cesses as a visiting fire fighter. If quality improvement involved calling in experts, we could simply hire engineers and be done with it. The real development of organizationwide capabilities in quality improvement is slower but in the end immensely more powerful. For many employees, participation in a quality improvement team is their first introduction to the skills of process improvement that will, from then on, be part of their daily work.

6

The Diagnostic Journey

If we think of quality improvement itself as a process consisting of *defining the problem, making the diagnosis, administering the remedy,* and *holding the gains,* then we may ask, as we may of any process: What does each step receive from the previous step—its supplier—and provide to the next step—its customer? We have seen how the teams in the National Demonstration Project arrived at a statement of the problem they would be focusing on. Using the process described and analyzed in Chapter Four, each group "supplied" a clear, well-informed problem statement that could then serve as the starting point for the next step: making the diagnosis.

As we will see, the "diagnostic journey" is in many cases a long and tortuous one; complicated issues often underlie the simplest diagnostic questions. Quality improvement has its own "black bag" of many specialized instruments—process flow diagrams, histograms, Pareto diagrams, Ishikawa diagrams, and more—for making the diagnosis (these tools are described in detail in Resource B). Despite the many new terms and tools introduced in the diagnostic enterprise, however, making the diagnosis is a conceptually simple procedure. It consists of two basic steps:

1. Defining and understanding the existing process.
2. Analyzing the existing process to determine where the flaws—and thus the opportunities for improvement—lie.

Despite the variety of problems and areas of health care tackled by the teams in the NDP in the course of their diagnostic journeys, we will see that their experiences have much in common. The departure point is a well-framed problem statement. The group begins the diagnostic journey by making a map of the process as it currently exists, continues by analyzing the process to generate hypotheses about possible flaws, and then moves to gathering data to test whether *those* particular flaws exist at *those* particular points in the process. Their work is an exact analogue of the scientific method used in modern research: ask a question, formulate hypotheses, and then, using data, seek to confirm, reject, or modify the hypotheses.

By the end of the diagnostic journey, a project team should be able to answer the following questions: "What goes wrong?" "Where?" and "How do we know?" The answers to these questions, in turn, are key inputs to the next step—the remedial journey—which has its own agenda for hypothesis generation ("What might *help* here?") and testing ("Does the remedy *work*?").

This chapter analyzes the diagnostic journeys of several project teams, using them to demonstrate three key stages of the journey: first, understanding the existing process; second, generating and testing hypotheses about where and why the process is flawed; and third, displaying and understanding data.

Defining and Understanding the Existing Process: Two Case Studies

In quality improvement efforts it is always helpful to ask, "what is the process that we intend to improve?" Defining existing processes as clearly as possible is a key step in the work of a team, as two specific cases from the NDP illustrate.

Kaiser-Permanente. The experience of the team at the Kaiser-Permanente Medical Care Program in Northern California shows the two-step nature of diagnosing the problem. The Kaiser-Permanente team had identified an apparently simple procedural problem that was affecting the quality of care in a variety of ways: When a patient appeared for a follow-up appointment at one of the Kaiser-Permanente medical office buildings subsequent to receiving hospital emergency treatment, the patient's emergency records were frequently unavailable. As we saw in Chapter Four, the group wisely chose to narrow the scope of its inquiry, in the words of its carefully articulated problem statement, to "the timely transfer of appropriate emergency department clinical data from the Santa Clara Medical Center to its two satellites, [the medical office buildings in] Milpitas and Sunnyvale."

The Kaiser-Permanente team began to tackle the problem by *examining the existing process* for transferring emergency treatment data to the medical office building at which the patient received subsequent treatment. To accomplish this, they used a basic tool of quality improvement, the *process flow diagram*. (For a complete discussion, see Resource B.) The process flow diagram is useful in many ways. First, it converts what may seem like a vague, complicated collection of events into a graphic display of a clear, visible series of steps. More important, the very exercise of creating the flow diagram forces the group to see a "problem" in terms of a process, and it highlights the various inefficiencies and obstructions that may lie within it. The flow diagram helps people recognize that, underlying even the most haphazard, unplanned collection of events in work, there *is* a process. A process does not have to be efficient, or well conceived, or even *intentional;* it is merely the way work is done. Forcing oneself to commit to paper a description of the existing process—becoming aware of it *as* a process—is an important step toward improving quality.

The act of creating the process flow diagram was particularly revealing in the Kaiser-Permanente case. Based on their collective perceptions of what happened to a patient's medical record from the time of entering the emergency room to the follow-up office visit, the members of the Kaiser-Permanente team created the flow diagram displayed in Exhibit 6.1.

The diagram shows the following steps: When the patient arrives in the emergency room, the receptionist asks for his or her "home facility." Depending on the answer—MIL (Milpitas), SUN (Sunnyvale), or SCL (Santa Clara)—the receptionist tags

Exhibit 6.1. Initial Description of Record Transfer Process: A Process Flow Diagram (Kaiser-Permanente Project).

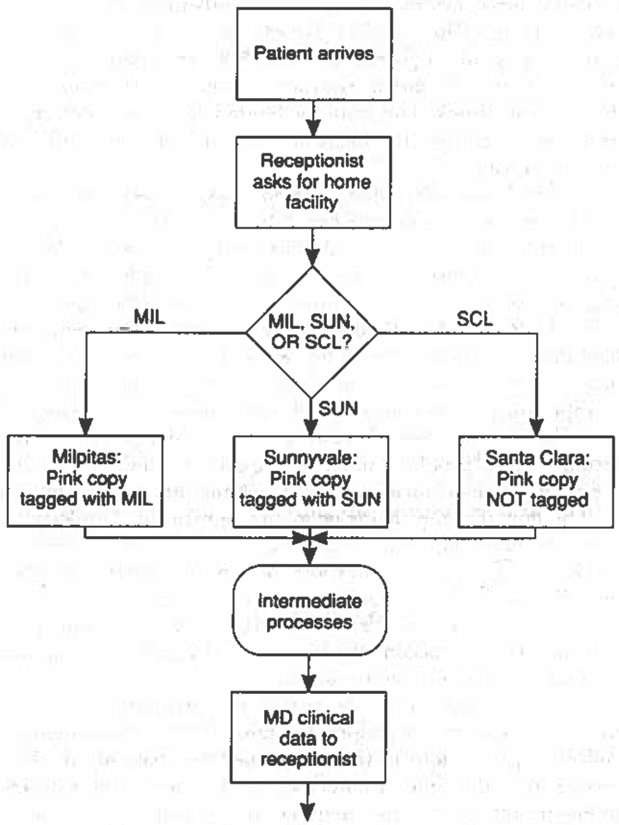

Exhibit 6.1. Initial Description of Record Transfer Process: A Process Flow Diagram (Kaiser-Permanente Project), Cont'd.

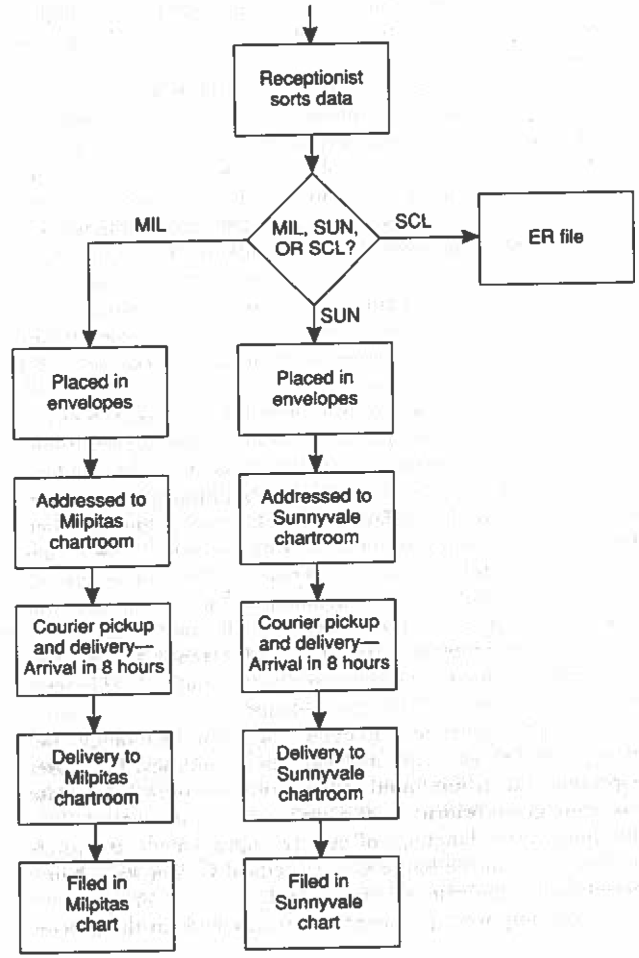

the pink copy of the patient's record, indicating the home facility. After the patient's emergency room treatment, clinical data are given to the receptionist, sorted, and delivered through a series of steps to the chart room of the facility at which the patient receives follow-up care.

As the team began its analysis of the process flow diagram to identify specific points at which the process could go wrong, problems at various steps immediately became apparent. For example, the group realized that at the apparently simple step, "Receptionist asks for home facility," a number of mistakes might occur, any one of which would derail the process from the start. As the group's report explains: "When the emergency room was very busy, the receptionists did not have time to get the information. Patients under stress, or because of language barriers, gave erroneous answers. At least one receptionist expected patients to volunteer the information because a sign on the wall asked that they do so."

The group found that, as is often the case, *when the process was under stress*—whether the stress of a busy emergency room, the stress of a patient in pain, or the stress of a language barrier—*it failed to work as it was designed.* Obtaining accurate information about the home facility might be easy enough when the emergency room was not busy, with patients who were not distracted and who could speak and read English; but the process was meant to work for *all* patients at *all* times, and as it was presently functioning, it was clearly inadequate to do so.

Another problem surfaced at a later step in the process: the sorting of all copies of emergency visit data. At this step, copies were distributed to several in-house departments—including pediatrics, obstetrics and gynecology, ophthalmology, and allergy—as well as to the medical office buildings. On closer inspection, the group found that distribution of in-house data was being given priority over distribution to the medical office buildings (in the language of quality improvement, the internal customers of the Santa Clara Medical Center were being served before the external customers).

Not only were problems at various steps in the process

immediately apparent; the team members were surprised to find that the *actual* process was in fact far more complicated than they had thought it to be. As they came to a clearer and clearer understanding of the process, the group put the original flow diagram through a series of revisions. Exhibit 6.2 is a later version of the same process represented in Exhibit 6.1. It indicates, for example, that (1) the receptionist may or may not ask for home facility and shows how the process proceeds in each case. Next, it reflects that (2) even if the receptionist does ask for home facility, the patient may or may not provide an answer, and indicates what happens in each case; later, (3) it more accurately reflects the path information takes from the emergency room to the site of follow-up care.

Not only did drawing the flow diagram help the team discover the actual process, it also made them realize in practice what they had heard of in theory: *the importance of customer-supplier relationships.* In the case of proper transfer of patient records, that meant recognizing that at each step along the way, there is a "customer" and a "supplier" and that the job of the supplier is to pass on to the customer *what,* and *only what,* he or she needs. The team reported: "The analytic process revealed unnecessary delivery of information to some departments and satellites, and inadequate delivery to others. During this process of revision, the actual information needs of the clinical departments and of the medical office buildings (all internal customers) were clarified and reclarified."

In the process flow diagram (Exhibit 6.2), for example, the receptionist is the *supplier,* whose job it is to determine accurately the site for follow-up care for every patient treated in the emergency room and to pass that information on to the next step along the way. The very exercise of creating the flow diagram helped clinical departments and medical office buildings clarify their roles and requirements as *customers;* what information did they need and what information was unnecessary? This, in turn, helped the emergency room understand its responsibility as *supplier* of information necessary for subsequent treatment of the patient.

Exhibit 6.2. Revised Description of Record Transfer Process: A Process Flow Diagram (Kaiser-Permanente Project).

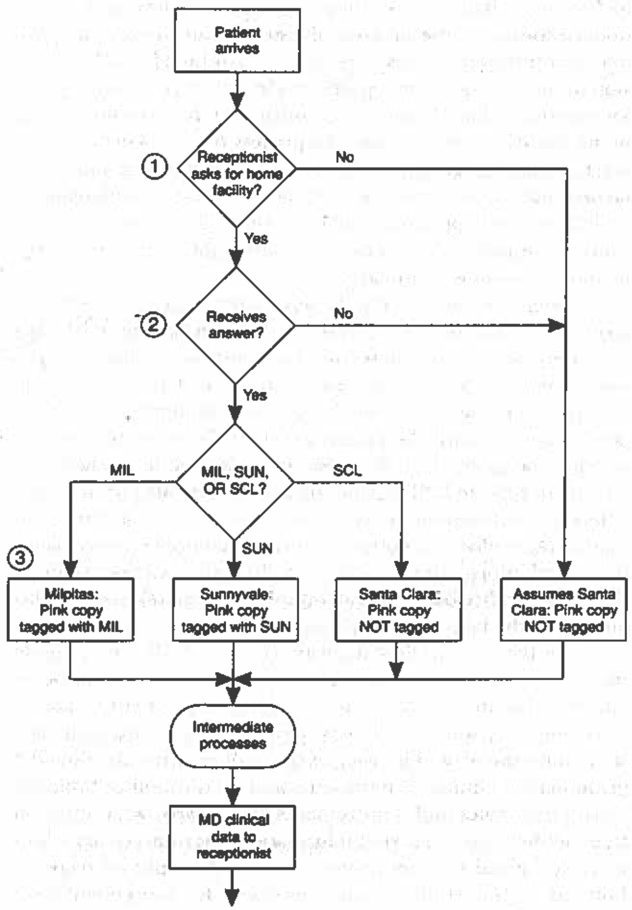

Exhibit 6.2. Revised Description of Record Transfer Process: A Process Flow Diagram (Kaiser-Permanente Project), Cont'd.

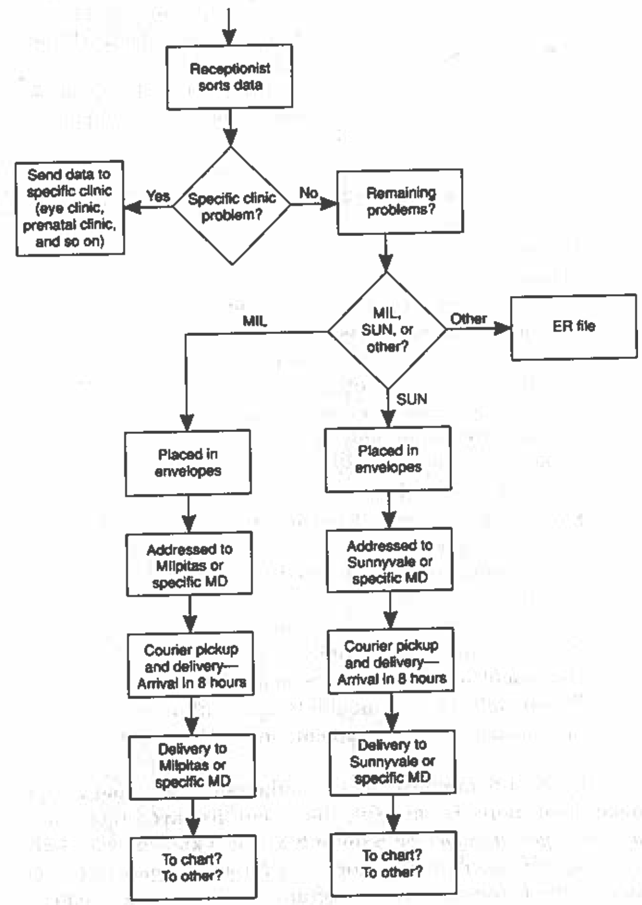

Once the Kaiser-Permanente team had created, analyzed, and then revised its process flow diagram, the next step was to generate hypotheses about the factors that inhibited the timely transfer of hospital data to the appropriate medical office building. Where and why might the process be failing?

The group began by brainstorming a list of hypotheses among themselves, and in short order produced the following list.

File lost in mailroom
Copy illegible
File sorted to wrong facility
White tag not used
No clarified procedure for labeling
Wrong information given
File delivered to wrong location
File sorted in wrong order
Patient not asked in effective way
File sent to wrong address
Duplicate copies received
Original record illegible
File sorted to wrong department
No copy made
Patient information not asked for
Duplicate copies made
Labeling information incomplete
Patient changes mind later
Document misplaced before being labeled
Patient refuses or is unable to give information
File mishandled by recipient, misfiled, or lost

In order to organize their hypotheses in a way that would make them more meaningful, the group displayed them in a *cause-and-effect diagram* (see Exhibit 6.3). Also known as an *Ishikawa diagram* (after its inventor) or a *fishbone diagram* (after its shape), the cause-and-effect diagram can be used to create a visual representation of the possible causes of a problem. (For a complete discussion, see Resource B.)

Exhibit 6.3. Cause-and-Effect Diagram for Timely Transfer of Emergency Room Records to Satellite Medical Office Building (Kaiser-Permanente Project).

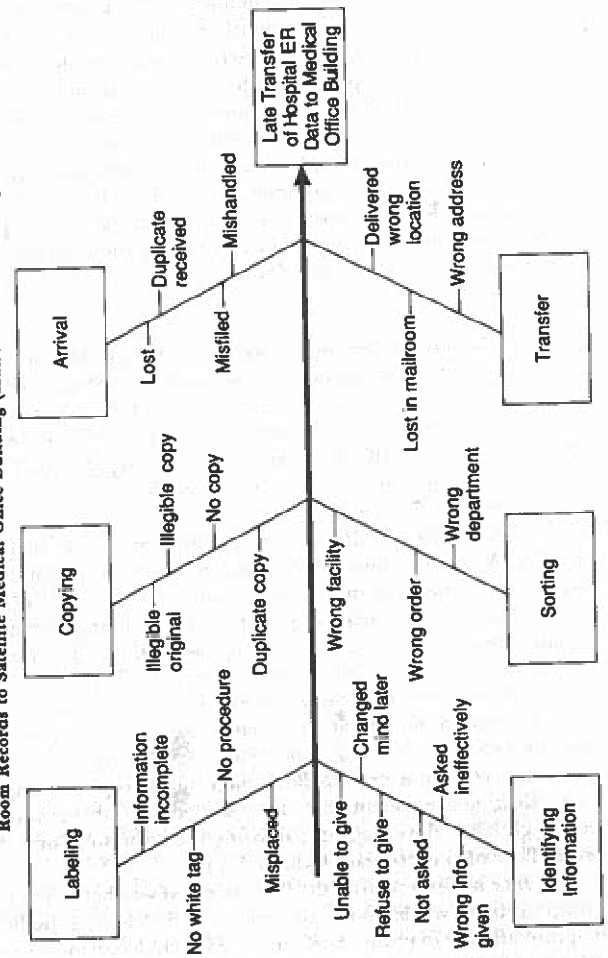

The diagram identifies six major possible points at which the process could go wrong—identifying information, sorting, transfer, labeling, copying, and arrival—and within each of these, a number of explanations for each particular error.

What had the Kaiser-Permanente team learned so far? First, the team had a clear and detailed understanding of the existing process, and second, they had analyzed that process and developed specific hypotheses about where it was flawed. At this point, the team was ready to move to the next step in the diagnostic journey: systematically testing their hypotheses to determine the root cause of the problem.

Massachusetts General Hospital. When the team from Massachusetts General Hospital chose to tackle its Medicare billing system, it began by calculating the "cost of poor quality," and determined that the problems of delayed and inaccurate billing were costing the hospital approximately $360,000 per year. In the language of quality improvement, they began by clearly "proving the need."

The first step in diagnosing the problem was to inspect all of the Medicare inpatient rejected claims for a one-month sample, collecting data on the most common reasons for rejection. At this stage, the team did not have to spend much energy documenting the symptoms, as was the case with Kaiser-Permanente; the reasons for rejection were clearly indicated by Medicare, and required only simple tabulation (see Exhibit 6.4).

In order to understand at what points, and for what reasons, the process was failing to provide acceptable bills, the team began by creating a process flow diagram of steps leading to the production of a patient's bill. They started with what is known as a "high-level" flow diagram, showing the basic steps and the broad flow of the process (Exhibit 6.5).

Once again, creating the flow diagram was helpful to the group in many ways: "Both the process of developing the flow diagram and the diagram itself enhanced each department's understanding of internal customer-supplier relationships—that is, how one area affects the next in the production of a 'clean bill.'

Exhibit 6.4. Common Reasons for Rejection, 10/29/87–11/20/87 (Massachusetts General Hospital Project).

Reasons for Rejection	*No. of Claims*	*Percent of Claims*
Invalid/Missing HIC No.	32	20.0
Medicare Secondary Payer (MSP)	19	11.9
Excess Ancillary Lines	13	8.1
Covered/Noncovered Days	15	9.4
Invalid MD ID No.	10	6.2
Invalid Total Charges	9	5.6
ICD-9-CM Dx Code	8	5.0
LOA/No Match	7	4.4
FY Split	7	4.4
Other	40	25.0
TOTAL	160	100%

Total billed in this time period = 850

Source: Rejection notification, November 1987.

The flow diagram clearly identified areas where the process can falter and where one activity's output does not match the next activity's requirements."

Massachusetts General's project team was truly multidepartmental, encompassing representatives from the medical staff, patient care services (including admitting), medical records, utilization management, financial services, and management information services. For them, the very exercise of having to sit down together and draw up a process flow diagram for the production of a patient bill was a revelation: It was brought home to them that the production of a "clean bill" is a process that crosses departmental boundaries. At each step along the way, the person involved is both the *customer* of the previous step, receiving exactly the information needed in order to proceed, and the *supplier* to the next step, providing exactly what the next step requires. Of course, the group had been aware of the interdepartmental nature of producing a patient bill, but the discipline of creating the flow diagram, and the explicitness of its display of the process, moved the group from awareness to active insight.

On the basis of the data collected from the examination of rejected bills, the Massachusetts General Hospital team identified three major process flaws and translated them into three goals:

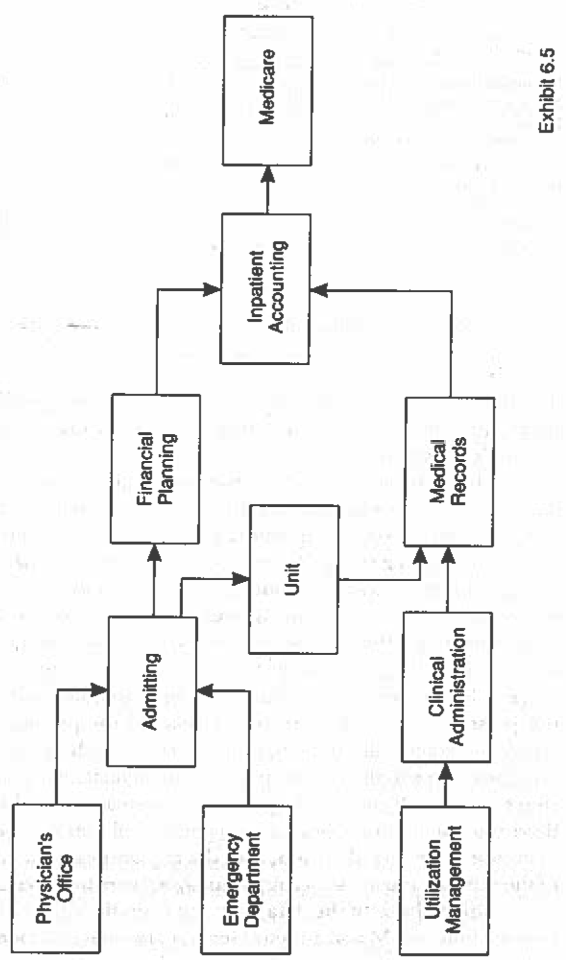

Exhibit 6.5. Medicare Inpatient Billing Process: High-Level Flow Diagram (Massachusetts General Hospital Project).

1. Improve the accuracy of HIC (health insurance claim) number information at admission.
2. Improve the accuracy of MSP (Medicare as secondary payer) information at admission.
3. Eliminate rejections due to "excess ancillary lines" on the claim.

The team then used the process flow diagram in yet another way: to identify the appropriate department to take responsibility for each of these issues. The team had begun with an elephant-sized problem—the high rate of rejected patient bills—and was able to divide it into manageable parts.

Generating and Testing Hypotheses: Two Case Studies

Once a team has used the process flow diagram to understand the existing process, the next step is to generate and test hypotheses about flaws in the process—in the jargon of quality management, "to identify and isolate the root cause" of the problem. The Evanston Hospital and Butterworth Hospital cases provide useful examples of hypothesis generation and testing.

Evanston Hospital. The team from Evanston Hospital began by attempting to find out what was causing delays in the operating room. They suspected that an increase in the volume of ambulatory surgery was the primary cause of scheduling delays, but the team wanted to conduct a systematic investigation to track down the root cause of the delay *before* proceeding with a cure.

Like the Kaiser-Permanente and Massachusetts General teams, the Evanston group began by examining the process as it was currently configured, creating a process flow diagram to follow the patient's progress from initial contact with the surgeon through preoperative preparation, surgery, discharge, and subsequent follow-up (Exhibit 6.6). The process flow diagram was useful to the Evanston team not only in developing its awareness of an underlying process (as in the Kaiser-Permanente and

Exhibit 6.6. Ambulatory Surgery Patient Flow: A Process Flow Diagram (Evanston Hospital Project).

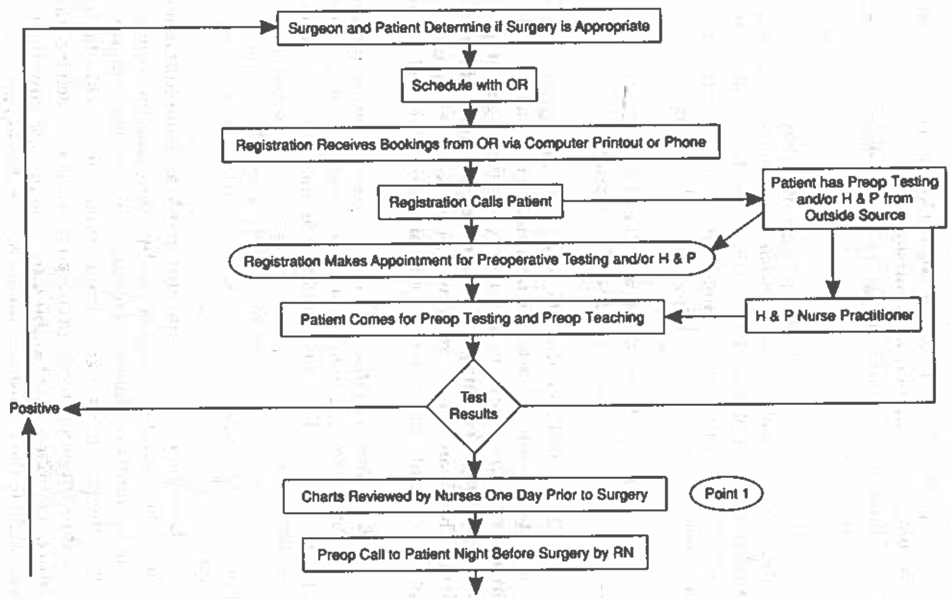

Exhibit 6.6. Ambulatory Surgery Patient Flow: A Process Flow Diagram (Evanston Hospital Project), Cont'd.

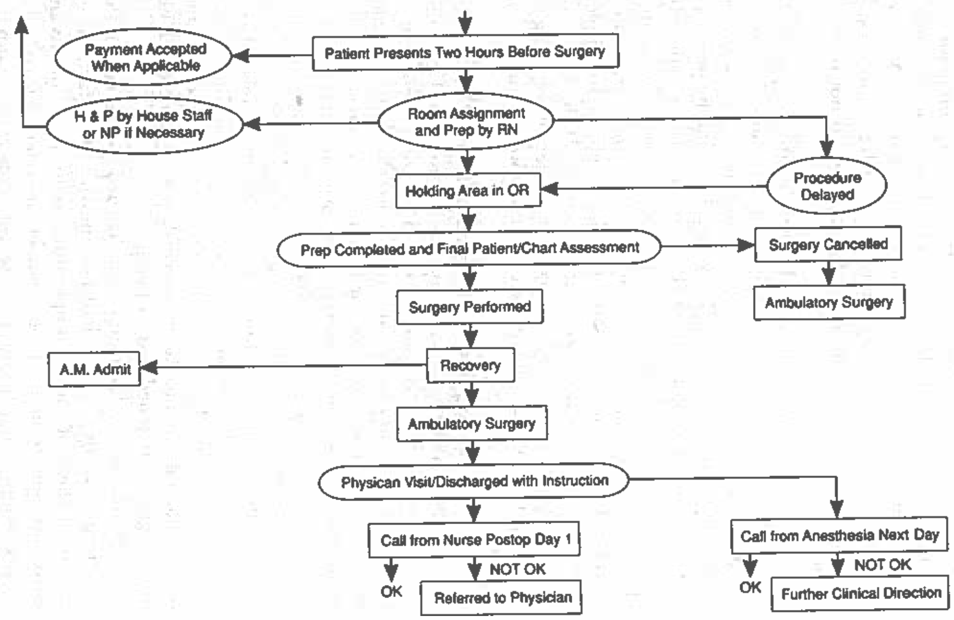

Massachusetts General cases), but *in guiding its development of hypotheses and collection of data.* First, the flow diagram made clear the major steps in a multistep process. The group was able to identify five key points as possible locations of delay. Data were collected for one month on the duration of each key process step; the results proved the *prevalence* of delay but did not reveal its *cause.*

Second, the flow diagram helped the group hypothesize where the process was most likely to malfunction. Examining the diagram, the team decided by consensus to review the process of data flow prior to surgery. They surmised that the preoperative stage was a likely point for stress, in that it was vulnerable to flaws in the process because of the many information sources that converged at that one point. Further, they ingeniously hypothesized that a prime cause of delay might be incomplete charts for the very first cases each morning. If the first cases of the day were delayed while preoperative data were being assembled, then all subsequent cases that day would similarly be delayed.

To test the hypothesis, the team collected data at point 1 on the process flow diagram during the month of December 1987, assessing both the timeliness and the completeness of information necessary to initiate surgery for the first cases each morning. Data collection took the form of medical record audits, conducted at noon on the working day prior to surgery. Their findings: "The data indicated that 16 cases [12 percent] were delayed for an average of 31 minutes. The key items missing that resulted in delay included the history and physical, lab results, and the electrocardiogram (EKG)." Their findings not only confirmed the hypothesis but indicated a course of future action: making sure the patient's records were complete well before surgery was scheduled to begin.

What had the Evanston team learned by the end of the "diagnostic journey" that it had not known at the start? In its search for the root causes of operating room delay, the Evanston team went through a process sometimes referred to as "peeling the onion." They began with the fact of operating room delays, and asked, "Which delays?" The answer: Preoperative.

They asked, "Why?" The answer: Data not in chart. They asked, "Which data?" The answer: History and physical, lab results, electrocardiogram. The next logical question, of course: "Why are those data missing?" With each successive "why," the team came closer and closer to an understanding of exactly at what points, and for what reasons, the process failed to function efficiently or at all. (One Japanese quality improvement expert specifically advises "asking *why* five times" at points like this in the diagnostic journey as a trick for "peeling the onion" of causes.)

It is especially interesting to notice that in Evanston's diagnostic journey the causal hypothesis with which they began when first defining the problem—that delays in the operating room were due to increased *volume* of ambulatory surgery—was replaced, as they accumulated facts and gained knowledge, with documented proof of a faulty *process.* As was noted in Chapter Four, people in organizations often believe that the solution to a problem lies in more staff, or more space, or more money. Evanston's employees could have protested, as have many other hospitals, that the operating room simply could not cope with increasing demand—that more operating suites, or more surgeons, were required to solve the problem. Indeed, the Evanston report stated clearly how strongly they believed at the start that ambulatory surgery volume was the cause of surgical delays.

However, careful investigation of the system as it really functioned indicated that the problem was not one of demand, but a flaw in the flow of work—*a flaw in the process*—that was responsible for significant operating room delays. That process could be improved substantially without either more staff or more money. In fact, adding staff or facilities *would not have solved Evanston's problem;* it would only have patched it over for a time. Until the process problem was addressed and the process redesigned to ensure that a patient's information would be complete before the scheduled surgery, delays would persist. In short, the Evanston group would have been treating the patient for the wrong disease.

Butterworth Hospital. Like Evanston, Butterworth Hospital chose to tackle the problem of an apparent inability to meet the demand for services — in this case, the services of the respiratory care department. The team began in the usual way: by examining the processes by which respiratory care was ordered in the hospital.

As in the case of Kaiser-Permanente's procedure for ensuring that patients' hospital records become a part of their regular medical files, the exercise of creating the process flow diagram was a revelation. As we saw in Chapter Five, "The team noticed a significant difference between how this process was working and how it was designed to work."

Like other effective teams, the Butterworth group began its diagnostic work by trying to describe the process underlying the problem it chose. Because of the multidepartmental nature of the Butterworth team, however, the members were able to recognize promptly that in fact there was not a single, consistent process for ordering respiratory care services but rather *multiple processes*. In health care, because "processes" often are not consciously designed to begin with, it is not unusual to find that they are inconsistent across units and across time. The team realized that this very multiplicity was a source of flaw.

In order to generate hypotheses about *additional* sources of flaw, the Butterworth team used several techniques to gather information about possible causes of the respiratory care department's inability to meet the demand for its services. Brainstorming within the team resulted in a long list of possible causes of process failure, which the group proceeded to organize in eight major categories of factors in a cause-and-effect diagram (Exhibit 6.7).

To test their hypotheses, the group then conducted a survey of the entire respiratory care department, asking members to respond to the list of the eight major factors identified in the cause-and-effect diagram. The results of the survey were displayed in the form of a *Pareto diagram*, which sorts out and presents variables in order of importance. (For a complete discussion, see Resource B.) In the Pareto diagram, six possible causes inhibiting the provision of therapy stood out as the most frequently mentioned (Exhibit 6.8). Of the six major factors

Exhibit 6.7. Reasons for Missed Respiratory Therapy: A Cause-and-Effect Diagram (Butterworth Hospital Project).

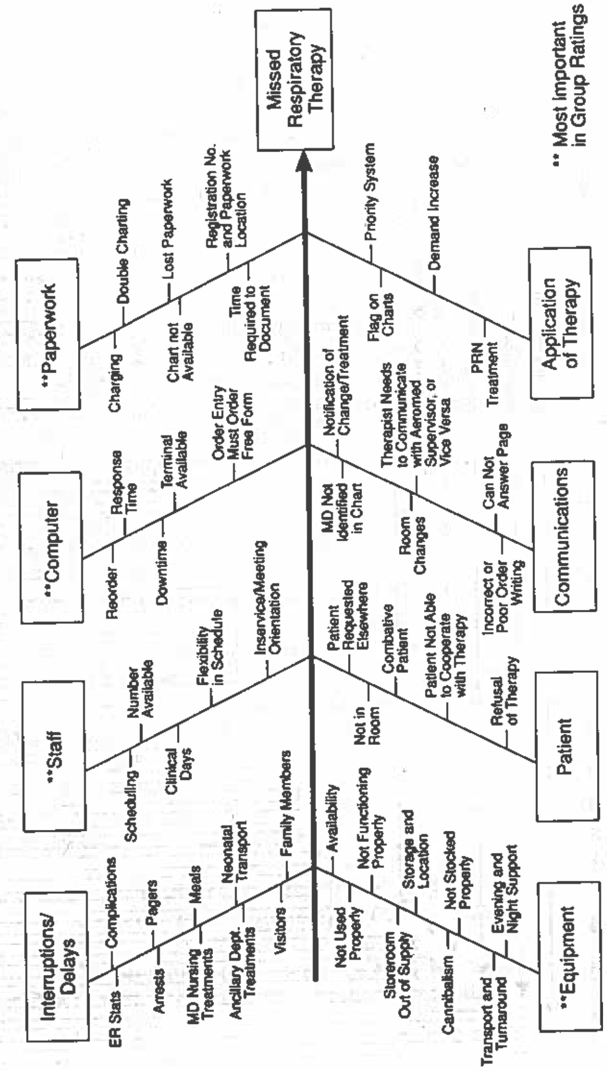

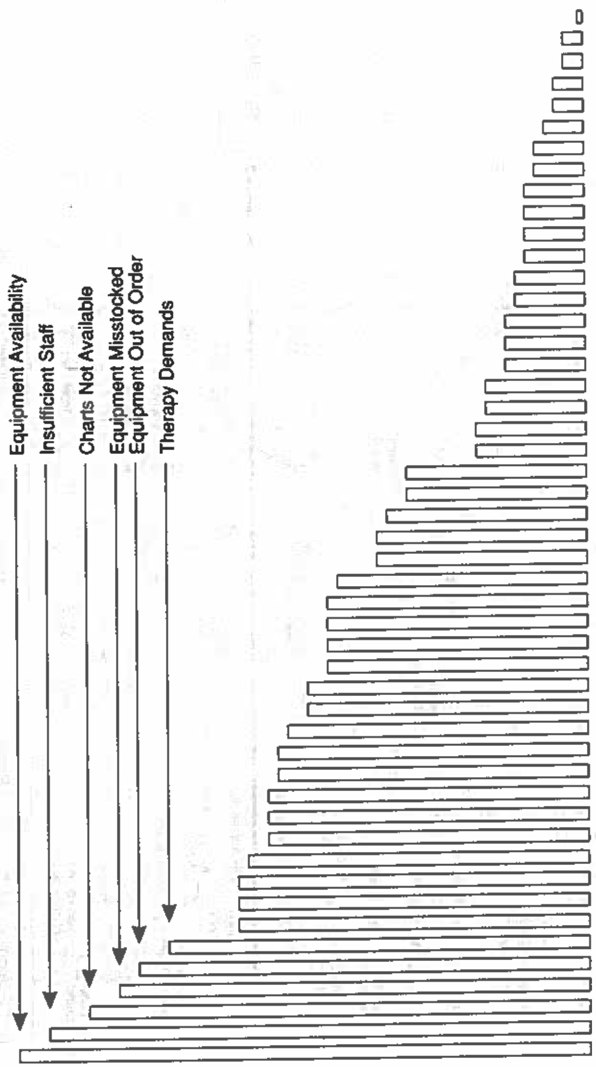

Exhibit 6.8. Reasons for Missed Respiratory Therapy: A Pareto Diagram of Survey Results (Butterworth Hospital Project).

in the Pareto diagram, the group noticed that three were related to equipment: "equipment availability," "equipment misstocked," and "equipment out of order." Consequently, the team *localized* its efforts; equipment was chosen as the issue to attack first. One final survey of the respiratory care department identified three specific equipment problems to be corrected: flowmeter unavailability, oxygen analyzer downtime, and oximeter unavailability.

The progress of the Butterworth team through its diagnosis is, like Evanston's, a classic example of "peeling the onion." Initially, the team knew only that the respiratory care department was not meeting the demand. By the end, they had a clear and specific agenda: solving three specific equipment problems. How had they gotten from here to there? They had followed a very useful pattern in the diagnosis of a problem: the alternating use of *divergent* and *convergent* thinking—alternately accumulating as many ideas as possible and then narrowing them down to the "vital few." They began by brainstorming lots of possibilities (divergent thinking) and then arranging them under eight major sources of flaw (convergent thinking); next they surveyed the entire department for priorities (divergent thinking) and then arranged the survey responses in a Pareto diagram that indicated six major factors (convergent thinking), three of which had to do with a single problem: equipment (convergent thinking).

The balanced use of divergent and convergent thinking is a valuable asset in problem solving. Yet, it tends not to come naturally to most people. More commonly, individuals tend to use only one mode—either convergent all of the time ("Let's make a *decision*, for Pete's sake!") or constantly divergent ("Have we really considered *all* of the options?"). Balancing the two modes of thought is a skill to be learned and honed in quality improvement.

Collecting, Displaying, and Understanding Data: Two Case Studies

In the preceding four cases, we have witnessed groups beginning to discover problems and think about them in a new way as they embarked on the diagnostic journey. Diagnosis begins with seeing that a process exists; carefully mapping and remapping it; then using that map to guide the development of hypotheses as to where the process malfunctions, and why. These early stages of the journey include sharing ideas and perceptions, generating hypotheses, and generally casting a wide net in thinking about processes and their possible flaws. As we have seen, the techniques associated with this stage—creating process flow diagrams and cause-and-effect diagrams, brainstorming, and surveying—are ways of soliciting and accumulating as many ideas as possible.

The next stage is quite different, both in purpose and technique. As groups move from *hypothesis generation* to *hypothesis testing,* they subject their ideas to rigorous scrutiny. Turning from reflection to action, from gathering ideas to winnowing them, they gather information—useful facts about problems and their causes—in the effort to hunt down the root cause of the problem. They move from the divergent enterprise of generating hypotheses to the convergent one of testing them.

Children's Hospital. The Children's Hospital project provides a telling example of the use of data to understand problems and test hypotheses. The Children's team chose a problem in which time was of critical importance to the quality of care: "As neonatal and pediatric intensive care units developed in tertiary care pediatric hospitals, it became clear that the time associated with transporting critically ill infants and children was an important factor in contributing to the morbidity and mortality of referred patients." In order to determine exactly where in the process of transporting children the most severe delays were occurring, the Children's team started by collecting data. First, they suspected that significant delays were occurring at Children's, *before* the ambulance even left the hospital. In order to test their hypothesis, they designed a time study measuring "elapsed time"—defined as the difference between the time a call requesting emergency pediatric transport was received and the time at which the ambulance was dispatched from Children's Hospital. Data indicated that, "in about 95% of instances,

it took at least 35 minutes for the team to leave Children's Hospital." Their hypothesis was indeed correct.

The actual graphic device the Children's team used was a type of *histogram*—essentially, a bar chart of frequency counts. Studying the data in a histogram, despite its apparent simplicity, can be a powerful way to generate ideas about localizing process variations. For example, if a histogram of waiting times shows two "humps," it is natural to ask why. The data may be telling us that there are *two* separate processes at work, each with its own characteristic waiting time. (For a complete discussion of histograms, see Resource B.)

Next, they wanted to know what proportion of the total delay in transporting children was due to "elapsed time." In order to find out, they did a study comparing "elapsed time" to the total amount of time from the initial call requesting transport to the arrival of the ambulance at the requesting hospital. Results indicated that "elapsed time" was responsible for between 35 and 90 percent of the total time, with a median proportion of 59 percent. *Data collection and analysis localized the problem* for the Children's team: Most of the delay was occurring right at Children's Hospital, before the ambulance pulled away.

The next step was to "peel the onion" and pinpoint where the in-house delay was occurring and why. In order to do so, the team refocused its examination of the data, analyzing the time delays to see if they could discover any correlation between excessive delays and certain work shifts or days of the week. In terms of quality management theory, they were looking for *special causes of variation*—that is, factors that are not intrinsic features of the process but are due to special, identifiable circumstances.

Every process has some degree of variation; in this case, for example, elapsed time might be forty minutes on one day and sixty-three minutes on another. But if, for instance, data analysis showed that every Sunday morning elapsed time was extremely high (technically speaking, usually defined as more than three standard deviations from the statistically expected value), the team would be alerted to a special cause of variation such as a smaller staff or fewer ambulances on Sunday mornings. Their next step would be to remedy that special problem, before going on to address common causes that affect the variation in the process as a whole. It is of critical importance to determine whether variation is due to a special or common cause, because the remedy for each is different. Special cause variation can be controlled by identifying and removing the causes. Common cause variation, being an inherent property of the process itself, can be reduced only by changing the process.

Using the data it had already collected, the group began to search for special causes by looking for correlations between excessive delays and certain work shifts or certain days of the week. The analysis of the data uncovered no such correlations. The group then decided to display the data differently, in a plot of elapsed times in order of occurrence (Exhibit 6.9). That plot showed extremely wide variations in elapsed times, although there did not seem to be a single explanation for the variability. The plot showed clearly that "the median elapsed time was not only too long, but was often out of statistical control." "Out of statistical control" means, by definition, that special causes of variation did exist, or, put otherwise, that *there was no stable transport process at all.*

The Children's team had used data to answer a series of questions: Where was the delay occurring? What proportion of the delay was attributable to "elapsed time"? Were there any special causes of variation in the process? The team was then able to design a strategy for remedy on the basis of the answers to those questions.

Park Nicollet Medical Center. Park Nicollet Medical Center also chose to work on a problem at the core of which was *time*—patient dissatisfaction with telephone access—and to use data collection and analysis to isolate the root cause of the problem. As was the case in Children's Hospital's investigation of reasons for delay in the transport of critically ill infants, the Park Nicollet team attempted to localize the problem by collecting time data on various points in the process flow diagram. Unlike the Children's Hospital team, the Park Nicollet team was

Exhibit 6.9. Elapsed In-House Transport Times in Order of Occurrence: A Control Chart (Children's Hospital Project).

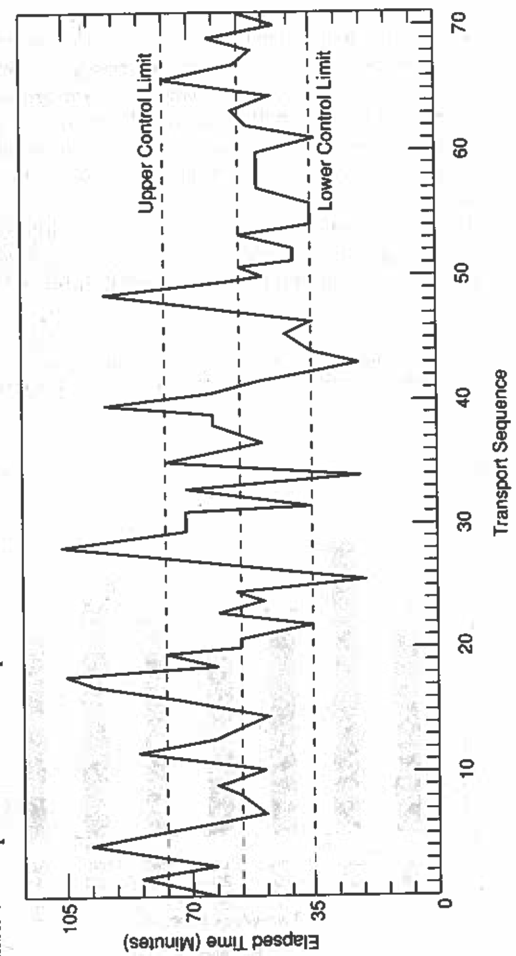

able to make use of existing data: The telephone system already in place tracked data on call volumes, calls abandoned, and average time to answer a call.

Looking for particular times and places when telephone access was especially difficult, the team reformatted existing data into charts tracking "average time to answer" over time. At first, the data were of little use in illuminating the problem: "Initially the charts were developed and displayed for two-week blocks, which resulted in a 'blurring' of a good deal of useful detail." (See Exhibit 6.10.)

However, when the data were displayed differently, to compare "average time to answer" during *one given hour* over a succession of days, the data came to life (Exhibit 6.11).

Exhibit 6.10. Average Time for Receptionists to Answer Telephones (Two-Week Averages) (Park Nicollet Project).

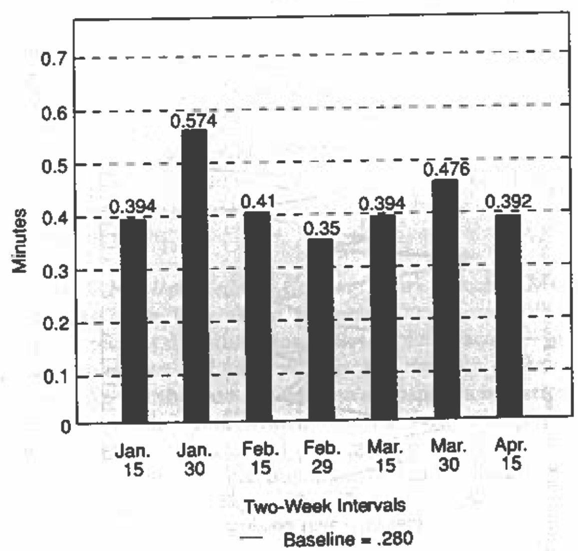

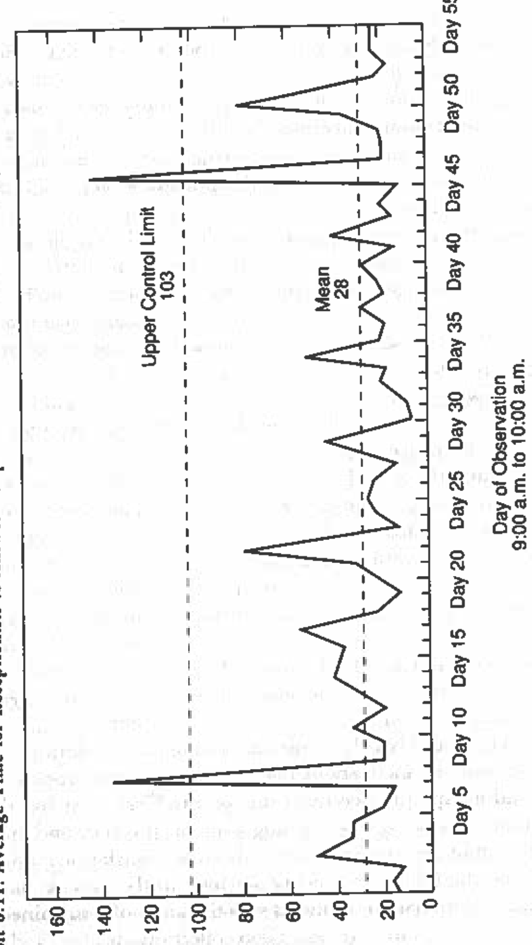

Exhibit 6.11. Average Time for Receptionists to Answer Telephones Between 9 A.M. and 10 A.M. (Park Nicollet Project).

Exhibit 6.11 indicated two points at which the difficulty of telephone access had spikes; at these points, the process was out of statistical control, indicating that special causes of variation were likely at work. When the group investigated these days and times, they discovered that they corresponded *exactly* to the days and hours during which two receptionists were called away to important training meetings. In this case, the data had brought to light a *special cause* of variation in the process—that is, a problem that is not associated with the process in general but with a particular time, place, or reason outside of the process itself. Further, the discovery pointed to a relatively simple solution: Park Nicollet changed the training policy immediately to avoid taking receptionists away from their positions during business hours. The pattern of response times for the medical information (MI) nurses showed high variation but no specific episodes of special cause. (See Exhibit 6.12.)

Having eliminated a special cause of variation in receptionist response times, the team had cleared the way for reexamining the process as it routinely functioned. On reexamining the data, the team found that, over a specific ten-day period, while the average time for the *receptionist* to answer the phone was fifteen seconds, the average time for the *medical information nurse* to answer was between sixty and ninety seconds. *The team had discovered the "capability" of the medical information nurse telephone answering process:* The medical information nurses were taking an average of four to six times as long to answer the telephones as the receptionists. On the basis of these data, the team chose to focus its efforts on the medical information nurse problem—in other words, to improve the process capability.

The Park Nicollet team was able to use data to bring to light two key facts about the telephone access process, facts that had escaped notice until this point: First, that receptionists' training sessions were causing a major, if special and limited, disruption in the process (a special cause of variation); and second, that the phone answering capacity of the process used by the medical information nurses was the locus of a sustained disruption in the overall process. Data collection, display, and analysis enabled the team to make a specific and informed diagnosis of the problem.

Lessons from the Diagnostic Journey

As we have seen, different diagnostic approaches work best for different types of problems. In *all* cases, however, we have observed the usefulness of the process flow diagram. Flow diagrams literally force people to see things in a new way, by arranging an amorphous collection of events into a visible map of interrelated steps. These diagrams emphasize graphically the interdependence of events that may have seemed unrelated and of departments that may have thought they functioned in isolation. One person's or one department's perception of the steps in the process may be very different from the next person's or department's. Having to commit one's understanding of the process to paper enables a group to discover those differing perceptions—often a crucial first step in solving the problem. In the case of the NDP teams, the flow diagrams made apparent, as no amount of explanation could, the nature and importance of the customer-supplier relationships.

We have seen that the process of diagnosis is most powerful when it involves a rhythm of *generating hypotheses* (whether by brainstorming or surveys)—and thereby broadening the scope of inquiry; and then *testing those hypotheses* (whether by surveys or data collection)—and thereby narrowing the scope of inquiry. Through this process, a group "peels the onion," getting closer and closer to understanding precisely where, and why, the process fails.

The collection of data at the stage of diagnosing the problem is a way of testing hypotheses to pinpoint the root cause of the problem. The display of data helps teams *see* what the data mean—to understand the lessons the data hold. Cause-and-effect diagrams are a useful way of picturing the factors contributing to a given problem. For example, Kaiser-Permanente's sixty possible reasons why a patient's emergency room data fail to find a place in the patient's regular file are difficult to draw conclusions from, when presented as a list; however, when the same sixty causes are rearranged into a cause-and-effect diagram, they provide a handy way of visualizing both the major causes of the problem and the specific contributing factors to each of those major causes. For another example, when telephone access

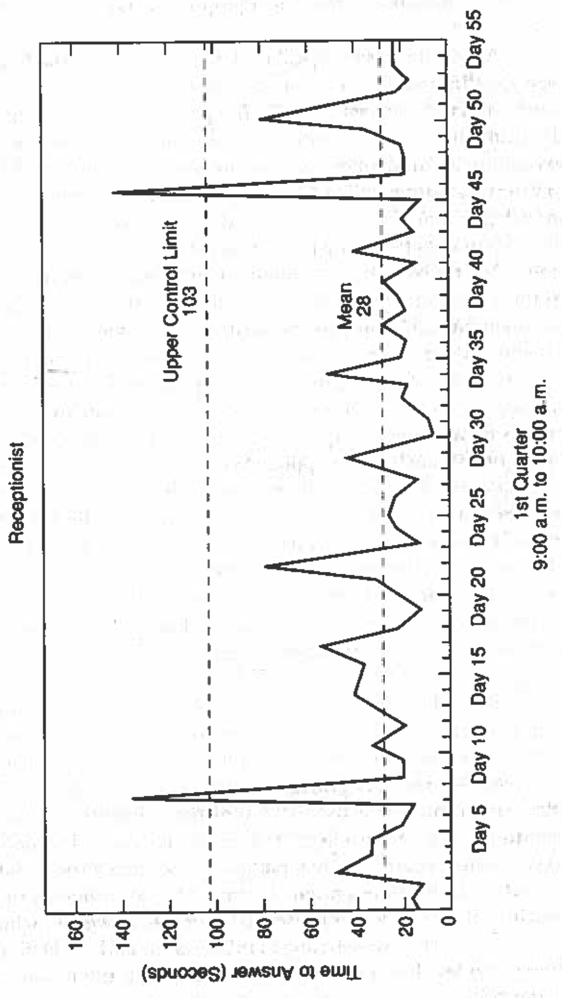

Exhibit 6.12. Average Time to Answer Telephone Between 9:00 A.M. and 10:00 A.M. (Park Nicollet Project).

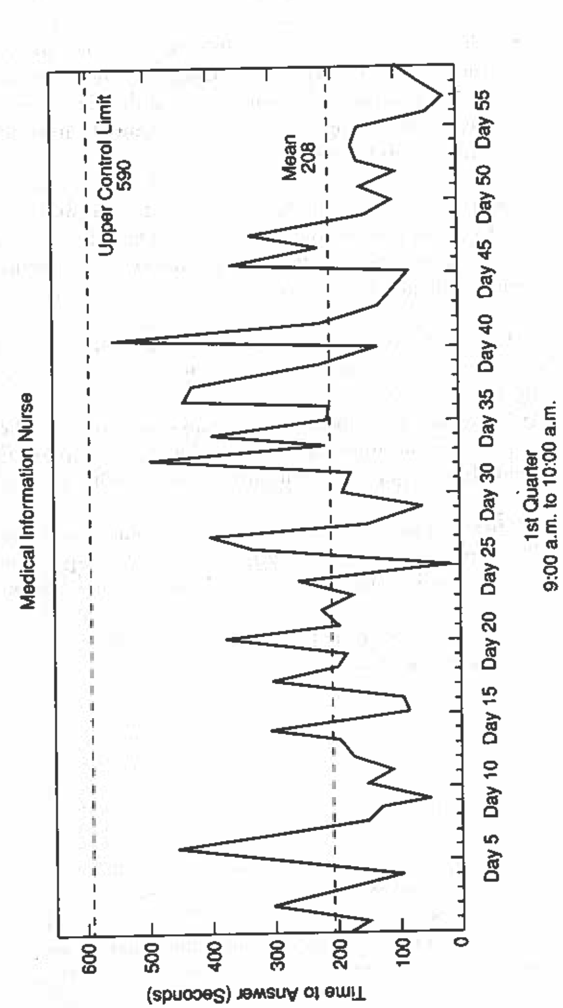

Exhibit 6.12. Average Time to Answer Telephone Between 9:00 A.M. and 10:00 A.M. (Park Nicollet Project), Cont'd.

data were displayed in terms of two-week intervals, they told the Park Nicollet team nothing; when the same data were displayed differently, in terms of hour-long intervals, the lesson almost jumped off the page. *Data suddenly became information.* Perhaps the most impressive theme we see in the diagnostic journey is the element of surprise—as groups come to new understandings of old problems:

- At Kaiser-Permanente: So *that's* why the hospital records aren't getting into the patients' regular files—because when the emergency room gets too busy, the receptionist doesn't have time to get the correct information, or the patient can't give it.
- At Evanston: So *that's* why the operating room gets backed up every day—because the charts for the first surgeries of the day aren't ready in time.
- At Butterworth: So *that's* why we can't seem to meet the demand for respiratory care—not because the demand is too great, but because the equipment isn't readily available.

These "aha!"s are moments of understanding *where* the problem is really occurring and *why*. Further, they supply needed information for the creation and testing of remedies—not throwing more time, money, or staff at an intractable problem, but remedying processes based on a careful and systematic diagnosis of the problem.

7 Implementing Successful Remedies

Once a quality improvement team has made a complete diagnosis of the problem and identified the root causes of flaws in the process, they are ready to begin the "remedial journey." The remedial journey consists of three well-defined steps:

1. Developing the remedy
2. Implementing and testing the remedy
3. Dealing with resistance to change

Developing the Remedy

The first step in developing an effective remedy involves considering a variety of alternative solutions. As was the case at the diagnosis stage, the different perspectives of team members are an especially valuable resource at this stage; since each team member has a unique understanding of the process, each may suggest different solutions.

In choosing among alternative solutions, a team must anticipate and weigh the *costs* associated with each, the length of *time* required to implement each, and the means of *evaluating the effectiveness* of each solution. The quality improvement project team, rather than one department, is best suited to make these decisions.