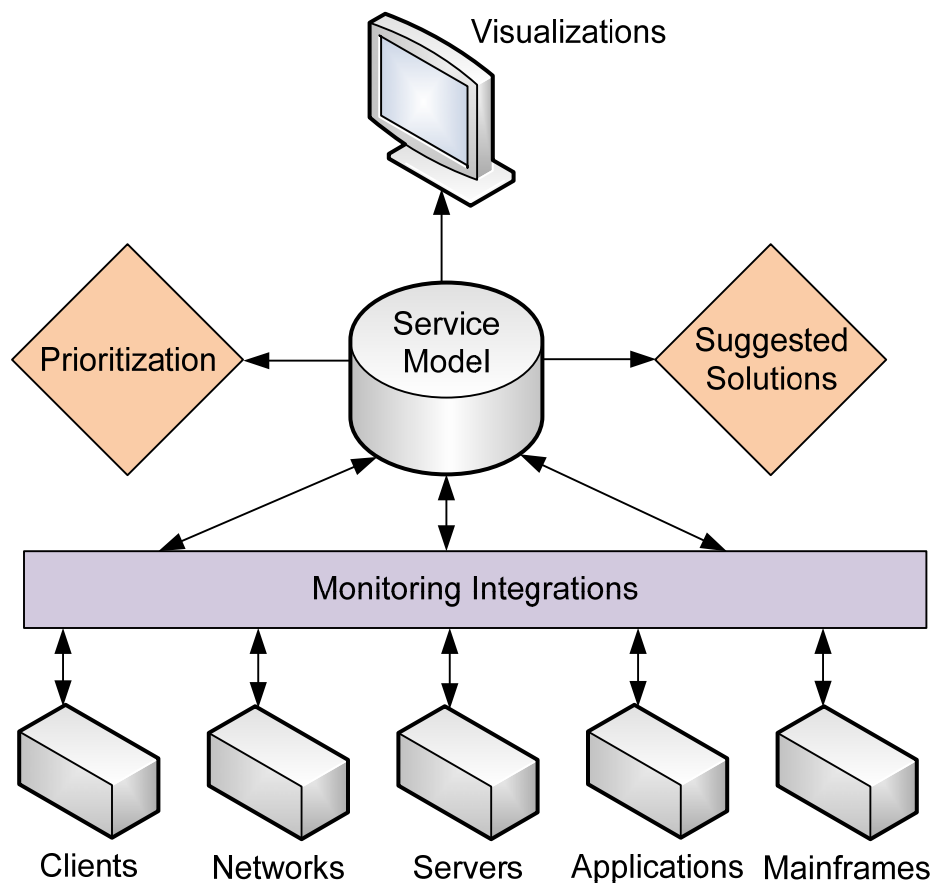
### Defining APM, More than “On vs. Off”

Fixing IT’s former “on versus off” approach to service management is therefore a critical step. As such, smart organizations are looking to accomplish this through a more comprehensive approach to defining their services, the quality of those services, and their ability to meet the needs of users. Application Performance Management (APM) is one systems management discipline that attempts to provide that perspective. Consider the following definition:

*APM is an IT service discipline that encompasses the identification, prioritization, and resolution of performance and availability problems that affect business applications.*

Organizations that want to take advantage of APM must lay in place a workflow and technology infrastructure (see Figure 10.2) that enables the monitoring of hardware, software, business applications, and, most importantly, the end users’ experience. These monitoring integrations must be exceptionally deep in the level of detail they elevate to the attention of an administrator. They must watch for and analyze behaviors across a wide swath of technology devices and applications, including networks, databases, servers, applications, mainframes, and even the users themselves as they interact with the system.

Figure 10.2 shows an example of how such a system might look. There, you can see how the major classes of an IT system—users, networks, servers, applications, and mainframes—are centered under the umbrella of a unified monitoring system. That system gathers data from each element into a centralized database. Also housed within that database is a logical model of the underlying system itself, which is used to power visualizations, suggest solutions to problems, and assist with the prioritization of responses.



**Figure 10.2: An APM solution leverages monitoring integrations and service model logic to drive visualizations, prioritize problems, and suggest solutions.**

With its monitoring integrations spread across the network, such a system can then assist troubleshooting administrators with finding and resolving the problem’s root cause. In situations in which multiple problems occur at once—not unheard of in IT environments—an APM system can assist in the prioritization of problems. In short, an effective APM system will drive administrators first to those problems that have the highest level of impact on users.

## Part 2—How APM Aligns with the Business

For an organization to efficiently make use of the kinds of information that an APM solution can provide, it must operate with a measure of process maturity. IT organizations that lack configuration control over their infrastructure don't have the basic capability to maintain an environment baseline. Without a baseline to your applications, the quality of the information you gather out of your monitoring solution will be poor at best and wrong at worst.

But how does an IT organization know when they've got that right level of process in place to best use such a solution? Or, alternatively, if an organization recognizes that they don't have the right level, how can an APM solution help them get there?

One way to evaluate and measure the "maturity" of IT is through a model that was developed in 2003 as part of a Gartner analysis titled *Transforming IT Operations into IT Service Management* (Source: Gartner IT Management Process Maturity Model, *Transforming IT Operations into IT Service Management, Data Center 2003, Deb Curtis and Donna Scott*). This groundbreaking white paper defined IT across a spectrum of capabilities, each relating to the way in which IT actually goes about accomplishing its assigned tasks. An IT culture with a higher level of process maturity will have the infrastructure frameworks in place to make better use of technology solutions, solve problems faster, plan better for expansions, and ultimately align better with the needs and wants of the business they serve.

Process maturity within an organization is defined as quite a bit more than simply having the ability to solve problems. Within Gartner's maturity model, the capacity of IT to solve—and prevent—ever more complex problems was defined largely by its level of process maturity.

Secondly, and arguably more importantly, *smart organizations can leverage an APM solution itself to rapidly develop process maturity in an otherwise immature organization*. By reorganizing your IT operations around a data-driven approach with comprehensive monitoring integrations, you will find that you quickly begin making IT decisions based on their impact to your business' applications. You will better plan for augmentations based on actual data rather than the contrived anticipation of need. You will better budget your available resources based on actual responses you get out of your existing systems.



## What Changes with APM?

With IT's movement from one stage to the next, the entire culture of the organization changes as well. IT at higher levels of maturity has the capacity to accomplish bigger and better projects. But IT at higher levels of maturity also thinks entirely differently about the tasks that are required:

- **The ways IT looks at itself.** In earlier stages of maturity, IT sees itself as a fully-segregated entity from the business. In many cases, IT can see itself as a different business entirely! Individuals in IT find themselves concerned with the daily processing of the servers and the network, to the exclusion of the data that passes through those systems. As IT matures, the natural culture of IT is to begin thinking of itself as a partner of the business, and ultimately as the business itself.
- **The ways IT looks at data & applications.** Data and applications in the immature IT organization are its bread and butter. These are the elements that make up the infrastructure, and are worked on as individual and atomic elements. IT in earlier stages will find itself leveraging manual activities and shunning automation out of distrust for how it interacts with system components. Applications in early-stage IT are most often those that can be purchased off the shelf, with customization often very limited or non-existent. Later-stage IT organizations needn't necessarily build their own applications; however, they do see applications as solutions for solving business processes as opposed to fitting the process around the available application.
- **The ways IT looks at the business.** Immature IT organizations are incapable of understanding how their activities impact the business as a whole. Lacking a holistic view of their systems, they focus on availability as their primary measure of success. Yet business applications require more than a ping response for them to be truly available to users. More mature IT organizations find themselves implementing tools to measure the end user's experience. When that level of experience is better understood, IT gains a greater insight into how their operations impact business operations.
- **The tools IT uses.** The tools of IT also get more mature as the culture grows in maturity. IT organizations with low levels of maturity are hesitant to incorporate holistic solutions often because they can't see themselves actually using or getting benefit from those solutions. As such, immature IT organizations lean on point solutions as stopgap resolutions for their problems. The result is that collections of tools are brought to bear while unified toolsets are ignored. Mature IT organizations have a better capability to understand the operational expense of an expanding toolset, while being more capable—both technically and culturally—of leveraging the information gained from unified solutions.

As the maturity of IT's tools grows, so does the predictive capacity of those tools. It was discussed in Chapter 1 that solution platforms such as those that fulfill APM's goals extend their monitoring integrations throughout the technology infrastructure of a business. Because APM's reach is so far into each of a business application's components, it grows more capable than point solutions for finding the real root cause behind problems or reductions in performance.

### Part 3—Understanding APM Monitoring

Part two discussed the concepts of IT organizational maturity. Although that conversation has little to do with monitoring integrations and their technological bits and bytes, it serves to illuminate how IT organizations themselves must grow as the systems they manage grow in complexity. As an example, a Chaotic or Reactive IT organization will simply not be successful when tasked to manage a highly-critical, customer-focused application. The processes, the mindset, and the technology simply aren't in place to ensure good things happen.

To that end, IT has seen a similar evolution in the approaches used for monitoring its infrastructure. IT's early efforts towards understanding its systems' "under the covers" behaviors have evolved in many ways similar to Gartner's depiction of organizational maturity. Early attempts were exceptionally coarse in the data they provided, with each new approach involving richer integrations at deeper levels within the system.

IT organizations that manage complex and customer-facing systems are under a greater level of due diligence than those who manage a simple infrastructure. As such, the tools used to watch those systems must also have a higher level of due diligence. As monitoring technologies have evolved over time, new approaches have been developed that extend the reach of monitoring, enhance data resolution, and enable rich visualizations to assist administrative and troubleshooting teams:

- Simple availability with ICMP
- Richer information with SNMP
- Device details with the agent-based approach
- Situational awareness with the agentless approach
- Application runtime analysis for deep monitoring integration
- Complete recognition of the end user's experience

Chapter 3 discusses how this evolution has occurred and where monitoring is today. As you'll find, APM aggregates the lessons learned from each previous generation to create a unified system that leverages every approach simultaneously.

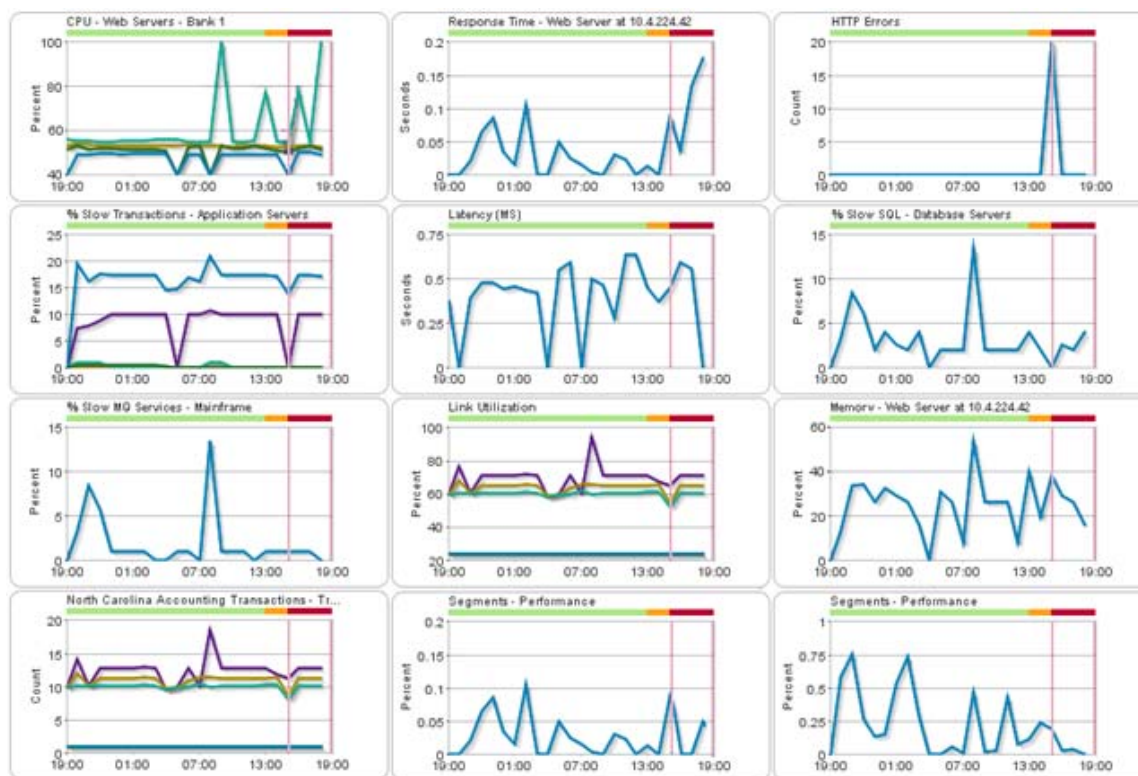
## Part 4—Integrating APM into your Infrastructure

Integrating an APM solution into your environment is no trivial task. Although best-in-class APM software comes equipped with predefined templates and automated deployment mechanisms that ease its connection to IT components, its widespread coverage means that the initial setup and configuration are quite a bit more than any “Next, Next, Finish.”

That statement isn’t written to scare away any business from a potential APM installation. Although a solution’s installation will require the development of a project plan and coordination across multiple teams, the benefits gained are tremendous to assuring quality services to customers. Any APM solution requires the involvement of each of IT’s traditional silos. Each technology domain—networks, servers, applications, clients, and mainframes—will have some involvement in the project. That involvement can span from installing an APM’s agents to servers and clients to configuring SNMP and/or NetFlow settings on network hardware to integrating APM monitoring into off-the-shelf or homegrown applications. As a result, an APM solution enables a level of objective analysis heretofore unseen in traditional monitoring.

The realities of that objective data are best exemplified through APM’s mechanisms to chart and plot its data. Figure 10.3 shows a sample of the types of simultaneous reports that are possible when each component of an application infrastructure is consolidated beneath an APM platform. In Figure 10.3, a set of statistics for a monitored application is provided across a range of elements.

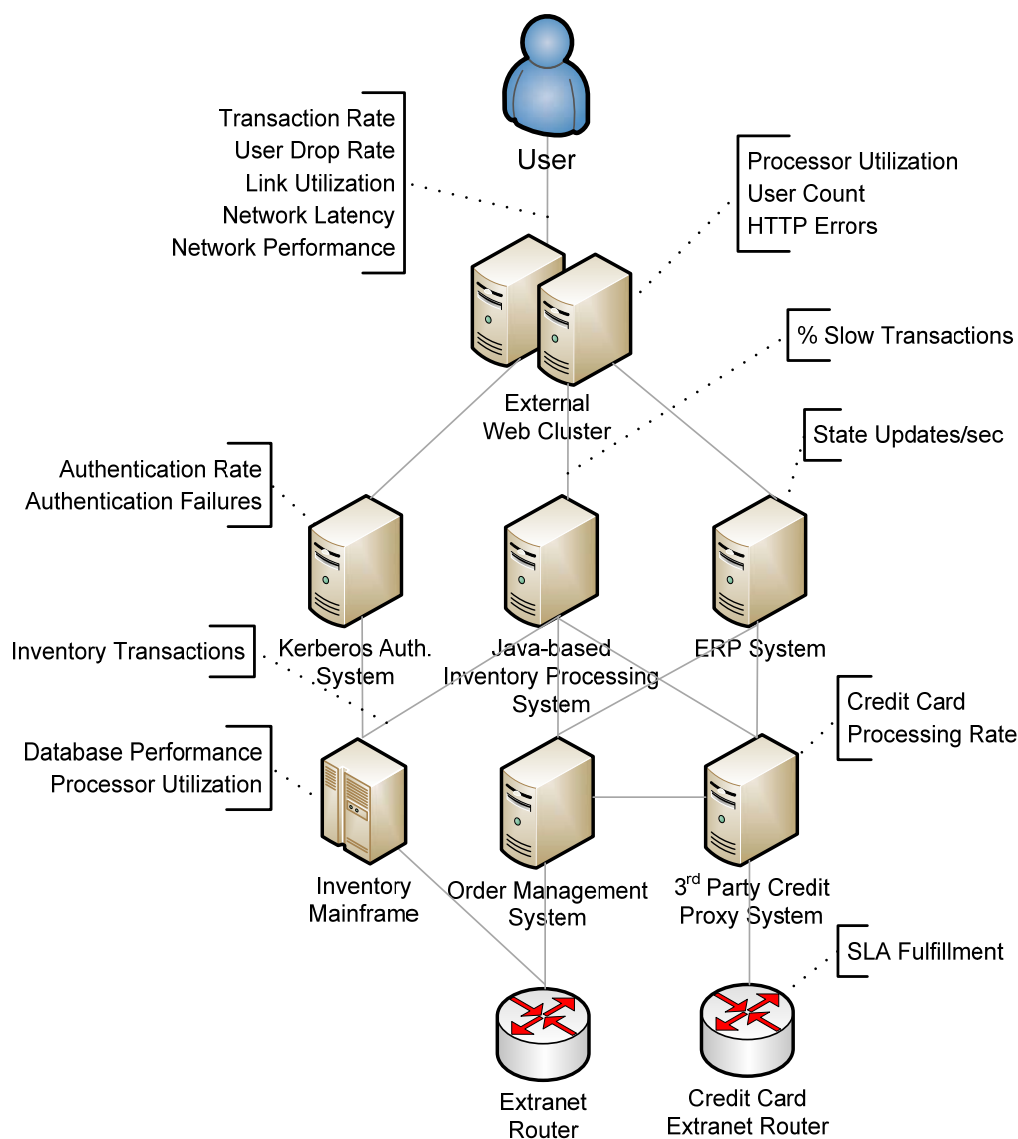
Take a look at the varied ways in which that application’s behaviors can be charted over the same period of time. Measuring performance over the time period from 10:00 AM to 7:00 PM, these charts enable the reconstruction of that application’s behaviors across each of its points of monitoring.



**Figure 10.3: APM's integrations enable real-time and historical monitoring across a range of IT components, aggregating their data into a single location for analysis.**

With the data you see in Figure 10.3, consider the points of integration where you might want monitors set into place. You will definitely want to watch for server processing. You'll need to record your network bandwidth utilization and throughput. You need to know transaction rates between mainframes and inventory processing.

All these monitors illuminate different behaviors associated with the greater system at large, and all provide another set of data that fills out the picture in Figure 10.3's charts and graphs. Now take a look at Figure 10.4, which shows how some of these monitoring integrations can be laid into place for an example customer-facing business service.



**Figure 10.4: Overlaying potential monitoring integrations onto a complex system shows the multiple areas where measurement is necessary.**

One end goal of all this monitoring is the ability to create an overall sense of system “health.” As should be obvious in this chapter, an APM solution has a far-reaching capability to measure essentially every behavior in your environment. That’s a lot of data. A resulting problem with this sheer mass of data is in ultimately finding meaning. Essentially, *you can gather lots of data, but it isn’t valuable if you don’t use it to improve the management of your systems.*

As a result, APM solutions include a number of mechanisms to roll up this massive quantity of data into something that is useable by a human operator. This process for most APM solutions is relatively automatic, yet requires definition by the IT organization who manages it.

The concept of “service quality” is used to explain the overarching environment health. Its concept is quite simple: Essentially, the “quality” of a service is a single metric—like a stoplight—that tells you how well your system is performing. In effect, if you roll up every system-centric counter, every application metric, every network behavior, and every transaction characteristic into a single number, that number goes far in explaining the highest-level quality of the service’s ability to meet the needs of its users.

Consider the graphic shown in Figure 10.5. Here, a number of services in different locations are displayed, all with a health of “Normal.” This single stoplight chart very quickly enables the IT organization to understand when a service is working to demands and when it isn’t. The graph also shows the duration the service has operated in the “normal” state, as well as a monthly trend. This single view provides a heads-up display for administrators.

Locations		Service Quality				
Service Title	Current	Duration	During Period	Available	Monthly Trend	
Atlanta, GA - Internap	Normal	14m 25s				
Houston, TX - Internap	Normal	14m 25s				
Internet	Normal	56m 17s				
Mesa, AZ - AT&T	Normal	27m 54s				
Philadelphia, PA - Level3	Normal	27m 55s				
St. Louis, MO - SAVVIS	Normal	27m 57s				
Washington, D.C. - Verizon	Normal	14m 25s				

**Figure 10.5: The quality of a set of services is displayed, showing a highest-level approximation of their abilities to serve user needs.**

Yet actually getting to a graph like this requires each of the monitoring integrations explained to this point in this chapter. The numerical analysis that goes into identifying a service’s “quality” requires inputs from network monitors, on-board agents, transactions, and essentially each of the monitoring types provided by APM.

## Part 5—Understanding the End User’s Perspective

With APM solutions, you’ll hear the term “perspective” used over and over in relation to the types of data that can be provided by a particular monitoring integration. But what really is perspective, and what does it mean to the monitoring environment?

It is perhaps easiest to consider the idea of perspective as relating to the orientation of a monitor’s view, which determines the kinds of data that it can see and report on. Although the computing environment is the same no matter where a monitor is positioned, different monitors in different positions will “see” different parts of the environment.



Consider, for example, a set of fans watching a baseball game. If you and a friend are both watching the game but sitting in different parts of the stadium, you're sure to capture different things in your view. Your friend who is sitting in the good seats down by the batter is likely to pick up on more subtle non-verbal conversations between pitcher and catcher. In contrast, your seats deep in the outfield are far more likely to see the big picture of the game—the positioning of outfielders, the impact of wind speed on the ball, the emotion and effects of the crowd on the players—than is possible through your friend's close-in view.

Relating this back to applications and performance, it is for this reason that multiple perspectives are necessary. Their combination assists the business with truly understanding application behaviors across the entire environment. An agent that is installed to an individual server will report in great detail about that server's gross processing utilization. That same agent, however, is fully incapable of measuring the level of communication between two completely separate servers elsewhere in the system.

### Monitoring from the End User's Perspective

Thus far, this guide has discussed how the vast count of different monitors enables metrics from a vast number of perspectives: Server-focused counters are gathered by agents, network statistics are gathered through probes and device integrations such as Cisco NetFlow, transactions and application-focused metrics are gathered through application analytics; the list goes on. Yet, it should be obvious that this guide's conversation on monitoring remains incomplete without a look at what the end users see in their interactions with the system.

This view is critically necessary because it is not possible—or, at the very least, exceptionally difficult—to construct this experience using the data from other metrics. Relating this back to the baseball example, no matter how much data you gather from your seat in the outfield, it remains very unlikely that you'll extrapolate from it what the pitcher is likely to throw next.

For the needs of the business application, end user experience (EUE) enables administrators, developers, and even management to understand how an application's users are faring. First and foremost, this data is critical for discovering how successful that application is in servicing its customers. Applications whose users experience excessive delays, drop off before accomplishing tasks, and don't fulfill the use case model aren't meeting their users' needs. And those that don't meet user needs will ultimately result in a failure to the business.

This line of thinking introduces a number of potential use cases where EUE monitoring can benefit an application's quality of service. EUE monitoring works for valuating the experience of the absolute end user as well as in other ways:

- Quantifying the performance characteristics of connected users as well as differences in performance between users in different geographic locales
- Simulating user behaviors through the use of robots for the purpose of predicting service quality degradations
- Identifying where internal users, as opposed to the absolute end user, are seeing a loss of service
- Keeping external service providers honest through independent measurements of their services

### Where Does EUE Fit?

It should be obvious at this point that there are a number of areas where EUE provides benefit to the business and its applications. Yet this chapter hasn't yet discussed how EUE goes about gathering its data. If end users are scattered around the region or the planet, how can an EUE monitoring solution actually come to understand their behaviors? Simply put, *the metrics are right at the front door.*

Think for a moment about a typical Internet-based application such as the one being discussed in this chapter. Multiple systems combine to enable the various functions of that application. Yet there is one set of servers that interfaces directly with the users themselves: the External Web Cluster. Every interaction between the end user and the application must proxy in some way through that Web-based system. This centralization means that every interaction with users can also be measured from that single location.

EUE leverages transaction monitoring between users and Web servers as a primary mechanism for defining the users' experience. Every time a user clicks on a Web page, the time required to complete that transaction can be measured. The more clicks, the more timing measurements. As users click through pages, an overall sense of that user's experience can be gathered by the system and compared with known baselines. These timing measurements create a quantitative representation of the user's overall experience with the Web page, and can be used to validate the quality of service provided by the application as a whole.



It is perhaps easiest to explain this through the use of an example. Consider the typical series of steps that a user might undergo to browse an e-commerce Web site, identify an item of interest, add that item to their basket, and then complete the transaction through a check out and purchase. Each of these tasks can be quantified into a series of actions. Each action starts with the Web server, but each action also requires the participation of other services in the stack for its completion:

- **Browse an e-commerce Web site.** The External Web Cluster requests potential items from the Java-based Inventory Processing System, which gathers those items from the Inventory Mainframe. Resulting items are presented back to the External Web Cluster, where they are rendered via a Web page or other interface.
- **Identify an item of interest.** This step requires the user to look through a series of items, potentially clicking through them for more information. Here, the same thread of communication between External Web Cluster, Inventory Processing System, and Inventory Mainframe are leveraged during each click. Further assistance from the ERP system can be used in identifying additional or alternative items of interest to the user based on the user's shopping habits.
- **Add that item to the basket.** Creating a basket often requires an active account by the user, handled by the ERP system with its security handled by the Kerberos Authentication System. The actual process of moving a desired item to a basket can also require temporarily adjusting its status on the Inventory Mainframe to ensure that item remains available for the user while the user continues shopping. Information about the successful addition of the item must be rendered back to the user by the External Web Cluster.
- **Complete the transaction through a check out and purchase.** This final phase leverages each of the aforementioned systems but adds the support of the Credit Card Proxy System and Order Management System.

In all these conversations, the External Web Cluster remains the central locus for transferring information back to the user. Every action is initiated through some click by the user, and every transaction completes once the resulting information is rendered for the user in the user's browser. Thus, a monitor at the level of the External Web Cluster can gather experiential data about user interactions *as they occur*. Further, as the monitor sits in parallel with the user, any delay in receiving information from down-level systems is recognized and logged.

A resulting visualization of this data might look similar to Figure 10.6. In this figure, a top-level EUE monitor identifies the users who are currently connected into the system. Information about the click patterns of each user is also represented at a high level by showing the number of pages rendered, the number of slow pages, the time associated with each page load, and the numbers of errors seen in producing those pages for the user.

User name/Aggregation block	Usage		Performance			Availability			Availability
	Pages	Slow pages ▼	Application performance	Page load time	Server time	Pages stopped	TCP errors	Errors	
combi004	997	137	86.3 %	3.63 s	897 ms	27	0	17	100 %
wmpdv004	430	128	70.2 %	7.48 s	1.57 s	21	0	2	100 %
doney005	4.33 k	119	97.2 %	1.52 s	488 ms	656	0	25	100 %
frabj005	449	118	73.7 %	5.96 s	1.21 s	147	0	5	100 %
bhcac001	575	93	83.8 %	4.3 s	1.22 s	11	0	9	100 %
dowbl004	896	88	90.2 %	3.29 s	883 ms	19	0	14	100 %

**Figure 10.6: User statistics help to identify when an entire application fails to meet established thresholds for user performance.**

Adding in a bit of preprogrammed threshold math into the equation, each user is then given a metric associated with their overall application experience. In Figure 10.6, you can see how some users are experiencing a yellow condition. This means that their effective performance is below the threshold for quality service. Although this information exists at a very high level, and as such doesn't identify *why* performance is lower than expectations, it does alert administrators that degraded service is being experienced by some users.

An effective APM solution should enable administrators to drill down through high level information like what is seen in Figure 10.6 towards more detailed statistics. Those statistics may illuminate more information about why certain users are experiencing delays while others are not. Perhaps one server in a cluster of servers further down in the application's stack is experiencing a problem. Maybe the items being requested by some users are not being located quickly enough by inventory systems. Troubleshooting administrators can drill through EUE information to server and network statistics, network analytics, or even individual transaction measurements to find the root cause of the problem.

## Part 6—APM's Service-Centric Monitoring Approach

This guide has spent a lot of time talking about monitoring and monitoring integrations. It discussed the history of monitoring. It explained where and how monitoring can be integrated into your existing environment. It outlined in great detail how end user experience (EUE) monitoring layers over the top of traditional monitoring approaches. Yet in all these discussions, there has been little talk so far about how that monitoring is actually manifested into an APM solution's end result.

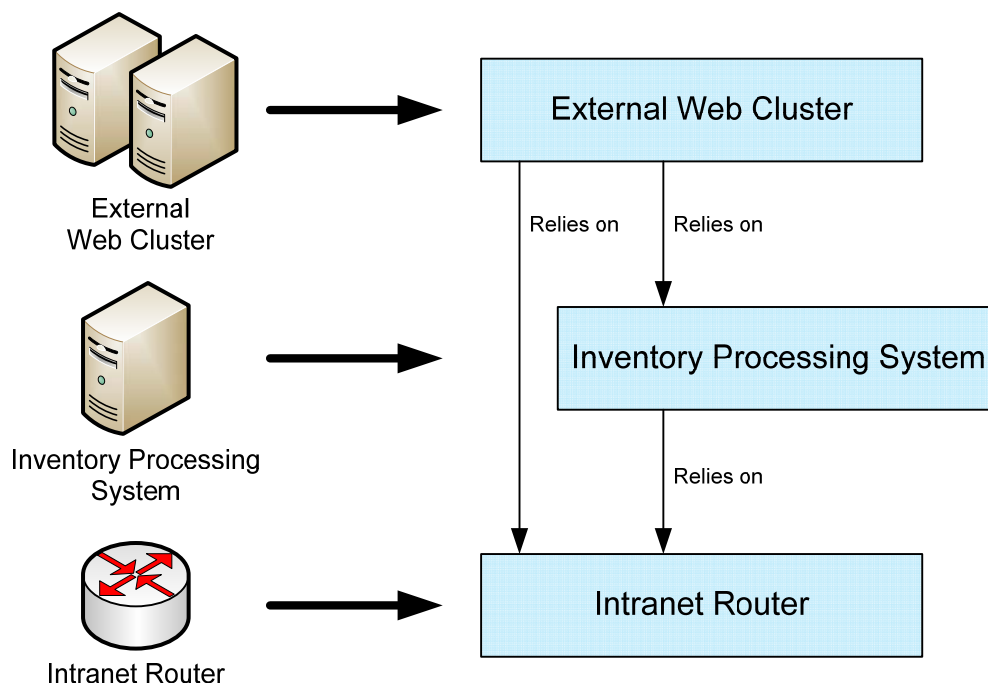
It is this process that requires attention at this point in our discussion. In reading through the first five chapters of this guide, you've made yourself aware of where monitoring fits into your environment. The next step is in creating meaning out of its raw data. As John mentioned earlier and as you'll discover shortly, the real magic in an APM solution comes through the creation and use of its *Service Model*.

To fully understand the quantitative approach to Service Quality, one must understand how the different types of monitoring are aggregated into what is termed a Service Model. This Service Model is the logical representation of the business service, and is the structure and hierarchy into which each monitoring integration's data resides. The Service Model is functionally little more than "boxes on a whiteboard," with each box representing a component of the business service and each connection representing a dependency. It resides within your APM solution, with the sum total of its elements and interconnections representing the overall system that the solution is charged with monitoring.

But before actually delving into a conversation of the Service Model, it is important to first understand its components. Think about all the elements that can make up a business service. There are various networking elements. Numerous servers process data over that network. Installed to each server may be one or more applications that house the service's business logic. All these reside atop name services, file services, directory services, and other infrastructure elements that provide core necessities to bind each component.

Take the concepts that surround each of these and abstract them to create an element on that proverbial whiteboard. This guide's External Web Cluster becomes a box on a piece of paper marked "External Web Cluster." The same happens with the Inventory Processing System and the Intranet Router, and eventually every other component.

By encapsulating the idea of each service component, it is now possible to connect those boxes and design the logical structure of the system. This step is generally an after-the-implementation step, with the implemented service's architecture defining the model's structure and not necessarily the opposite. Figure 10.7 shows a simple example of how this might occur. There, the External Web Cluster relies on the Inventory Processing System for some portion of its total processing. Both the External Web Cluster and the Inventory Processing System rely on the Intranet Router for their networking support. As such, their boxes are connected to denote the dependency.

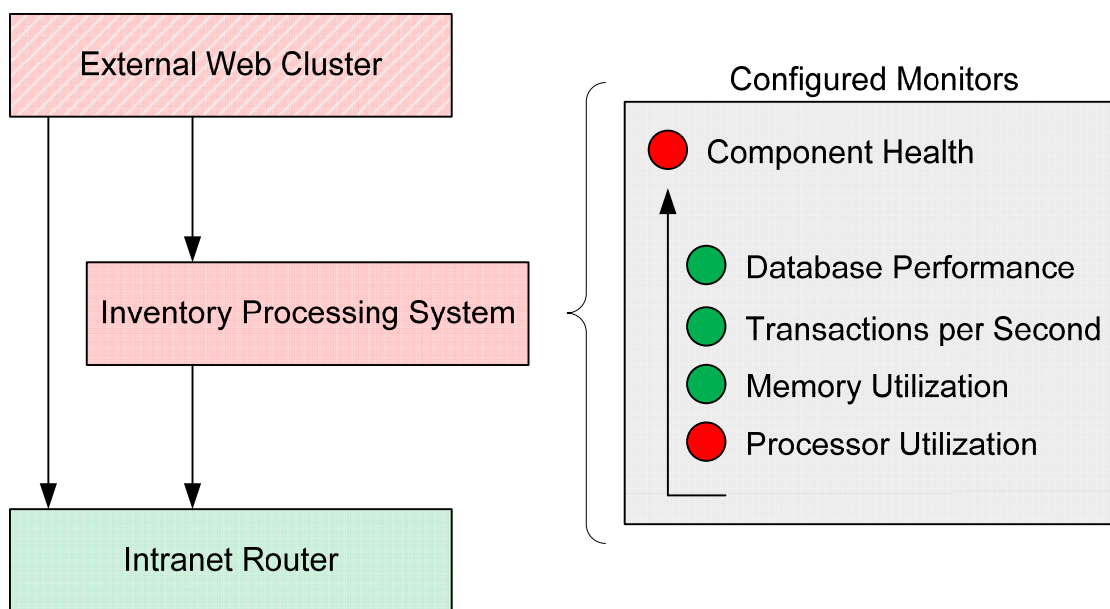


**Figure 10.7: Abstracting each individual component to create connected elements on a whiteboard.**

This abstraction and encapsulation of components can grow as complex or as simple as your business service (and your level of monitoring granularity) requires. One simplistic system might have only a few boxes that connect. An exceptionally-complex one that services numerous external customers—such as the one used by TicketsRUs.com—might require dozens or hundreds of individual elements. Each element relies on others and must work together for the success of the overall system.

This abstraction and connection of service components only creates the logical structure for your overall business service. Internal to each individual component are metrics that value the internal behaviors of that component. As you already saw back in Figure 10.4, those metrics for a network device might be Link Utilization, Network Latency, or Network Performance. An inventory processing database might have metrics such as Database Performance or Database Transactions per Second. Each individual server might have its own server-specific metrics, such as Processor Utilization, Memory Utilization, or Disk I/O. Even the installed applications present their own metrics, illuminating the behaviors occurring within the application.

With this in mind, let's redraw Figure 10.7 and map a few of these potential points of monitoring into the abstraction model. Figure 10.8 shows how some sample metrics can be associated with the Inventory Processing System. Here, the Database Performance and Transactions per Second statistics arrive from application analytics integrations plugged directly into the installed database. Agent-based integrations are also used to gather whole server metrics such as Memory Utilization and Processor Utilization.



**Figure 10.8: Individual monitors for each element are mapped on top of each abstraction.**

You'll also notice that the colors of each element are changed as well. At the moment Figure 10.8 is drawn, the Inventory Processing System's box is colored red. This indicates that it is experiencing a problem. Drilling down into that Inventory Processing System, one can identify from its associated metrics that the server's Processor Utilization has gone above its acceptable level and has switched to red.

Each of the metrics assigned to the Inventory Processing System's box are themselves part of a hierarchy. The four assigned metrics fall under a fifth that represents the overall Component Health. This illustrates the concept of rolling up individual metrics to those that represent larger and less granular areas of the system. It enables the failure of a down-level metric to quickly rise to the level of the entire system.

### Flow Up, Drill Down

Drilling down in this model highlights the individual failure that is currently impacting the system, but that specific problem is only one piece of data found in this illustration. As you drill upwards from the individual metrics and back to the model as a whole, you'll notice that the individual boxes associated with each component are also active participants in the model. Because the overall Component Health monitor associated with the Inventory Processing System has changed to red, so does the representation of the Inventory Processing System itself.

Going a step further, this model flows up individual failures to the greater system through its individual linkages between components that rely on each other. In this example, the External Web Cluster relies on the failed Inventory Processing System. Therefore, when the Inventory Processing System experiences a problem, it is also a problem for the External Web Cluster. The model as a whole is impacted by the singular problem associated with Processor Utilization in the Inventory Processing System.

It is the summation of all these individual threshold values that ultimately drives the numerical determination of Service Quality. A business service operates with high quality when its configured thresholds remain in the green. That same service operates with low quality when certain values flip from green to red and is no longer available when other critical values become unhealthy. The levels of functionality between these states become mathematical products of each calculation.

In effect, one of APM's greatest strengths is in its capacity to *mathematically calculate the functionality of your service*. Taking this approach one step further, IT organizations can add data to each element that describes the number of potential users of that component. Combining this user impact data with the level of Service Quality enables the system to report on which and how many users are impacted by any particular problem.

## Part 7—Developing & Building APM Visualizations

This guide's growing explanation of APM has introduced each new topic with an end goal in mind. That end goal—both for this guide as well as APM in general—is to gather necessary data that ultimately creates a set of visualizations useful to the business.

It is the word “useful” that is most important in the previous sentence. “Useful” in this context means that the visualization is providing the right data to the right person. “Useful” also means providing that data in a way that makes sense for and provides value to its consumer.

The concept of *digestibility* was first introduced in this book's companion, *The Definitive Guide to Business Service Management*. In both guides, the digestibility of data relates to the ways in which it can be usefully presented to various classes of users. For example, data that may be valuable to a developer is not likely to have the same value for Dan the COO. Dan's role might care less about the failure of an individual network component compared with how that component impacts the system's customers. Each person in the business has a role to fill, and as such, different views of data are necessary.



No visualization is effective unless it is created first with its consumer in mind. If that consumer can't digest what's being presented to them, the information being displayed is valueless. Think about the types of consumers who in your business today might benefit from the data an APM solution can gather:

- Service desk employees and administrators gain troubleshooting assistance and an improved view into systems health.
- IT managers are assisted in positioning troubleshooting resources to the most crucial problems as well as plan for expansion based on identified problem domains.
- Business executives gain a financial perspective and better quality data that is formatted specifically for their needs.
- Developers are able to dig into specific areas where code is non-optimized or requires updating.
- End users are proactively notified when problems occur, maintaining their satisfaction with your services.

A fully-realized APM implementation will include visualizations that provide the right kind of data to each of these stakeholders. Technical stakeholders get readouts on the stability of their devices and applications. Business leaders get a financial perspective. But each class of individual receives the information they need, which has been calculated from APM's singular database.

#### Note

With a picture really being worth a thousand words, consider turning back to Chapter 7 to see examples of APM visualizations for each of these classes of consumer. There, you'll see how APM's graphical representation of your business services enables a much improved situational awareness of their inner workings and impacts to the business.

## Part 8—Seeing APM in Action

Environments that benefit from APM's data-driven approach consolidate the problem resolution process into six very streamlined steps. This new process consolidates many steps from the traditional approach, while at the same time adding a few new ones that improve the overall communication between teams and to the rest of the business. Consider the following six steps as best practices for an APM-enabled environment.

### Visibility

Behaviors that occur outside expected thresholds are alerted via high-level visualizations. Through drill-down support, the perspective and data found in that high-level visualization can be narrowed to one or more systems or subsystems that triggered the failure. Using tools such as service quality metrics and hierarchical service health diagrams, triaging administrators can be quickly advised as to initial steps in problem resolution.

### Prioritization

Counts of affected users are predefined within an APM solution's interface, enabling triaging teams to identify the actual priority of one incident in relation to others that are outstanding. As a result, those with higher numbers of affected users or greater impacts on the business bottom line can be prioritized higher than those with lesser affect.

### Problem & Fault Domain Isolation

Triaging teams then work with troubleshooting teams, often through a work-order tracking system, to track the root cause of the problem. The same visualizations used before in the visibility step are useful here. Different from the unmanaged environment is that all eyes share the same vision into environment behaviors through their APM visualizations. As such, details about the problem can be very quantitatively translated to the right teams to assist in their further troubleshooting.

### Troubleshooting, Root Cause Identification, & Resolution

Using health metrics, the problem is then traced to the specific element that caused the initial alarm. That alarm describes how the selected element is not behaving to expected parameters. Here, troubleshooting administrators can work with other teams (networking, security, developers, and so on) to translate the inappropriate behavior into a root cause and ultimately a workable resolution.

### Communication with the Business

During this entire process, business leaders and end users are kept apprised of the problem through their own set of APM visualizations that have been tailored for their use. Business leaders see in real time who and how many people are affected by the problem as well as how much budget impact occurs. End users are notified through notification systems that give them real-time status on the problem and its fix.

### Improvement

Throughout the entire process, the APM solution continues to gather data about the system. This occurs both during nominal as well as non-nominal operations. The resulting data can then be later used by improvement teams to identify whether additional hardware, code updates, or other assistance is needed to prevent the problem from reoccurring. By monitoring the environment through the entire process, after-action review teams can identify whether the resolution is truly a permanent fix or if further work is needed.

It should be obvious to see how this six-step process is much more data driven than the earlier traditional approach. Here, every team remains notified about the status of the problem and can provide input when necessary through the sharing of monitoring data. When problems occur that cross traditional domain boundaries, those teams can work together towards a common goal without the need for war rooms and their subsequent finger-pointing.



**Note**

For a fictional narrative of the entire six-step process, consider turning back to Chapter 8. There, a made-up storyline is used to show how a fully-realized APM solution can and does improve the process of triaging, troubleshooting, provisioning resources, and eventually solving what would otherwise be an exceptionally painful problem.

## Part 9—APM Enables Business Service Management

Thus far, this chapter has shown how effectively resolving problems requires a data-driven approach, one with a substantial amount of granular detail across multiple devices and applications. Using this approach, it is possible to trace a system-wide performance problem directly into its root cause. By integrating into databases, servers, network components, and the end users' experience itself, a fully-realized APM solution is uniquely suited to gather and calculate metrics for entire business services as a whole.

Yet the topics in this chapter's story so far have been fundamentally focused on the technologies themselves, along with the performance and availability metrics associated with those technologies. Its resulting visualizations were heavily focused on the needs of the technologist:

- Service desk employees were able to track the larger issue directly into its problem domain.
- Network administrators were able to identify whether metrics for network utilization were within acceptable parameters.
- Administrators were able to use health and performance metrics to identify symptoms of the problem.
- Developers were able to ultimately identify the failing lines of code and quickly implement a fix.

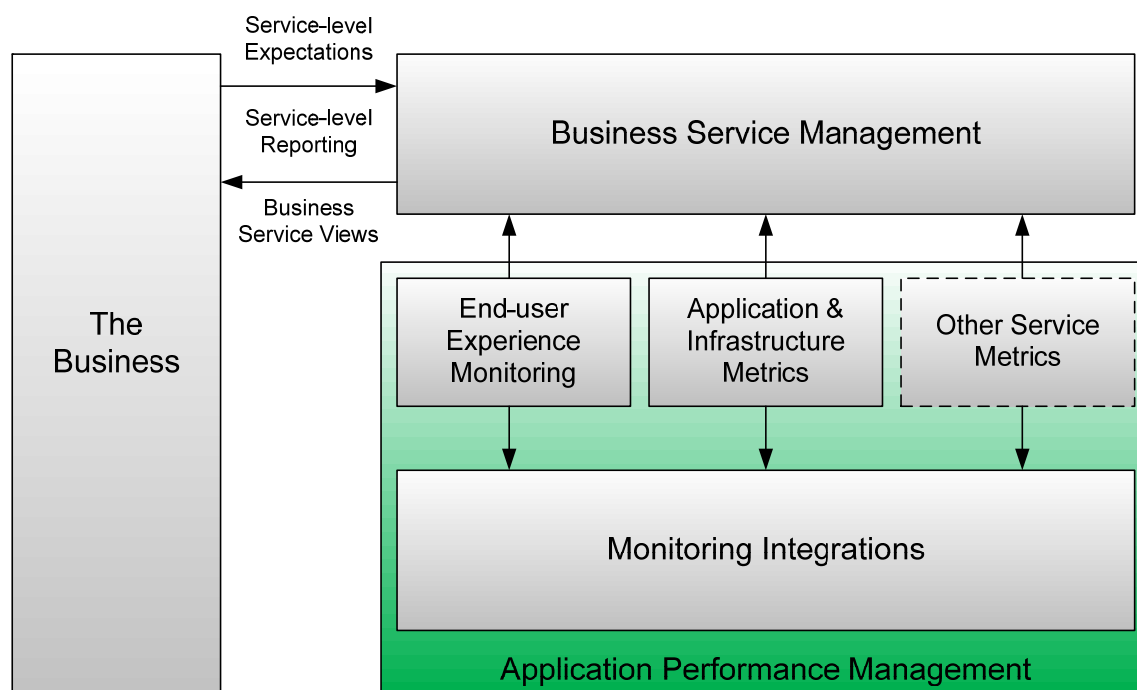
Missing, however, in the previous chapter's story is another set of business-related metrics that convert technology behavior into useable data for business leaders. This class of data tells the tale of how a business service ultimately benefits—or takes away from—the business' bottom line. It also creates a standard by which the quality of that service's delivery can be measured. It is the gathering, calculation, and reporting on these business-related metrics that comprise the methodology known as *Business Service Management (BSM)*.

## Linking BSM to APM

The IT Information Library (ITIL) v3 defines BSM as *an approach to the management of IT services that considers the business processes supported and the business value provided. Also, it means the management of business services delivered to business customers.* Businesses that leverage BSM look at IT services as enablers for business processes. They also look at the success of IT as driving the ultimate success of the business.

BSM and APM are two methodologies that are naturally linked by their requirements for data. The information gathered through an APM solution's monitoring integrations directly feed into the requirements of a BSM calculations engine. Performance, availability, and behavioral data of the overall business service and its components are all metrics that aid in calculating that service's overall return. These metrics also provide the kind of raw data that helps identify how well a business system is meeting the needs of its customers.

Figure 10.9 shows a logical representation of where BSM links into APM. Here, APM begins with the creation of monitoring integrations across the different elements that make up a business service. Those monitoring integrations gather behavioral information about the end users' experience. They collect application and infrastructure metrics as well as other customized metrics from technology components. APM's data by itself is used primarily by the IT organization for the problem resolution and service improvement processes discussed to this point in this guide.



**Figure 10.9: BSM converts technology-focused monitoring data into business-centric metrics.**

The addition of BSM creates a new layer atop this APM infrastructure. Here, the business itself becomes a critical component of the monitoring solution. Business processes and service level expectations are encoded into a BSM solution, with the goal of creating business service views that validate and report on how well the technology is meeting the needs of the business.

The metrics gained through a BSM implementation are also useful when fed into management frameworks such as ITIL or Six Sigma. Like BSM's roots in APM data, these frameworks are often highly data-driven in how they accomplish and improve upon the tasks of IT. One of the common limitations, however, in successfully implementing ITIL and Six Sigma framework processes is in gathering enough data of the right kind to be useful. The data gathering and calculation potential of a BSM/APM solution enables greater success with both frameworks.

BSM provides a substantial added value to this process through its identification and quantification of service quality. This quantification enables improvement teams to very discretely identify areas of gap in service delivery, develop appropriate solutions, and visibly see how well those solutions impact the overall quality of service delivery. In essence, using BSM's metrics, service improvement teams can measure the difference in as-is and to-be levels of service quality, proving that their improvement activities have indeed brought about improvement.

## APM Is Required Monitoring for Business Services

If this short "Cheat Sheet" has piqued your interest in the technologies and the business relevance of an APM solution for your business, consider turning back to its other chapters. Through both regular reporting as well as narrative storytelling, this Definitive Guide attempts to relate the tale of why APM should be required monitoring for business services. Through its comprehensive approach, you'll quickly find that an APM solution can and will bring vast amounts of value to your critical business services.

## Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.