

Realtime
publishers

"Leading the Conversation"

The Definitive Guide[™] To

Business Service Management

sponsored by

COMPUWARE[®]

Greg Shields

Chapter 4: Implementing BSM.....	69
BSM Provides a Business Focus to IT Operations	70
Three Reasons to Implement BSM.....	71
Understand the Critical to Quality Services.....	71
Manage Daily Risk and Improve Business Decision Making	71
Initiate Service Improvement Activities	71
The Seven Steps of a BSM Implementation.....	72
Step 0 – Preparation.....	72
Identify Key Project Members.....	72
Identify Stakeholders and Build the Project Plan	73
Step 1 – Selection.....	74
Identify Critical and Measurable Business Services.....	74
Assess Services	75
Assess Cost to the Business	75
Step 2 – Definition	76
Define Services	76
Define Service Requirements	78
Define Problems and Opportunities.....	79
Define Critical Success Factors	79
Step 3 – Modeling.....	79
Model Defined Services and Dependencies	80
Model Associated Metrics	81
Build the Service Model	81
Step 4 – Measurement.....	83
Implement Data Collection.....	84
Measure Services & Gaps.....	85
Step 5 – Data Analysis.....	86
Analyze Returned Monitoring Data.....	86
Validate Measurements & Costing Assumptions	86
Build Fault Tree Analyses	87
Build Impact Analyses	88
Step 6 – Improvement.....	89
Locate Problem Domains.....	90

Identify and Resolve Gap.....90

Revise the Service Model90

Step 7 – Reporting91

 Implement Dashboards91

 Implement Notification.....91

 Hand-off to Operations92

A Carefully Planned Implementation Is a Successful Implementation92

Copyright Statement

© 2007 Realtimepublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimepublishers.com, Inc. (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 Realtimepublishers.com, Inc or its web site sponsors. In no event shall Realtimepublishers.com, Inc. 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.

Realtimepublishers.com and the Realtimepublishers 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 Realtimepublishers.com, please contact us via e-mail at info@realtimepublishers.com.

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

Chapter 4: Implementing BSM

Its 11:15am the very next day and John and Dan find themselves sitting across the desk from each other in Dan's office, each thoroughly annoyed with the other.

"If you don't want the alerts, we can take you off the notification list," suggests John, "Or even better yet, we can configure the notification not to alert you in the middle of the night. You'll still get all of the accumulated alerts at 8:00am."

Dan realizes that John just isn't getting his point, "That's not where I'm going with this. There are two problems here. First, we shouldn't be getting these alerts in the middle of the night anyway. If you need a new router to replace the aging one, we'll get you a new one. But know that we should have a better capability to plan and budget for when these things occur. Second, if we take me off the notification list, we're back to where we were before the outage a month ago. Something goes 'pop' in the middle of the night and I don't hear about it."

"...but if we get an outage on a redundant system, we're still operational! The site is still up. Do you really care about every single device outage," John queries. He's similarly irritated by this conversation. For the past few weeks, every time his pager beeps, a phone call invariably follows from Dan. He himself wouldn't mind a week's worth of uninterrupted sleep.

Dan fires back, "I just don't want to take that risk right now. The outage a month ago cost us so much in overtime pay and customer givebacks that we blew our numbers for the quarter. We're thinking now it might impact the annuals too. I don't want to end up in McWilliams' office any more than you do," he says, referring to FCG's CEO Mike McWilliams, "but I will agree with you that all these alarms are interfering with the other work I need to do. You know, the COO work."

Dan and John stare at the wall opposite each other for a minute, both unable to think of what next to say. The problem here is evident in the minds of both individuals. Back about 2 years ago, FCG recognized the need to know when systems went down. They spent a not insubstantial amount of capital that year to implement a monitoring system intended to provide them with just the data they're getting paged on now.

What they're at the same time also realizing in Dan's office that day is that monitoring data is just that—data. Alerting on the outage of every device is doing more towards deluging them with notifications and less towards understanding the needs of their customer-facing systems.

Dan sits back in his chair, "What we need here is some way to turn that data into useable product. You need to know when the router goes down. I need to know when we're not servicing our customers. If one of our major customers can't get their supplies, it's my head that ends up in Mike's office and not yours. I want to know when their buyers are cursing our online system instead of loving it.

“Heck, here’s what I’d love. Get me another little monitor I can sit right here on my desk that just shows me how our systems are doing—what our customers are feeling when they’re doing business with our Web site. Something that’ll give me the warm fuzzy that our systems are up, we’re still meeting our numbers, and our customers are happy. Can you get me one of those?”

John groans silently to himself, “Performance data? Now he wants systems performance data too? How am I going to get that to him as ‘useable product’?”

Dan sees John’s concerns with his line of thinking, but he also recognizes the need for both John and his IT department to start thinking strategically. Maybe he can turn this challenge into an advantage for FCG, “Here’s what I’m going to do. It’s just about time we start setting performance goals for next year. I’m going to set a goal for you for next year to figure out the answer to this problem. I’ll take care of finding the funding and any business analysts you need. You just figure out the technology.”

John stands up to leave the office, wondering how he’s going to figure this one out. Dan stops him with a grin, “Oh, and John. Do it fast. My wife hasn’t been too happy either with your 4am D-E-N-R-T-R’s.”

BSM Provides a Business Focus to IT Operations

As you can see, our continuing story on First Class Glass (FCG) gives us a glimpse into the maturation of their organization. That maturity aligns with their need for ever-better data quality. Two years ago, FCG found that by implementing a monitoring system, they would be notified when systems go down rather than waiting for customers to tell them about it. After 2 years of middle-of-the-night pages, they have gotten better with parsing through the data provided by that notification system. But one thing is still missing from that data: business relevance.

This is validated by the way notifications are arriving into Dan’s mobile inbox. He is receiving data that doesn’t make sense to him. Dan embodies the “business” side of the business, while John embodies the technology side. John recognizes FCG’s need for a new router, but his focus on the tight IT budget and cost avoidance has led him towards running that device in production well past its useful life.

Throughout this guide, we’ve discussed how a well-defined BSM service model with links into the correct business systems can augment monitoring data with value. BSM provides a quantitative measure of the quality of a service by measuring it against financial rules specific to the business. Chapter 2 discussed how IT organizations must endure a process of maturation for them to recognize the need for data quality. Chapter 3 analyzed how that organizational maturity links to the evolution of IT Service Management and service management targeting. Our historical look there helped us better understand the gap filled by BSM.

In this chapter, we begin the process of implementing our BSM solution and its surrounding framework. We will begin by assuming that the implementing organization has the will and the way to incorporate the necessary software and processes to successfully complete the installation. The evolution of the implementing organization’s service management has elevated them to recognizing the need for BSM’s data quality measurements within their organization. All that is left is laying down the structure. At the conclusion of this chapter, you should be fully cognizant of the tasks and activities necessary to implement the BSM solution that is right for your environment.

Three Reasons to Implement BSM

Before we get into the technical discussion, let's frame this chapter's discussion around three critical reasons why an organization might want to implement BSM. We've talked about these reasons in generalities up to this point, but it is important for us to restate them here so that we recognize the underlying reasons BSM is important to our implementing organization.

Understand the Critical to Quality Services

A BSM implementation provides data to an organization's decision makers on how individual elements affect the whole. In our chapter example above, this is shown by the single server whose outage impacts the total operational standing of the customer-facing system. These elements we call *Critical To Quality*, meaning that the quality of our overall ability to service our customers is impacted by their reduction in service. As we'll learn later on, this is an important differentiation for us to properly scope which services should and should not be a component of our BSM service model.

Manage Daily Risk and Improve Business Decision Making


Also in our chapter example is the struggle of the business executive to make sense out of an unnecessarily complicated system as it is presented to them. There, Dan was unable to make the best business decisions in relation to servicing FCG's customers because "the system" wasn't presented in a way that makes sense. By providing digestible data to business leaders, BSM alleviates them from tactical decision making and focuses them on forward-looking, strategic initiatives. Conversely, John's daily operational focus on the IT infrastructure simultaneously requires data that helps him shift resources as necessary to manage systems based on business impact and solve problems. He requires a different set of data to help him reduce the daily risk to operations. BSM provides that data to him and his administrators.

Initiate Service Improvement Activities

Lastly, as was identified by the problem with the new router purchase, a well-designed BSM service model assists with the budgetary and planning process. That forecasting process, when fully recognized and populated with pertinent data, will facilitate better decision making for improving existing services. BSM is not intended to be a static implementation, but a rolling set of constant data review and remediation. The data provided through the BSM system aligns expenditures with where that money can be best allocated.

The Seven Steps of a BSM Implementation

Throughout the rest of this chapter we'll be discussing the seven steps involved with a BSM implementation. You'll note that we also include a first, eighth, step, titled *Step 0 – Preparation* that is needed to set up the teams, stakeholders, and project plan associated with the project. As you'll see in the sections below, the processes necessary to implement a BSM framework in an organization are non-trivial. Incorporating BSM into existing business processes is not an install-the-software-and-go procedure. In fact, installation of the enabling software doesn't even occur until *Step 4 – Measurement*. Let's begin with setting up the necessary teams and project plans associated with the preparatory step in the process.

 There are two important cautions regarding the information in the rest of this chapter. First, here we're attempting to show the process involved with the incorporation of a non-specific BSM system. There are multiple enabling software packages available on the market. Some incorporate the steps below and some may use different steps. The steps used here are included as an example of one way to set up BSM within your environment.

Secondly, in relation to the first point above, the information listed below is in no way comprehensive to the entire process. As you'll soon see, this process will likely take an extended period of time to complete – and in some ways is never truly complete. You will find tools and techniques that work well in your organization that may not in others. So be aware that these processes are malleable and should be customized to fit your particular organization.

Step 0 – Preparation

Before beginning any project, the identification of team members and stakeholders is critical for the division of responsibility within the project. Here within this step are a few key points necessary to ensure that the project begins down the correct path.

Identify Key Project Members

Firstly, as was discussed in Chapter 1, one of the most critical components of identifying a project team is the assurance of non-technocentrism. Although at first blush a BSM implementation can involve much impact from the IT organization, BSM in and of itself is a process-centric tool. The incorporation of too much technical input into the project team at the outset can have the tendency to turn a BSM implementation into little more than an IT Service Management implementation (e.g. with an inappropriate focus on IT elements).

That being said, a BSM project team should include the following members:

- **Executive Sponsor**—The role of the executive sponsor is to fund the project and ensure that that project stays within scope, budget, and relevance to its needs within the organization. The Executive Sponsor will likely not be a regular contributing member to the team, other than to provide overall guidance.
- **Business Service Manager**—Generally also the project manager, the Business Service Manager is tasked with ensuring the overall success of the project as well as reporting its status upwards to executive management. From a technical standpoint, the Business Service Manager is responsible for defining the business services of relevance and assisting with the development of their requirements.
- **Business Service Analyst(s)**—In conjunction with the Business Service Manager, any Business Service Analysts assigned to the project team have the responsibility for identifying and isolating individual business services, their requirements, and linkages between business services. Their job here is to create the business service model and populate that model with the necessary risks, linkages, and controls. Once the service model is built and implemented and *Step 5 – Data Analysis* has begun, the role of the Business Service Analyst is to monitor and interpret the data being generated vis-à-vis the model. This individual need not be of technical background, but rather a background with deep understanding of underlying business processes.
- **IT Manager**—Similar to the Executive Sponsor, this individual’s responsibility lies with overall guidance as well as to provide the liaison between the IT Specialists below and the Business Service Analysts above.
- **IT Specialist(s)**—Once the service model is identified, that model must be connected into data gathering and service monitoring tools. This function may be a part of the BSM system itself or more likely may be components of existing monitoring and management tools. The role of IT Specialists is to facilitate the proper connection of those tools into the BSM system.

Identify Stakeholders and Build the Project Plan


Also important in this first step is the identification of the ultimate stakeholders for the project. Often, when executive leadership is driving the implementation of the system, they become the stakeholders. As BSM tends to layer on top of existing monitoring systems, its incorporation tends to add greater value to leadership than IT itself (which may already have the necessary device monitoring tools in place). Looking at BSM’s tenet of service quality, one very obvious silent stakeholder in the project is the customer of the systems under management and ultimately the business.

Once stakeholders are identified, the project leaders must begin by creating a project plan that outlines each phase in the project. The next seven steps will assist with creating that project plan.

Step 1 – Selection



Step 1 embodies the identification of services that will ultimately become a part of the service model. Here the analysts on the team will analyze business services from a process focus and identify the lines of demarcation between individual services. Important here is that Step 1 is merely an inventory and identification function. We are not yet defining services and their representation. Here, we are merely getting our hands around those services that are in-scope, of value to us, and out-of-scope for this iteration of the project.

 The lead-in to this section mentioned that the BSM implementation can be a process that is “never truly complete”. The project team must be very careful at the outset of the project – indeed at this phase to keep the initial scope aligned to services that are **low hanging fruit**.

You need not necessarily identify all the services and processes in your organization during your first pass through the seven steps. Greater success is actualized by running through the steps more often with fewer elements in the model than the inverse. Iterating through the steps with a smaller model, especially during the initial adoption, provides early wins for the implementation.

Identify Critical and Measurable Business Services

The team must identify those services that are critical to quality. These will be core business services that provide measurable value to the company. “Measurable” such that components within the service can be quantified and a dollar value to the value of the component’s functionality can be assigned. At this phase that assignment is not yet done. However, services that have the capability of assignment are inventoried here in Step 1.

Shown as Figure 4.1 is a copy of the two service models we originally looked at back in Chapter 1. Here we see examples of a correct and an incorrect categorization of services as what must be accomplished during this phase. Recognize that services to be inventoried at this phase should align with a function of a business process rather than a function of the IT department. Later, the team will link these services to the resources that provide functionality to the service. A critical juncture here is not to get too technically “deep” during this activity.

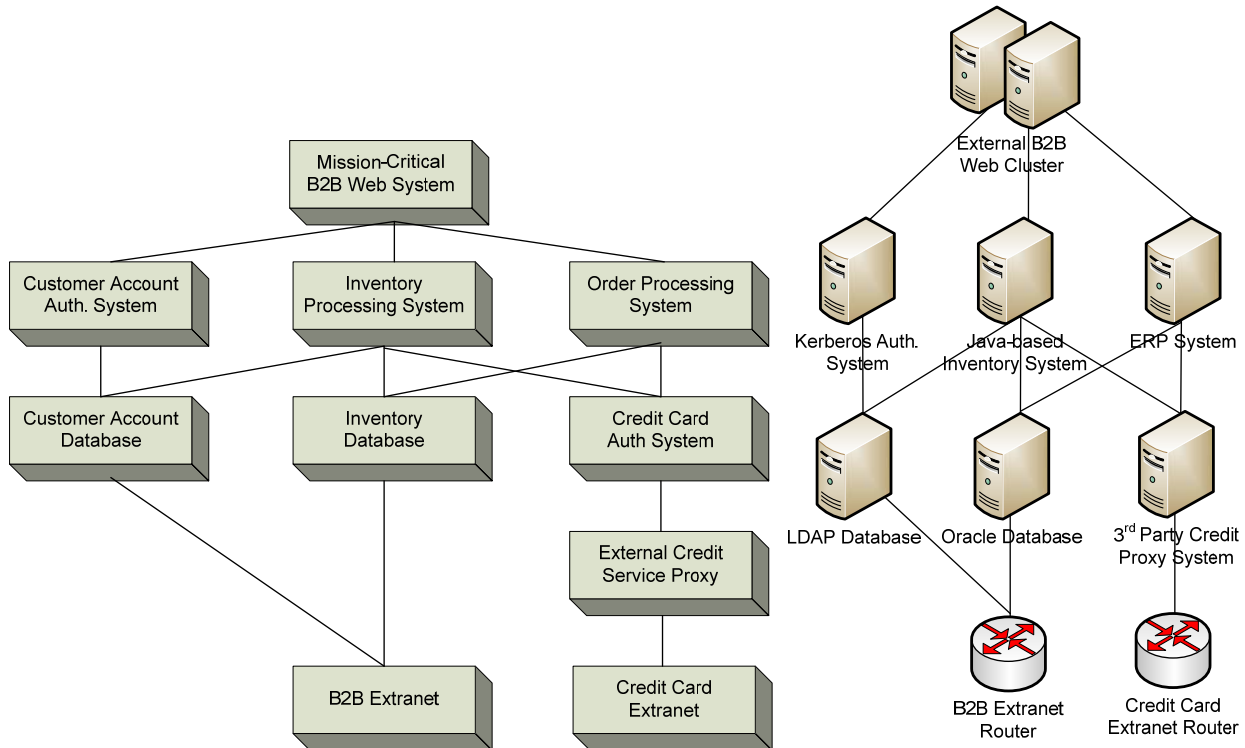


Figure 4.1: Copied here from Chapter 1 are our representations of a good service model breakdown on the left based on the interrelation of business processes. On the right is an incorrect model breakdown focused on individual devices.

Assess Services

While inventorying the services that the business provides, the feasibility of each service's ability to be easily categorized and quantified is also completed. Some business services are only tangentially related to the established Key Performance Indicators of the business, and so will be more difficult to quantify during early passes through the seven steps. The idea for our first pass is to find those services that are most critical to the business and yet are easy to incorporate into the service model.




Priority one here is to pick services that are most central and most critical to the business. Priority two is to choose those easiest to work with. The reason for this is that the action of quantifying "easy" services within an organization iteratively reveals new touch points for the later quantification of the "hard" services in later cycles.

Assess Cost to the Business

Once the rough inventory of gross business services is completed, a ranking of those services based on their business dependency is next. Here, the team will rank each of the business services by the level of impact to the business that could occur based on an outage, failure, or reduction of the service. This identifies that services' **Cost of Poor Quality**.

Step 2 – Definition




Once the inventorying of services is complete and the selection of candidate services for initial inclusion into the model has been made, those services need to be defined in terms relative to BSM. Within this, *Step 2 – Definition*, the team will identify and solidify the boundaries of the services of interest.

Define Services

The first step here is to gain as much knowledge as possible about the structure, behavior, necessity, and relevance of the service. This service might have ties into other services unknown to the project team or may have elements that make it more or less difficult to define as later steps begin deconstructing its dependencies. So by defining each service as comprehensively as possible in this initial step, much is learned about their inputs and outputs.

One mechanism for best documenting the service characteristics is to use a spreadsheet. Identify categorizations of interest about the service that will assist in later plugging this service into the BSM model. Some of those categorizations could relate to those in Table 4.1 below.

 When creating this spreadsheet, ensure that each cell within the spreadsheet is atomic. Hybridizing data within an individual cell means that that category has not been defined as elemental as is necessary.

Another handy tool to use in helping to visualize individual services is to draw up use cases and the associated data flow or process flow associated with that use case. As an example, for a purchasing system, the use case might include each of the components of a purchase, from browsing, to inventory validation, to shipping cart population, to checkout, to item delivery.

Categorization	Description / Utility
Service Name and Description	A unique identifier for the service, specifically one that can be easily identified both within the BSM system and in external documentation. Also, a short descriptor associated with the service and its functionality.
Business Purpose	What is the reason for the existence of this service? What is it intended to do?
Users	Who are the individuals who use this service? This information will help identify the business impact later on.
Service Hours	The hours of operation of the service. Does this service only run during business hours or must it be continuously operational. This information will help build out the business calendar. Also, what are the hours of support for this service? Is staff on-site continuously to support it? Are people on pager notification? This will help identify resolution time metrics.
Location of Service	Also useful for the business calendar, this information identifies the time zone and location of operation for the service.
Code Ownership	Is this is a home-built service or one purchased from another vendor? For customized services, what is the underlying code that drives it (Java, CORBA, C++, .NET, etc.)? This information is useful for pluggable end-point monitor devices.
Outage Impact	If this service goes down, who is impacted, why are they impacted, and what are they unable to accomplish. Are there mitigating factors (such a redundancy or lack thereof) to an outage that either exacerbate or alleviate the outage? This information later helps quantify the cost of the outage.
Abnormal Operation Impact	Much more difficult to measure, but what is the impact to the business if the service incurs non-nominal operations. What if it is slow? What if occasional hiccups in service cause the service to operate in an unpredicted manner yet without a failure? This information further quantifies the impact to the business when these situations occur. Part of this category is also the identification of which abnormal operations are to be measured.
RTO / RPO	This category identifies the service's Recovery Time Object and Recovery Point Objective, or how soon must it be brought back on-line and how much data can be lost as part of an outage. This information helps feed into disaster recovery metrics as well as individual service data loss metrics.
SLAs and OLAs	Have any Service Level Agreements been assigned to elements that make up the service? Or have any Operational Level Agreements been assigned across platforms or teams to legislate expectations?
Dependencies	Lastly, what other services does this service rely upon? This information will be heavily used in creating the service model.

Table 1.2: The table above provides list of possible characteristics that could be used to identify a service in the model.

Define Service Requirements

Now that we've come to understand the nature of the service in a narrative format, we need to translate the requirements of that service into quantifiable metrics we can use to measure its quality. This component may be one of the most important activities to be done for each individual service as this activity identifies the “numbers” by which the BSM system's mathematical logic uses to translate a loss of service quality into a numerical result. Three elements must minimally be identified and values assigned:

- **Availability**—The easiest of the three to quantify, during what hours of the day/week/month must this service be available for consumption by its users? Is it a 7x24x365 service, or is this service required only for operation between the hours of 8:00a and 5:00p? For some systems, like those that perform report generation or occasional data manipulations this number may even be merely a few minutes or hours per day – though at certain highly specified times. *Metrics to use here include:* Hours of operation, days and times of operation.
- **Reliability**—Slightly more difficult is the scoping of how often this service can become inoperative or undergo a loss in service performance. Some services have a greater tolerance for an outage. Some services include redundancy features that limit the scope of an individual element outage. Depending on how the service was categorized, those redundancy features may or may not be included in this calculation. Be aware that this information will be used heavily in identifying and measuring the quality of the service as compared with the desired level of service. *Metrics to use here include:* Acceptable mean time between failures, acceptable mean time to repair.
- **Performance**—The most difficult to identify, much of this quantification involves a qualitative look at the service's ability to perform to the needs of its consumers. Performance metrics identify the bars by which the operational service is measured. If reliability identifies how often the service can go down, then performance measures how poorly the service can operate and remain within desired specifications. *Metrics to use here include:* Per-action or per-transaction measurements of time delay, acceptable and unacceptable time-wait.

Define Problems and Opportunities


Related to the metrics identified above, this service was chosen for incorporation into BSM for a reason. Some component of its operation in some way causes pain to the organization. That reason should be identified in order to populate the project plan with data about the potential future success of the project in relation to solving a business problem.

Some potential questions to ask to help with the population of this information: Does the outage of this service cause an unexpected increase in cost to resolve (similar to the unexpected problems caused by the outage in our chapter example)? Is the outage of this service a locus for the outages of other services? Does this service's outage cause pain on the part of customers who may go elsewhere for their business? The problems and opportunities associated with this service, quantified into specific financial guidelines, helps further identify the service and its characteristics when built into the service model.

Define Critical Success Factors

Lastly for this step is the identification of what we hope to achieve by augmenting the monitoring of this service with BSM. Here, we want to provide metrics that document the improvements to service quality or availability we hope to achieve by going through this activity. As we discussed in Step 1, our first run through the model will likely be with those services that involve the most pain in our environment – and thus the most potential for return. The information here will be used later on in Steps 5 and 6 to help us improve the process and recognize benefit from the activity.

Step 3 – Modeling



Once an inventory of the desired business services has been collected, the connection of those services can begin. Looking above in Table 1.2, each service should have a list of dependencies. Those dependencies will go far in helping the team identify the connections between services.

The resulting service model will be a top-down decomposition of the business service in relation to its constituent components. One artifact of this process will be the creation of ever more detailed hierarchical diagrams identifying business processes in relation to the processes and resources that support it.

Model Defined Services and Dependencies

In Figure 4.1, the image on the left shows a series of 10 disparate system components that make up the top-level *Mission-Critical B2B Web System*. Each of these disparate elements feeds components above it, while each also requires information, processing, or resources from elements below it. In generating this model, we create a hierarchy of dependent components that (along with accompanying metrics) will eventually be plugged into the BSM system for logical and mathematical processing. To create this top-down diagram, four formalized tools are often used to assist with the visualization:

- **Failure Model Effect Analysis (FMEA)**—This is a tool used to identify and categorize the risks associated with potential failures within a system or a process. This tool identifies the possible failures that can occur within a system and prioritizes these failures by the seriousness of their potential consequences, their frequency of occurrence, and the ease in detecting them. FMEA is most often used as a bottom-up approach to failure detection. This augments the top-down approach to generating our BSM model. Here, the FMEA tool assists with identifying how the failure of a dependent component can impact the quality of the top-level service.
- **Component Failure Impact Analysis (CFIA)**—This tool, a component of ITIL, identifies individual ITIL Configuration Items and their potential for failure. Many of the same characteristics as FMEA are examined through its analysis. The difference here is that CFIA specifically analyses the potential for backups and redundancy between elements in an attempt to find flaws in a system design.
- **Fault Tree Analysis (FTA)**—A more top-down approach is the completion of a Fault Tree Analysis against the system. This thorough system for documenting the probability of fault amongst various logically linked situations helps in categorizing the risk of a system and where that risk may manifest. FTA is handy for adding numerical values into the BSM system.
- **CCTA Risk Analysis & Management Methodology—CRAMM** is yet another formalized mechanism, specifically suited for identifying security issues within a system. CRAMM breaks down the analysis into three stages: Asset identification and valuation, threat and vulnerability assessment, and countermeasure selection and recommendation. Complimenting the fault-based nature of the other two tools, CRAMM assists the team with identifying where security-based issues may impact the system.

Model Associated Metrics

As the tools are utilized, the service model “picture” begins to take shape. Once the picture has evolved to the point where it is realized to the satisfaction of the team, metrics associated with model elements are then attached to that picture. Attached Key Performance Indicators may relate to user wait time, business metrics, or IT metrics like systems and transaction performance. Mature organizations will likely already have many of the IT KPIs in existence, though their data population may not be automated. Business metrics may be similarly available. For organizations that don’t yet have KPIs in place to measure success, this may require an additional activity to find and document relevant metrics that make sense and provide value specific to the business.



This metrics assignment is another of the highly important steps in our process. One implements a BSM system because they are interested in obtaining these metrics and notifying when they go out of specifications. As we discussed in Chapter 2, organizations at lower levels of maturity have more work to do in implementing BSM, primarily because they need to develop the necessary Key Performance Indicators.

Build the Service Model

The process above continues the process of filling out the “picture”, adding metrics of interest to its structure. The concluding step in Step 3 is the finalization of the service model with all its requisite components. Earlier we’ve discussed how the service model shouldn’t look like the graphic on the right of Figure 4.1. But at some point, the business processes that make up the model must link into the data-generating tools that feed each process. It is here in the final step where individual processes and functions are mapped into IT functions. So where an actual database may drive the *Inventory Database* function or where a set of network devices may enable the *B2B Extranet* function, those linkages are created at this point.

Depending on the enabling BSM software chosen, this process may be an automated process or a manual one. It is not unreasonable to expect that some automation of the service model creation can be handled through the enabling software. After all it is that software that inevitably “runs” the model, so auto-discovery features are a good starting point for model creation. At the same time, it is unreasonable to assume that the software will be fully able to realize and draw the model without operator input. There are just too many possibilities for processes and their ties into network devices and applications.

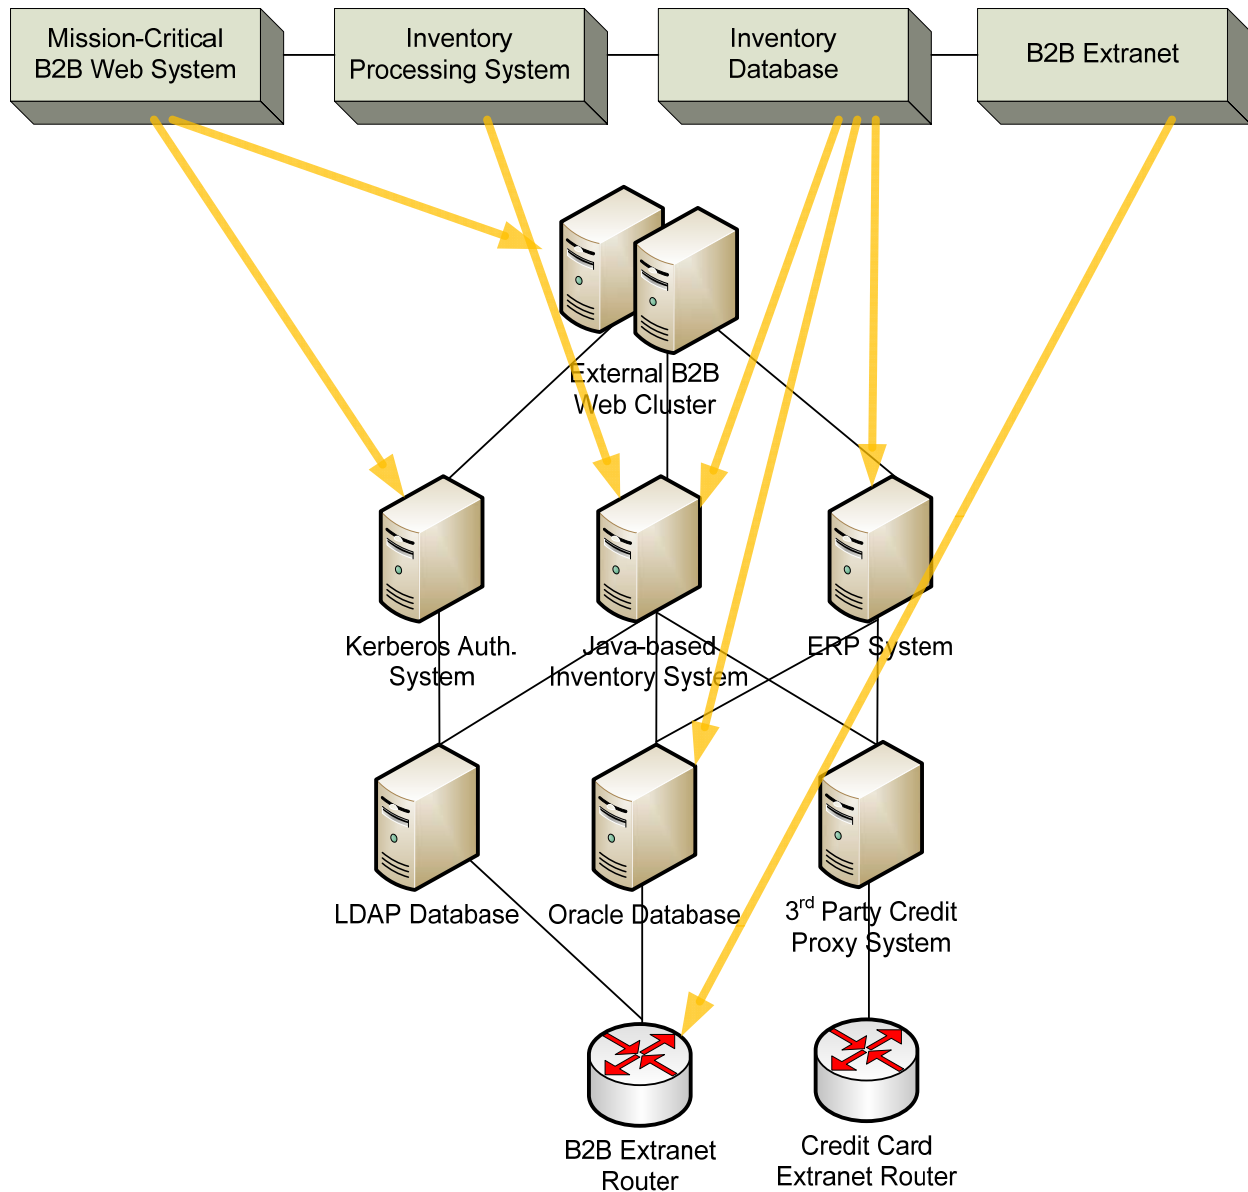



Figure 4.2: Once the service model is fully realized, the next step is to connect its processes to the IT functions that drive its data. This mapping is used by the BSM system to populate the model with metrics information.

Step 4 – Measurement



Once the modeling is complete, our next step is to link the designated monitoring and measuring tools into the service model. It is within this step where much of the effort within the enabling BSM software tool begins. Here, for those services and their associated metrics previously identified, categorized, and modeled, we'll begin the process of actually measuring the metrics we aim to obtain. In later steps we'll take this information, analyze it for gaps in service, and use it to drive change within the environment design.

In Step 0 of this process, we discussed how much of the BSM implementation process does not necessarily require heavy involvement with the IT organization. However with Step 4 comes much of the work needed by IT specialists. Here, the team will be implementing or otherwise coding the necessary connectors that pull data from disparate systems into the BSM software platform. Those skills are often highly-specialized and often are specific to the type of software platform the BSM system attempts to connect into. It is important for IT to be part of the process up until this point so they can prepare the systems from a technical standpoint for the monitoring plug-ins necessary to begin measuring.


 Remember too that BSM is not intended to “rip and replace” existing monitoring systems. Nor in many cases is it intended to be a systems monitoring system of its own. The organization likely already has monitoring tools in place that leverage technologies like SNMP, NetFlow, WMI, WS-Management, and other management protocols on which monitoring data is already being collected and stored. The BSM implementation can simply pull from that data for its metrics needs.

Implement Data Collection

The first step here is to tie the BSM system into infrastructure monitoring. This tie-in may be with built-in connectors that ingest into the BSM system or regularly output from the monitoring system. These data collection tools are configured to query for and collect the metrics identified in the model as designed in the previous step.

Some examples of existing data collection aggregators and elements that may already be present in your environment include:

- Java Message Service
- WS-* / WS-Management
- JDBC Data Access
- Log File Reading and Scripts
- Command-Line Interface Tools
- Messaging and Messaging Connectors
- File Transfer Statistics
- Port Services (HTTP, HTTPS, DNS, LDAP, etc.)
- Windows Management Instrumentation
- Storage Management Tools
- Enterprise Monitoring & Eventing Tools
- SNMP and other Network Monitoring Tools
- Databases and Database Monitoring Tools
- Inventory and Inventory Management Tools
- CMDB & Service Desk Connectors

 Recognize that this step is non-trivial. Whereas previous steps involved “moving around boxes on paper” to find the correct representation of the model, here we are actually manipulating computer code to enable the connections. This process should involve the same levels of Configuration Management hopefully already within the system – the process of which requires substantial testing and validation before integrating into a production environment.

The project team can eliminate one very important “gotcha” by ensuring enough time is built into the project schedule to acquire the proper talent and properly build, test, and implement these connectors.

End user experience monitoring tools may additionally be attached into the customer-facing interfaces of externally-facing systems. We'll talk at length on these types of tools in the next chapter. But for now know that the code frameworks that typically drive these customer-facing tools often have built-in monitoring toolsets that allow for the integration into the BSM system. Processes like synthetic queries and scripted actions can simulate the load of a particular user and determine their wait time (e.g. their "experience") while using the system. This information ties into KPIs associated with customer satisfaction.

The BSM environment may also tie into Service Desk applications to get a time-based understanding of how user experience drives incoming requests and complaints. One effective KPI for measuring the quality of an external service is to monitor for incoming tickets alongside end user experience monitoring. By doing this, an organization can discover what the pain points are with their particular brand of customer. Some customers may be more or less willing to handle elements of pain within systems. The rate of generated ticket workloads can drive a better understanding of how those users are experiencing the system.

Measure Services & Gaps

As the team begins to implement data collection tools around the network, the BSM system will begin measuring the quality of each listed service that makes up the model. Areas where data is not yet incoming will show as gaps in desired metrics. We are not yet to the step where we can begin implementing reporting and dashboards to visualize those metrics, so careful attention to system data as it arrives into the system will identify where KPIs are being measured and where gaps still exist.

At this point, a review of the existing data coming in as related to KPIs currently in place or desired to be in place by the organization is an excellent double-check against the service model. The service model, though considered "frozen" in its first iteration by the project team, may require additional work to pull the necessary data required by the system. This can also manifest as calculations that are lacking necessary data to properly represent loss as a measure of service quality.

Step 5 – Data Analysis



Once the initial connectors are in place, the BSM system begins collecting data from the various systems throughout the network. The BSM system, when configured with appropriate metrics and logic associated with those metrics will apply cross-device and cross-application computations to determine health and quality status.

Within Step 5, we will begin the process of analyzing the incoming information and trending that information to see if the data we've expected to receive aligns with the data we intended to receive. Once this process is complete and we are ensured the validity of the model, we can begin to analyze the system to see where gaps in service occur. This may be based on bad service quality or customer ratings, system overloading, element response time, or transaction throughput.

Two tools used to find these gaps that we'll discuss later in this section are Fault Trees and Impact trees. A component of the service model, these two tools identify where the root causes and overall impacts to a service degradation or outage occur.

Analyze Returned Monitoring Data

Initial incoming data arrives in a relatively raw format. This raw data often needs to be converted into a format useable by the calculations required elsewhere within the system. The process of converting this data may involve multiple steps as data may require multiple refactoring based on the target metrics required for it. As an example, performance data may arrive in a binary format and need to be refactored into a numerical format for analysis in comparison with outage metrics. It may require additional refactoring into a third format based on measured time to compute its data in comparison with performance data from other kinds of systems.

Validate Measurements & Costing Assumptions

During this "learning" mode for the model, it is also important to validate that the measurements as identified by the project team are to scale and include correct units. Converting KPIs into measurable statistics can be a complex mathematical and logical task which can involve unit comparison and conversion between multiple systems. This depends on each individual system's capability to supply the data in the correct units. It's also essential to validate measurements taken against the original data to make sure they are computationally correct.

One function of validating necessary measurements is the ratification of costing assumptions made during the model generation. The initial determination as to the cost associated with a loss in service quality or a loss of service altogether may have been based on faulty or misleading information. Or the data arriving into the BSM system may not confirm the assumptions. It is important here for the reliability of the system that the earlier costing assumptions are related to the actual data arriving in-system.

Build Fault Tree Analyses

Once the model is validated as correct the data within the model can be analyzed in comparison with desired metrics to help identify where individual components of the system are not performing to specifications. Now that the model is in place and functional, the reduction in quality of each individual element can be related to how that element affects the whole. For example, in Figure 4.3 the completed service model now shows how a change in performance of the *Inventory Database* impacts each of the services that rely on the *Inventory Database*.

Here we can see that the performance of our *Inventory Database* has gone above our desired specification of 5000 transactions per second. That reduction in quality directly impacts the *Inventory Processing System's* capability to respond to inventory requests fast enough. It also affects the *Order Processing System's* capability to fulfill orders as a fulfilled order will change the level of inventory. Ultimately, each of these metrics directly relates to the customers satisfaction or dissatisfaction with their experience.

We've discussed thus far about how a major component of the reporting step involves the digestibility of information specific to its consumer. Here we see how this information is immediately of value for multiple consumer classes, depending on how it is presented. Non-technical executives and business leaders can get a high level representation of the system and associated (lack of) service quality. They can use this data to make decisions about the business in general or additional purchases to augment the design. Administrators gain necessary information to help them quickly troubleshoot the problem.



It's worth noting here that prior to having this model in place the organization may not have been able to trace how unhappy customers were impacted by a loss of performance in a down-level system. BSM's built-in fault tree analysis tools provide both the IT department as well as the business leadership the data they need to make the right decisions. That decision may be to purchase a second load-balanced *Inventory Database* or vertically scale the existing one.




Figure 4.3: Our completed service model begins to show how a fault in the *Inventory Database* – in this case, the number of transactions per second going above specifications can impact the systems above that rely on that database.

Build Impact Analyses

In concert with the generation of fault trees are impact analyses. This tool illustrates to the viewer the anticipated impact to the service based on a component outage or reduction in service. Notice in Figure 4.3's fault tree above that the information provided helps in troubleshooting the problem down to the individual element level. From the perspective of an impact analysis, Figure 4.4 shows another representation of our service model with the same failure state. Here we gain a perspective on how that fault is impacting our operations.

You can see that the same out-of-spec transactions rate is impacting the *Inventory Processing System* to the tune of 37 inventory changes or updated being delayed beyond specifications. The *Order Processing System* similarly cannot keep up with the load of the *B2B Web System*, setting 17 orders to a wait state prior to approval and final purchase due to the problem. Ultimately, 20 Users elect to cease their transaction (e.g. the drop rate) due to this problem. Relating the drop rate further to lost revenue allows the BSM service model to recognize the loss of revenue associated with the *Inventory Database's* problem.

Whereas the fault tree visualizes the troubleshooting component of *fixing the problem*, the impact tree assists with recognizing *how the problem affects users*.

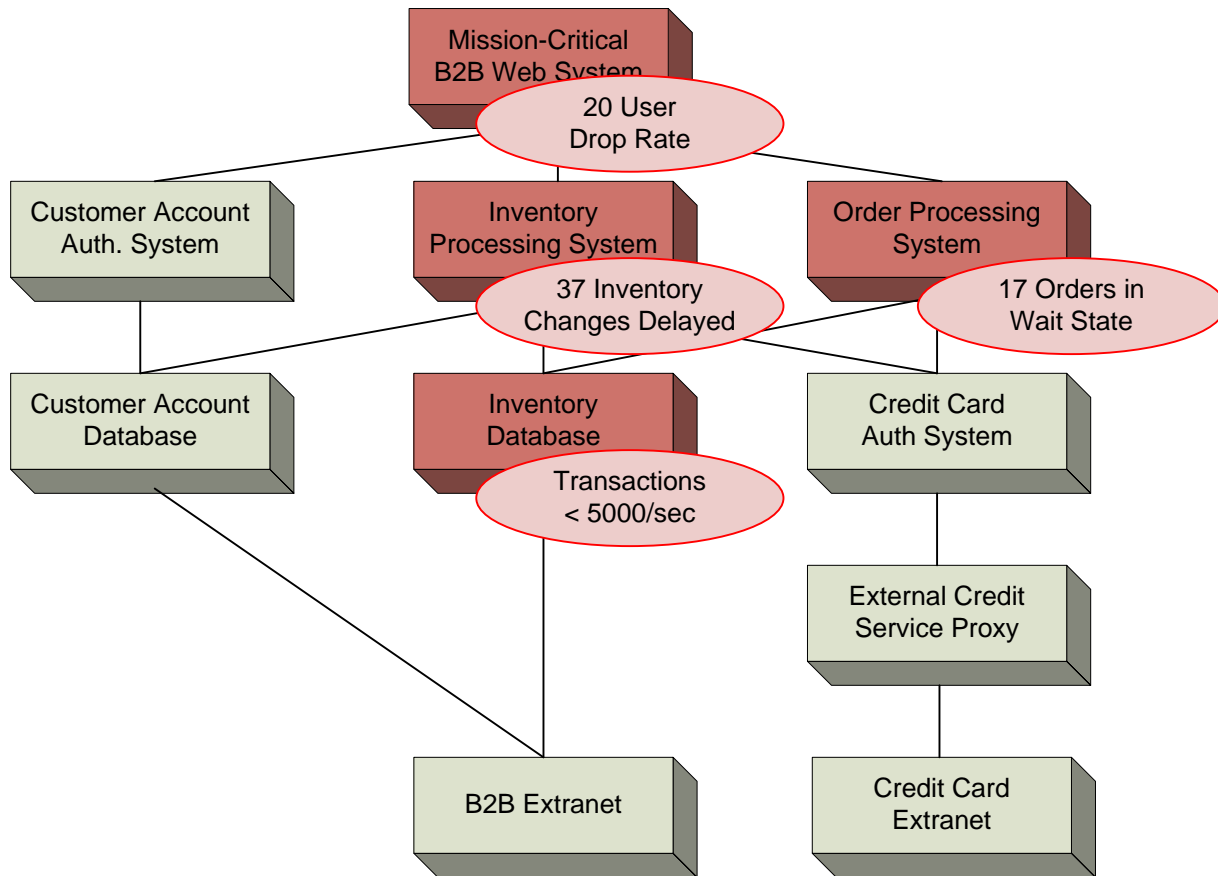




Figure 4.4: Relating the information within our service model to an impact analysis, we begin to see how the out of specification performance of our Inventory Database is directly impacting other functionality. Ultimately, our system is seeing a 20 User Drop Rate due to this problem.

Step 6 – Improvement



Step 6 in our process asks the question, “So, now what do we fix and why?” Up until this point our process has been driven by the need to populate the framework with data. Also a component up to this point is the analysis of the collected data to find where gaps exist in the best design of our system. Here in Step 6 we finally get to take what we’ve learned thus far and turn it into productive change for the service under monitoring.


 It's worth mentioning here that it is entirely possible that no improvement may be needed. Monitoring is, by nature, a long-term activity. Thus, the time delay between Step 5 and Step 6 may in actuality involve a period of time. Two elements can characterize this time delay.

First, there may be no problems whatsoever with the service. Though we all wish this was the case in our systems, the selection criteria for our first service was to find the one that was already causing us the most pain. So, the likelihood of this occurring is low. What is more likely is...

Second, finding no actionable data within our system means that something within our model is likely missing. That may be an unknown connection to an IT function, a missing step in the process, or a metric that is either missing or mischaracterized. If the project team finds themselves reaching Step 6 and finding little to action upon, circling back to Steps 2 or 3 may be in order for additional discovery.

Locate Problem Domains

Assuming that the service model in place is accurate and the metrics being collected provide value to the organization, visualization tools within the BSM system similar to the fault trees and impact trees discussed above will help the Process Analysts and the Systems Administrators locate problem domains. It is important here for these two halves of the team to work hand-in-hand in this process. Often the problems found are IT-related problems that manifest in a total reduction of service quality.

 Fault trees and impact trees are not the only visualization tools provided by BSM to aid in troubleshooting. Chapters 7 and 8 will both touch on additional tools that BSM can provide to assist with this step.


Identify and Resolve Gap

Less often, but arguably more critical, is the identification of process gaps. In these cases, the process itself is actually causing the problem. In these cases, however, the incorporation of BSM data into the justification provides quantifiable evidence of the need to change a process. For example, if the BSM system begins to find that an unnecessary step in the *Credit Card Authorization System* is consuming an out-of-spec time delay in completing a transaction, the team may decide that the solution is to remove that step from the process.


Revise the Service Model

Lastly, in some cases the service model itself may be incorrectly represented. It is probable, especially during early run-throughs of the model, that the service model may be lacking some necessary component, metric, or connection, for it to properly process incoming data. Recognition of this problem is arguably the most difficult to resolve as incoming data may mask the true problems with the model itself. Earlier in this chapter we discussed how the model itself should not be considered a static and immobile construct, but instead an organic representation. Networks change. Applications and services are added and removed to the environment. So a continuous review of the model in relation to its connection with various IT functions is critically necessary as its IT underpinnings or processes themselves may change over time.

Step 7 – Reporting



Our final step in running through the seven steps is setting up the dashboards and associated reports. This step is set to last as it gives the project team time to analyze data and refine the model before tying down the model to specific reporting functions. Once reporting is configured, it has the tendency to “freeze” the model. This is the case because model and characteristic or metric changes will impact reports.

 Once reports are enabled for consumers, they typically grow reliant on them. So, be aware that making reporting available is often a step that involves strong configuration control.

We will actually discuss in detail the process of creating reports and dashboards in Chapters 6 and 7. So for this chapter we’ll review only from a high level the steps necessary.

Implement Dashboards

Dashboards are often “skinned” web sites that dynamically update data as necessary. We say “skinned” because a visualization image is often used as the base layer. Atop this base layer are overlaid data representations like Stoplight charts, Control charts, Pareto charts, and other tools that represent numerical data in a graphical format. There are some best practices for developing dashboards that we’ll discuss in detail in Chapter 7. But for now know that this step will involve equal measures of graphic design, data manipulation and visualization, and knowledge of web platforms.

Implement Notification

Notification elements can also be added into the BSM system. Similar to the problems we saw within our chapter example, the executives and business leaders in the organization want to be notified when situations occur that they can understand and resolve. They aren’t necessarily as interested when individual IT functions (like John’s extranet router) have problems that they are unable to act upon. Once the BSM implementation is in place, notification components can be enabled that provide valuable data.

As a continuation of our chapter example, once FCG implements their own BSM system, they may consider removing Dan from all the device notifications he is currently receiving. Instead, he may want to know when the drop rate for the web site goes beyond a certain metric. The failure of an individual router means little to him and his position and skill set can’t add value to its resolution. But when the web site’s drop rate goes below acceptable parameters, that situation impacts FCG’s bottom line, which is a problem that he can identify with and help to resolve.

Hand-off to Operations

The final step in this process is the ultimate hand-off to operations. This step involves fully documenting the configuration, entering that configuration into the organization's configuration control tool, and dissolving the project team or redirecting it to another spin through the seven steps for another service.

Another component of this hand-off is the procedures necessary to monitor and maintain the model. Spinning up another project team for minor updates to the model can be a waste of resources. Thus, one final task for the project team will be to document those procedures in enough detail that individuals later on can affect change.

A Carefully Planned Implementation Is a Successful Implementation

This chapter has attempted to provide some real-world ideas and information on how best to begin the process of implementing BSM software and methodology within an organization. We've talked about how each of the seven steps iteratively brings the product team from an empty sheet of paper to a fully-recognized service model with attached metrics, IT functions, reporting, and notifications. Throughout the last three chapters we've talked about BSM's intent to providing information that is understandable to its consumer. In this chapter we've shown how implementing it actually realizes that expectation.

One element of BSM that we've glossed over to this point is the measurement of experience from the perspective of the end-user. Looking at a service's endpoints provides valuable information not typically gathered by other monitoring tools. The information gathered there specifically populates certain metrics within the service model. In Chapter 5, our next chapter, we will sidestep into this discussion. We'll talk about how to set up End User Experience Monitoring, how best to manage it, and where it ties into the greater BSM picture.

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>.