

Ten Years of Staffing Assessments

Andy Brazier

Christmas 2011



www.abrisk.co.uk

1 INTRODUCTION

In 2001 a document was published that introduced a method of assessing staffing arrangements for process operations in the chemical and allied industries. It was developed by consultancy company Entec UK Ltd. on behalf of the Health and Safety Executive (HSE) and published as Contract Research Report (CRR) 348/2001.

The method was developed because of concerns about the effect of staffing reductions and reorganisations on a company's ability to control the risks of major accident hazards. It was not intended to calculate a minimum or optimum number of staff, but instead to flag up when staffing levels were too low or the systems and organisation in place were not sufficient to ensure the people present could control risks effectively. Before its publication, there was very little available in the public domain to guide assessments of these critical but complex issues.

1.1 Method overview

The method is presented in two distinct parts as follows:

1. Physical assessment – a set of decision trees used to consider whether staffing arrangements are sufficient to handle credible scenarios and incidents;
2. Ladder assessment – a method of benchmarking individual and organisational factors that affect staffing arrangements.

The method can be used at any time. However, it is arguably of most value at times of change when an assessment can compare current staffing arrangements with the expected outcome following the planned change.

1.2 My experience of the method

I was not involved in developing the method. However, I joined Entec in early 2002 as principal human factors consultant and was given the task of promoting use of the method with clients. It is something I have continued having become a self-employed consultant in 2005.

I've lost count, but over the last 10 years I have been involved in at least 30 staffing assessment projects for more than 15 different clients. Also, even where the method is not formally used I often refer to elements of it as guidance for my other human factors and risk consultancy work.

In 2004 I lead the development of a user guide for the method on behalf of the Energy Institute (EI). The aim of this document was to provide best practice guidance on how to conduct assessments using the method based on practical experience.

1.3 Objective of this paper

It occurred to me that having spent 10 years using the method, it was a good time to stand back and reflect. In general, although I can point to some flaws in the method, I have found it to be a very good framework for assessing human and organisational factors. It prompts you to ask challenging questions and to be objective in your analysis. Also, I have found that the observations and recommendations I have made as a result of using the method have been very well received by my clients.

My main aim with this paper has not been to evaluate the method, but to summarise some of my key findings from using it. To do this I have reviewed every decision tree and ladder; and documented the typical arrangements I find at companies and the issues that have been raised by assessments. I believe that this approach gives a good summary of how human and organisational factors are being dealt with at the end of 2011 and hope it will be of wide interest, even to people who have no knowledge or experience of the method.

Something that has occurred to me whilst generating this paper is that what constituted best practice in 2001 is likely to now be out of date. Although I plan to propose some updates to the method separately in the future, I have summarised my thoughts about using the concept of high reliability organisations to define current best practice for the ladder assessments.

1.4 Useful references

The following references can all be downloaded for free from the internet. You may find them useful, particularly if you are not familiar with the method.

- The HSE report documenting the method can be downloaded from http://www.hse.gov.uk/research/crr_pdf/2001/crr01348.pdf
- The EI user guide can be downloaded from <http://www.energyinstpubs.org.uk/tfiles/1323709523/429.pdf>
- A presentation giving some background to the method can be downloaded from http://abrisk.co.uk/human_factors_course/11%20Staffing.pdf

Having completed this paper, I now plan to review the method and look for opportunities to improve. I will make my proposed update available at <http://www.abrisk.co.uk/papers-articles/298-ten-years-of-staffing-assessments>

2 PHYSICAL ASSESSMENTS

The physical assessment considers the staffing arrangements in place to deal with hazardous scenarios. The HSE report provides a set of eight decision trees that help to determine whether there are likely to be enough people, in the right place at the right time to detect, diagnose and respond to events so that are not able to escalate into major accidents. The EI report includes guidance on which scenarios should be assessed and how an assessment should be carried out. Typical scenarios include toxic releases, process fires, utility failures and major disturbances to the system (e.g. plant trip).

Each decision tree is discussed below. I have summarised my experiences of using the trees, including the typical issues that are raised.

2.1 Decision tree 1 – Control room manning

All the trees in the physical assessment assume that the control room is the hub of the operation. In my experience this is an accurate reflection of operations at most sites dealing with major accident hazards. However, this does not mean the method is only concerned with control room manning; although that is the view held by some people. I have always used the method to consider how the whole team functions.

2.1.1 TYPICAL ARRANGEMENTS

Most of the assessments I have carried out have been at sites where the control room is described as being permanently manned. One or more control room operators are assigned to the role for a shift. There have been some cases where the control room operator has also been required to work in the field as well, meaning they sometime leave the control room to check equipment or carry out operations. However, in most cases the control consoles have been located alongside others so that there is normally someone in the control room the whole time.

2.1.2 ISSUES RAISED BY ASSESSMENTS

The initial answer to the question “is the control room continuously manned?” is usually “Yes.” The decision tree suggests this is an immediate pass. However, further questioning can often identify a number of issues that sometimes lead me to record a fail.

I have a major concern with the way many companies make arrangements for control room operators to take breaks during their shift. Many have no formal arrangements in place meaning that they sometimes need to ‘pop out for a few minutes’ to use the toilet or get something to eat. There are two problems with this:

1. A delay in detecting the early signs of a scenario developing because no one is present to observe a change or receive an alarm;
2. Control room operator fatigue and/or loss of alertness because they cannot take a break.

Some companies will insist on the control room being continuously manned, but do not have any ‘spare’ control room operators. This means the person covering is not competent in the role. This may be acceptable if they have been trained to recognise problems and have a reliable method of calling the control room operator back to the control room. But in many cases this is only handled informally.

My view is that the control room operator role is highly critical, and so any person in the role must be competent, fit and alert. This means formal arrangements should be in place to make sure control room operators can be covered for breaks. Also, I feel that control room operators should be made to take adequate breaks. Many choose not to, and eating meals at the control console is a very common occurrence.

One of my clients has arranged it so that the control room operator will only spend half a shift in the control room, with the other half spent in a field operator role. This is something that I have, in the past, considered to be a bad idea because it effectively doubles the number of handovers that take place, which we know can contribute to major accidents. However, the client is satisfied that the arrangement is beneficial, and that the concerns about handover have proven to be unfounded. As a result I would urge all companies to consider this arrangement, bearing in mind that merchant and military navies have operated 'watch' systems for many years, typically of a maximum duration of four hours.

2.2 Decision tree 2 – control room operator activities

This decision tree makes the point that just because someone is in a control room, it does not mean they are always sat in front of the control console where early warnings will be received that something is going wrong.

2.2.1 TYPICAL ARRANGEMENTS

It is fairly normal for a control room operator to have responsibility for a number of process plants at the same time. With modern equipment it is easy to incorporate all the controls the operator needs into a single console. However, in many cases the control room and the role of the operator have evolved over time. Therefore, control room operators often have to work at more than one console to monitor and control the plants they are responsible for.

Other things can take the control room operator away from their main control console. They include hard-wired panels (e.g. emergency shutdown, fire and gas detection), communication devices and 'standard' office equipment (e.g. personal computer, photo copier, fax machine). In some cases the control room operator may get involved in completely separate activities such as dealing with contractors and issuing permits-to-work, although this appears to occur less often than it may have in the past.

2.2.2 ISSUES RAISED BY ASSESSMENTS

The problems highlighted by this decision tree have less to do with the activities performed by the control room away from the console, and have more to do with poor management of change in the control room and poor design of the Human Machine Interface (HMI).

In too many control rooms it has been considered acceptable to simply add new equipment wherever there is a space to accommodate it. Little or no thought is given to how the operator is going to use it and how that is going to affect their ability to continue their existing monitoring and control activities. I have been in control rooms where the operator is completely surrounded (full 360°) by panels and screens. Clearly, if they are working on one, they cannot be monitoring the ones behind them.

Although surrounding a control room operator with equipment is clearly poor, I have to acknowledge that they usually cope. The explanation is that when something goes wrong process alarms will alert them. The problem with this is that alarm systems are generally pretty poor and they cannot be relied on to provide an early warning. This can be particularly serious when emergency alarms (e.g. fire, toxic gas) are only indicated on the control system HMI (i.e. there is no separate alarm panel) because operators may not always respond immediately if busy because they assume the alarm is a nuisance process alarm. I do acknowledge that companies have been doing a lot of work to reduce the known problems with nuisance alarms and floods of alarms, but I have yet to see much evidence that any have actually solved these problems.

Reliance on alarms could be reduced if the HMI provided control room operators with useful overview displays. These allow the operator, with a quick glance, to determine what is going on. Unfortunately many HMIs only provide detailed displays that require some considerable effort to determine status, which makes it difficult to identify deviations. Even though large screen displays are now quite common in control rooms, their use has not been considered

adequately. Operators continue to rely on alarms to warn them of problems, even though we know most alarms systems are flawed.

2.3 Decision tree 3 – distractions in the control room

This decision tree asks what the control room operator does in addition to monitoring and controlling the process. The challenge is whether these distractions could delay detection of the scenario.

2.3.1 TYPICAL ARRANGEMENTS

I would say every control room operator has to answer telephones, talk to other people in the control room, carry out some administrative duties (including completing their shift log) and deal with nuisance alarms. Whilst we would like to avoid distracting operators, we have to acknowledge that some of these are necessary and cannot be avoided; and the job would be very boring without some variation and contact with other people.

2.3.2 ISSUES RAISED BY ASSESSMENTS

It appears to me that control room operators are distracted less often than they were in the past. More thought is given to keeping unnecessary visitors out of the control room and new technology means other personnel can often access process data without having to visit or phone the control room. I am not saying control room operators do not have complaints, but I have not had to fail many assessments for this decision tree in recent years because operators do not feel the distractions in the control room affect their ability to detect a scenario.

2.4 Decision tree 4 – information for scenario diagnosis

This decision tree talks about “additional information” required to diagnose the scenario and initiate a response. It asks whether this information is accessible, correct and understandable; and whether there is any back-up that can assist the control room operator.

2.4.1 TYPICAL ARRANGEMENTS

Control room operators do have a key role in diagnosing a scenario and initiating a response, but they rarely act alone. They will invariably ask field operators to investigate on the plant and for a supervisor to assist. Therefore, diagnosis is usually a group activity.

Diagnosis is usually made from information obtained from:

- Process graphics;
- CCTV images;
- Reports from the field.

2.4.2 ISSUES RAISED BY ASSESSMENTS

The effort put into diagnosing an event is often determined by the effort it requires to shutdown and restart a plant. If it is easy to initiate a trip and then easy to restart, operators are not so concerned that they may misdiagnose an event. However, if tripping a plant is likely to cause problems, they will take their time to make sure their diagnosis is correct and try to work out a way of containing the situation that does not have such a big impact on the plant (e.g. controlled shutdown, partial trip). Whilst this is sensible, operators often have to make complex decisions of this nature in stressful situations, often with minimal support.

At many sites, the large number of alarms that occur when something goes wrong means that, whilst they are useful at indicating a problem, they are of little use when diagnosing the cause. Also, even though control room operators will interrogate process graphics to diagnose the event, many are set-up for ‘normal’ operation and of limited use when things are ‘abnormal.’

A consistent finding in the assessments is that very little procedural support is provided for diagnosing scenarios. The only documents readily available are operating procedures or instructions. These are usually too long and wordy to refer to within the time available, and most provide little information about diagnosis or the initial response to a scenario. Also, there is a trend to only provide procedures electronically, which often adds to the complexity and time required to access written information. Some of my clients have developed prompt cards to assist operators dealing with scenarios that are readily available and easy to use. However, these are the minority.

2.5 Decision tree 5 – communication during diagnosis and initial response

I have already alluded to my belief that diagnosis is usually a group activity. Therefore, communication between individuals is important. This decision tree asks about the preferred method of communication between the control room operator and the rest of their team; and whether there is a reliable backup.

2.5.1 TYPICAL ARRANGEMENTS

The most commonly used method of communication between members of the operating team is the plant radio. Some sites have Personal Address (PA) or Tannoy systems, and a small number have plant telephones. Often, the sounding of the site evacuation alarm will result in field operators returning to or contacting the control room, which is often a last resort method of contact if plant radios fail.

2.5.2 ISSUES RAISED BY ASSESSMENTS

When I first started carrying out staffing assessments the reliability of plant radios was a serious concern at most sites. I am pleased to say that I hear a lot less complaints with radios. I think improved battery technology made the most impact.

A lot of sites actually rely on radios because they do not have a back-up system. In these cases the reliability of the system has to be ensured. Usually, any problems are limited to single hand-held units and during a scenario something can be done, such as asking someone to take a spare radio to the field or to act as runner.

The introduction of mobile phones rated for use in hazardous areas has provided an opportunity to improve communications whilst introducing some potential risks. During normal situations, being able to contact someone directly and hold a one-to-one conversation can be very useful, especially between a Supervisor and a member of their team. However, in an emergency the inability of a mobile phone to broadcast messages to groups of people can be a problem. Also, emergency services have the option to shutdown networks if they believe an incident may be terrorist related or lead to an overload. There is a role for mobile phones, but they should only be seen as an addition to the plant radio system.

2.6 Decision tree 6 – executing a response

This decision tree has always been a bit of a mystery to me because it only applies when the control room operator performs all the actions in a scenario. This is very rare, with most scenarios requiring a team response. I have tended to use this tree to identify who is involved in response and how well they are able to work together.

2.6.1 TYPICAL ARRANGEMENTS

It is inevitably the operating team that deal with the early stages of any scenario. If their actions are effective there is no need for others to get involved. However, if a scenario cannot be controlled other people will usually be called in to assist. On larger sites, these

people can come from other departments. Managers can assist at more strategic levels and the emergency services have an obvious role when people are at risk.

2.6.2 ISSUES RAISED BY ASSESSMENTS

I often feel that companies are not sufficiently prepared to handle the range of scenarios that can occur on their site. Emergency procedures often present only a high level overview of activities to be performed. Emergency exercises are carried out too infrequently. Sometimes a lot of effort is put into one big event per year, which means very few scenarios are actually tested.

There is often a degree of naivety in how people think they will be able to respond to a scenario. For example they assume:

- The full team will be available – no one will be incapacitated in the early stages of the scenario;
- The 'normal' team will be present – no one will be on holiday or called in sick that day;
- Emergency services will arrive quickly and be immediately ready to respond;
- Everyone present will know what to do;
- The worst case scenario cannot happen.

Some of these issues are addressed in the Ladder Assessments (see the next section of this paper), but this decision tree can be a good opportunity to challenge some assumptions.

2.7 Decision tree 7 – communication during incident response

This decision tree can appear to be a duplicate of tree number 5. Whilst there are obvious overlaps, this tree is actually concerned with communication during the main response to the scenario rather than the initial diagnosis. The main difference is that other parties are likely to be involved.

2.7.1 TYPICAL ARRANGEMENTS

The plant radio remains the preferred method of communication throughout a scenario. Other parties rarely have access to the same radio systems and channels, so communication links have to be established using different methods.

2.7.2 ISSUES RAISED BY ASSESSMENTS

Although the same issues apply to the plant radios that are addressed in tree 5, additional issues can arise as a scenario develops. For example, if people are required to use additional PPE (e.g. chemical suits) using radios can be difficult. Also, the amount of traffic on the radio system can be large and cause confusion, especially if people who usually use different channels switch to the same 'emergency' channel.

When the site's management team get involved in a scenario it is clearly important that they know what is going on so that they are able to assist in an effective manner. The challenge here can be communicating information about the scenario in a way that minimises disruption to the team working at the sharp end. There can be a tendency for managers to visit the control room, which can cause distraction. Phoning the control room for an update can be even more distracting. Also, it has been known for managers to go to the scene of the incident so that they can see what is happening for themselves, but this compromise their safety and delays them fulfilling their emergency role.

It is very rare for emergency services to be able to communicate directly with the operating team because their radios are usually incompatible. Sometimes they are given a plant radio, but this can cause problems if they are not familiar with the site's communication protocols. An alternative solution is for a member of the site personnel to join the emergency services,

and this can be an important part of establishing a ‘forward control point.’ However, this person may have been performing a valuable role in the response, and have to be pulled away. At some sites it is the Supervisor who goes to meet emergency services. I feel this is a poor arrangement because it is usually better to put Supervisors in a more strategic role, typically ‘Incident Controller.’

2.8 Decision tree 8 – activities during scenario response

This decision tree asks about the activities the control room operator performs during a scenario. The concern is the potential for these activities to cause the operator to miss the signs that a scenario is escalating.

2.8.1 TYPICAL ARRANGEMENTS

Control room operators can be very busy during scenarios. Some sites have formal arrangements so that other people assist. This may be a control room operator from another plant or part of the site, a field operator or the Supervisor. They can assist by monitoring parts of the plant not directly affected by the scenario or by performing tasks such as making phone calls, keeping a log or collating roll call.

2.8.2 ISSUES RAISED BY ASSESSMENTS

It is fairly normal for control room operators to activate the alarm that alerts people on plant that they need to evacuate. This does not create any undue demand on them providing they have the authority to do it when they see fit. In some cases they may require, or perceive they require, permission to sound the alarm; or have to ask someone else to do it. Where this is the case they can spend some considerable time making contact with the appropriate person, explaining the scenario etc. Not only can this delay the response, it distracts the control room operator from what is happening on the plant and adds to their stress.

Another major demand on control room operators can be the requirement to make phone calls to internal and external emergency services, other departments on the site, managers and neighbours. Whilst these calls have to be made, they can take some considerable time, especially if the people receiving the calls ask lots of questions. It is usually better if someone else makes these calls, but that person must be available; and have access to a phone and the information that needs to be communicated.

In the past an event would often require ‘muscle power’ in the field. With modern plant it is often the control room operator who is in more need of support because they are the communication hub and they are often responsible for a number of plants that have to be monitored in addition to managing the scenario. They often need assistance but there is not always someone with the necessary sufficient competences available. Also, control rooms are not always designed in a way that allows additional people to assist.

Field operators are usually required to assess the scenario, but can be subject to competing demands. Sometimes they may be key members of the on-site emergency team, which requires them to muster somewhere away from the scene and means they are no longer available to deal with the immediate issues. They may be responsible for others areas of the site, and so feel over stretched. This can be a particular concern if it means they feel obliged to make a quick response, which they may achieve by circumventing key safety rules such as going to the scene alone without a ‘buddy’ or failing to obtain required PPE, breathing apparatus etc. Whilst, on larger sites, arrangements may be in place for field operators from other areas to assist in an emergency they often do not know the plant or its procedures so their assistance is of limited value.

Incident Controller is a critical role in an emergency. Whilst they should be working at more strategic level, they can often get sucked into the detail of the incident. This has been a particular concern where the Incident Controller are expected to liaise directly with emergency services or get bogged down with relatively simple tasks such as phoning

emergency services and performing roll call which, although critical, can usually be performed by others.

Issues can occur when people are expected to fulfil a key role such as Incident Controller cannot be contacted or have been incapacitated. The obvious solution is for someone else to stand-in. However, this is rarely covered in emergency plans or exercises, and in practice can result in a person leaving one key role to fulfil another.

3 LADDER ASSESSMENTS

It is not clear to me where the concept of a 'Ladder' assessment came from, but I find they work quite well. Unlike the physical assessment, they are not applied to hazardous scenarios. Instead, each ladder considers a management or organisational arrangement that influences the ability of individuals and teams to deal with situations they may encounter.

The HSE report provides a set of 11 ladders, with another being added in the EI report. Each ladder consists of a number of statements that make up the 'rungs.' The assessment involves selecting a ladder, discussing some introductory questions before moving onto the ladder itself. The text making up the bottom rung of the ladder is discussed to determine if the arrangements in place meet or exceed what is described. If they do the assessment moves to the rung above and continues. A point on each ladder is identified as the minimum acceptable, whilst the top rung is considered to be industry best practice.

Each ladder is discussed below. I have summarised my experiences of using the ladders, including the typical issues that are raised.

3.1 Ladder 1 – Situational awareness

Situational awareness is concerned with the ability of people to keep track of what is happening. This awareness is used to detect when action is required and for determining what should be done. It is important that people have access to up to date, accurate information, otherwise they may miss important events or take the wrong action, either of which can allow small events to escalate to major incidents. Also, the quality of situational awareness can significantly affect efficiency and productiveness.

3.1.1 TYPICAL ARRANGEMENTS

For operating teams working shifts, their situational awareness will start with a shift handover. They will then usually carry out some form of plant check. For control room operators this will involve reviewing process data and alarms on their control console. For field operators it will usually involve a patrol of the plant, looking at local instruments and using their human senses to detect anything unusual. After an initial orientation, a team meeting may take place to discuss plant status and plans for the day.

As a shift progresses, individuals will update their knowledge of what is going on by repeating control room and field checks. They may receive information from other sources such as lab sample results or reports from other plants. Also, they need to have up to date information about planned and on-going maintenance and other non-routine activities.

Although the focus for situational awareness is often on operating teams, it is clear that everyone in any role needs to have an up to date and accurate knowledge of what is going on. For example maintenance personnel need to plan their work according to plant priorities, which can be affected by equipment breakdowns. Managers need to make decisions based on relevant information. Commercial people need to know what the plant can do before they make agreements with customers.

3.1.2 ISSUES RAISED BY ASSESSMENTS

Although I don't think I have ever come across an operating team working shifts that did not carry out a shift handover, there is great potential to improve quality. Many companies do little more than stipulating a handover should take place, without providing guidance of what should be covered. Personnel receive minimal training and very few companies monitor or audit the process. As a result, situational awareness at the start of a shift can at best be described as variable, and is probably quite often poor.

A shift log is often kept as a record of events, and will usually form the basis of a handover. Traditionally these have been handwritten. Some companies use pre-printed log sheets with specific headings that need to be completed, and these are an improvement over a blank sheet of paper. Some have moved to electronic logs. Where this has simply meant recording information in word processor or spreadsheet file it has made little difference to the information being recorded, and in fact operators often complain that keeping the log is more difficult and has resulted in a decline in quality. Some software tools (e.g. Opralog) have been developed that can help with logging data in a way that I believe will be very valuable for supporting the shift handover process. But, to date even companies using this type of software have not really addressed all the behaviour and communication issues related to handover.

Control room operators do spend a lot of time looking at the information available to them. Unfortunately, the move to computer based control systems has usually resulted in data being displayed over a large number of graphics pages, often totalling more than 100. Whilst control room operators will 'page through', it is often not practical to go through every one. Instead, they concentrate on the ones where there is current activity or where they experience the most problems. This would not be an issue if good overview displays were provided, but these are often poor. I consider this to be a major failure in the move to computer based control systems. In practice, there is a high reliance on alarms to indicate to operators that something requires their attention. However, I believe the activation of an alarm is an indication of poor situational awareness because it shows that the early signs of a problem have been missed. Also, the problems with alarm systems (e.g. frequent nuisance and alarm floods) are well known and apply to the majority of sites.

Field operators cannot spend their whole time walking around the plant. They have tasks to do that may require them to be in one place, they have to communicate with others and need to take breaks. Therefore, they tend to carry out a number of plant patrols during a shift. In some cases they may have check sheets to complete. These can be useful for making sure they do visit critical parts of the plant, but can result in their aim becoming to complete the sheet rather than taking an objective view of what is happening. Where check sheets are not provided there is often no definition of what needs to be covered in a plant patrol. Operators are trained 'on the job,' often with no standard being set or assessment taking place. This can lead to the perceived role of the plant patrol being devalued. There is a challenge with defining plant patrols between telling operators exactly what to do, which can result in a lack of initiative being used, vs. leaving it to the individual, which can mean key items will be missed. I don't think many companies have recognised this dilemma and are simply satisfied that patrols take place.

Permit to work systems perform an important role in communicating what maintenance is taking place, who is on site, where they are and what they are doing. Most companies do have well established systems in place that can make a positive contribution to situational awareness. However, some systems can be overly bureaucratic and the full purpose of the system has sometimes been forgotten, especially as reductions in staffing levels have increased the overall workload of individuals. One result is that the permit becomes a bit of paper that someone needs to obtain before work can start, and so all the effort is put into making sure this can happen with minimum effort. This can mean that less people are involved in the process and know what is going on. The introduction of electronic permit to work systems may provide the opportunity to overcome some of the potential problems, but I am not convinced this is actually happening in practice yet.

The use of Plant Information (PI) systems can help the situational awareness of people outside of the operating team (e.g. managers and maintenance personnel) by making data more readily available via computer networks. It can mean more people are monitoring what is happening, allowing them to identify problems that may have been missed by the operators. However, it can result in people visiting the control room and plant less often, which may mean there is less oversight of what the operators are doing.

Overall, I think the role of situational awareness has not been fully understood by companies. There seems to be an assumption that giving people access to data will mean they know what is going on. In reality, if that data is not presented correctly the effort required to maintain situational awareness means it cannot be achieved realistically.

3.2 Ladder 2 – Teamworking

This ladder relates to the fact that people being able to work together means they can achieve more than if they only ever consider their own requirements. Just because a group of people is called a team does not automatically mean it has good teamworking.

3.2.1 TYPICAL ARRANGEMENTS

Most organisations split themselves into discrete teams with fairly well defined boundaries. Operations and maintenance teams are usually separate, and each may be split into plant and/or discipline specific teams. Although there has been some experimentation in recent years with the concepts of 'self-managed' and 'multi-skilled' teams, these have not been very successful at most companies working with major accident hazards. Therefore, hierarchies remain, with people working at higher levels in the organisation tending to have responsibility for more than one operating or maintenance team. These people, typically working in engineering, technical and support roles may consider themselves to be working within a department rather than a team; and so they may not always recognise how teamworking applies to them.

3.2.2 ISSUES RAISED BY ASSESSMENTS

It appears to me that most teams are expected to be self-sufficient and able to deal with every eventuality. They rarely ask for external support and do not really expect to receive it. There is often a degree of competition between teams and, whilst this can be healthy, I think in many cases it has a negative influence. In particular, rivalry between operations and maintenance teams often results in a transfer of blame rather than sharing of responsibility. Also, there can be inequity between teams performing the same function at different parts of the business. Some plants see themselves as being more important because they are bigger, more profitable or have been around for longer. As a result they expect to be treated differently, do not believe they can learn from others and are reluctant to help. They become arrogant and complacent; whilst others feel resentful. These are not qualities that support good teamworking.

Many companies have reduced staffing levels by removing a layer of supervision. Often this has led to the loss of the person (e.g. Shift Manager or Site Supervisor) who would have had the authority to organise support for teams that were experiencing high demands, including during weekends and night shifts when day staff are not on site. Without that single point of authority it can be difficult for one team to ask another for support because no one is in the position to determine priority objectively.

3.3 Ladder 3 – alertness and fatigue

This ladder is mostly concerned with the hours people work and how the risks of fatigue are managed. It is particularly concerned with shift-work.

3.3.1 TYPICAL ARRANGEMENTS

Shift-work is prevalent through industries handling major accident hazards. It is an obvious requirement for operations teams whenever the process is required to operate 24 hours per day. Some companies have maintenance personnel working shifts, although it is more normal these days for them to only be on call during nights and weekends. Other personnel that may work shifts include security and laboratory.

8 and 12 hour shifts are common, with a large number of different shift patterns. The main variations include the number of shifts worked per cycle and the rotation between day, morning, afternoon and night shifts. Also, the provision for people to take holidays which is typically achieved by either including long breaks within the shift pattern or by shift-workers covering each other.

3.3.2 ISSUES RAISED BY ASSESSMENTS

There tends to be a general interest in shift patterns and their effect on fatigue risk. In my experience, most 12 hour shift patterns achieve a reasonable compromise because double shifts and 'quick turnarounds' are essentially impossible. However, whilst it is perfectly possible to design a reasonable 8-hour shift pattern, there are still some in use that are far from ideal.

Whilst the shift pattern is relevant, the actual hours people work is often quite different because of 'shift swaps' and overtime, which is used to cover sickness, vacancies that have not been filled and projects. Few companies have rules that define what hours can be worked and they are relatively poor at monitoring what happens in practice. They rarely look for instances where someone was exposed to a high risk of fatigue.

One of the reasons companies may not pay too much attention to working patterns and fatigue is because problems are reported rarely. This may be because shift-workers are happy to accept the risk if the shift arrangements fit in with their personal lives. Also, companies may be concerned that restricting working hours may remove flexibility that is required at busy times, especially as staffing levels are reduced.

The working environment can affect fatigue. This can be a particular issue with control rooms because of poor lighting and ventilation. Another issue that causes me serious concern is that control room operators often have minimal breaks on a shift. Eating meals at the control console is very common and suggests to me that little is being done to make sure this group of people, that arguably have the most critical role, are fit and alert at all times.

3.4 Ladder 4 – employee health

This ladder reflects the fact that people working under the influence of alcohol or drugs can create a significant risk to themselves and others. Companies working with major accident hazards have to have systems in place to ensure employees understand the rules and to deal with non-compliance if it occurs.

3.4.1 TYPICAL ARRANGEMENTS

I think every staffing assessment I have carried out has confirmed that the company concerned had some form of drugs and alcohol policy. They ban consumption at work and require employees to arrive at work in a fit state. People tend to be reasonably happy with the policies in place.

3.4.2 ISSUES RAISED BY ASSESSMENTS

Most drugs policies tend to focus on illegal drugs. The fact that prescription and over-the-counter medication may affect someone's ability to work safely is often glossed over.

This ladder does prompt discussion around health monitoring programmes. It is clear that cost cutting has resulted in on-site medical staff being less common. Where it remains, it is often contracted out. People sometimes complain that they do not have access to the health services they used to or that it has become difficult to attend medicals and assessments; with the expectation that shift workers will do this on their days off.

3.5 Ladder 5 – training and development

This tree examines how people learn their job initially and then develop their skills over time. It expects that people have structured development plans that ensure they understand the risks of working with major accident hazards, as well as being able to handle 'normal' situations.

3.5.1 TYPICAL ARRANGEMENTS

Most companies will take new employees through some form of induction training when they first join. The amount of formal training that follows varies significantly. Some have the resources to put people through class room and workshop training, whilst others rely almost entirely on 'on the job' training.

Most companies dealing with major accident hazards will carry out emergency exercises. The frequency varies. Quite a lot of emphasis is put on annual, large scale exercises involving emergency services. Some will supplement these with more regular small scale 'table top' sessions.

3.5.2 ISSUES RAISED BY ASSESSMENTS

Training is a function that goes through peaks and troughs at many companies. The number of people working in the 'training department' will vary from year to year. They are often a victim of their own success, because once a lot of training has been completed the need appears to have reduced and the cost of maintaining the department is considered to be difficult to justify. The department then declines to the point where training almost stops, and so a lot of effort is required to reinvigorate it. Also, I do believe people working in training departments sometimes become complacent, taking on an administrative role once they have a good system in place. They stop working with colleagues to identify specific needs and sometimes lose interest in providing hands-on training, being happier to send people on courses provided by others. It is probably fair to say that the longer someone spends in a training role, the less in-touch they are with what is happening in the business and so their effectiveness inevitably reduces.

Even at companies that provide formal training, I believe that most of the important training is carried out 'on the job.' This should not be a surprise, because this is probably the only way to learn a job properly. However, the methods companies use to plan and assess training rarely cover 'on the job' training, which I believe is a major deficiency.

People learning a new job often have to organise their own 'on the job' training in their 'spare' time. This is often difficult, and can result in people trying to learn jobs at times when they are not particularly receptive to learning (e.g. in the early hours of the morning). Training can take a very long time, which causes problems for the business if it creates skill gaps and calls into question whether people are able to retain sufficient levels of knowledge and skill by the time they finish their training.

Companies often take a very narrow view when it comes to refresher training. When asked about it they will often refer to first-aid and fork-lift-truck driving; where the requirements are clearly defined. However, for more site specific skills and knowledge related to critical aspects of the job and major accident hazards there is often little in place to manage refresher training.

Emergency exercises are often managed poorly. It seems that so much attention is put onto a major annual exercise that nothing else gets done. Given that only a relatively small proportion of the workforce can be involved and only one scenario can be covered; the learning experience is very limited. My view is that everyone should be involved in several emergency exercises per year. Most of these can be a 'table top' talk through. The aim should be to give people experience of different scenarios and requirements to implement

different elements of the emergency plan. Obviously this will require some effort, but if it becomes a habit I don't think it need be too onerous.

One aspect of training that has not really developed as may have been expected is simulation. This is because most companies have not invested in simulators. However, I believe there are many forms of low-tech simulation that can be effective. Interestingly, companies that do have high-tech simulators do not always use them very effectively. They are expensive to operate and maintain, and sometimes are too sophisticated for providing basic training. Also, they need to be modified when changes are made to the plant or process, which is an additional expense and does not always happen.

The good news is that companies are starting to do more to develop competency management systems that may result in improve training systems in the longer term. Unfortunately, many are focussing on keeping records of training and completion of assessments. Whilst an improvement, they seem to fall well short of a system that will drive business and safety improvements by making sure the necessary competencies within the company continually match the requirements.

3.6 Ladder 6 – Roles and responsibilities

This ladder examines how well roles and responsibilities are defined and understood. It requires that systems are in place to ensure core competencies are always fulfilled through recruitment, training and management of change.

3.6.1 TYPICAL ARRANGEMENTS

In the past, companies may have written quite detailed job descriptions, possibly tailored to individuals. They are now more often written at a general, higher level. This has become a requirement as organisations have de-layered and staffing levels have been reduced requiring people to be more flexible in their ways of working.

Although there have been experiments with 'self-managed teams,' which would have resulted in very wide ranging and hence general job descriptions, they have not been very successful in organisations dealing with major accident hazards. Keeping some levels of hierarchy has helped people know their place. However, terms such as 'dotted line responsibility' and 'matrix management' have entered the vocabulary, reflecting the fact that traditional hierarchies have their weaknesses.

3.6.2 ISSUES RAISED BY ASSESSMENTS

Most people do have job descriptions but refer to them infrequently. This is generally taken as an indication that they are of limited value, often because they lack detail. Whilst the trend is for only high-level job descriptions, it is probably fair to say that most people would prefer them to be more detailed and specific.

Although they do not read their job descriptions very often, people are usually comfortable that they know (or assume they know) what is expected of them in both normal and abnormal situations. People may interpret this to mean that they will do anything they are asked or told to do by their supervisor or manager, but I think most people working with major accident hazards are more proactive than this.

The main issue I have with the way roles and responsibilities are managed is that the focus is often on the individual, with less consideration given to team or company-wide requirements. Whilst teams are usually set-up to fulfil core competencies, there is rarely a robust system that ensures all are fulfilled all of the time. Few organisations have systems in place that trigger actions when the people present are unable to fulfil the core competencies. The suggestion of a 'traffic lights' approach appeals to me for directing activities on a day-by-day basis. Green indicates that the people at work fulfil all core competencies and activities can continue as normal. Amber indicates there are some potential weaknesses, and

activities should be restricted. Red indicates major weaknesses and the process should be put into its safest state until other people can be brought in to fulfil the core competencies.

3.7 Willingness to initiate recovery actions

Even when they manage their risks well, companies dealing with major accident hazards have to be prepared to respond to unplanned events. If actions are delayed the consequences can be serious. The problem is that the response usually involves stopping the process, which clearly has implications for production and profit. Also, there can be risks associated with responses such as 'crash' shutdowns, and so initiating such actions cannot be taken lightly.

3.7.1 TYPICAL ARRANGEMENTS

The normal response to any major event that may result in loss of control is to shutdown plant. The ways of doing this depend on the nature of the process. Some can be shutdown quite quickly and easily; and restarted with minimal disruption. Others cannot.

3.7.2 ISSUES RAISED BY ASSESSMENTS

I am happy to say that, in my experience, most operating teams are happy to initiate whatever response they feel appropriate if there is a reason to do so. In some companies the operators may prefer to receive approval from their Supervisor or equivalent if they can, but would rarely delay action if they could not make immediate contact.

Whilst I believe there is the willingness to initiate the appropriate actions as required, companies do not always prepare their operating teams to make such decisions. Diagnosis skills are rarely taught or assessed. Emergency exercises are often not as effective as they should be and there is rarely any form of useful procedural support such as diagnosis or decision aids. Therefore, the only time it is determined whether an operator is able to make the decisions is when they need to do it for real. This clearly adds to the stress caused by the situation and will inevitably lead to errors.

The perception that an operator may be reprimanded for shutting down a plant in a situation where it was subsequently deemed unnecessary is fairly widespread. This can be because it is felt managers are more interested in production than safety, based on what they say on a day to day basis. I have to say I very rarely hear of any situations where an operator has actually been reprimanded for such an error of judgement, although the perception persists.

One issue that is a little surprising, on the face of it at least, is that there is not a consensus on when it is appropriate to automate shutdown initiation or to leave it to the operators' discretion. It introduces the possibility of a number of operator errors that can result in delay or the wrong response. Also, it can lead to the accusation that management are worried about the cost of spurious initiation. Having looked at this issue on a number of occasions I am satisfied that in many cases the number of variables required to be considered when deciding what course of action to follow is so great it is better to leave it to people. Automated systems are not yet sophisticated enough to take account of what is happening on the plant, the location of personnel, weather conditions, availability of resources etc.

3.8 Ladder 8 – Management of procedures

Although the dictionary definition of 'procedure' indicates that it refers to a method of performing a task, for most companies a procedure is a written document. Other terms such as 'instruction' and 'operating practice' may be used, but they usually perform essentially the same purpose.

3.8.1 TYPICAL ARRANGEMENTS

I can safely say that every company I have ever worked at has had some procedures. They tend to fall into different categories including:

- Operations and maintenance;
- Safe systems of work;
- Safety, environmental and quality management system;
- Emergency;
- Business administration.

3.8.2 ISSUES RAISED BY ASSESSMENTS

Whilst every company may have procedures, most have problems with them. In my opinion, the root of these problems has been a misguided attempt to generate written procedures in the same format for every activity carried out. The result is a bureaucratic nightmare where the effort spent on managing the system outweighs the benefits. The procedures do not fulfil the users' requirements and means people will often perform highly critical tasks, where the potential consequences of human error can be significant, with nothing in place to define the correct method or to assist them in adhering to the safe method. I would go as far to say that the problems are so fundamental that there is no point in listing here the many issues I have observed when carrying out staffing assessments!

3.9 Ladder 9 – Management of change

Experience has shown that failing to manage change can have significant consequences. The main concerns being that the changed system may not be as safe as it was before or that risks may arise during transition from current to new arrangements.

3.9.1 TYPICAL ARRANGEMENTS

Every company I have worked with that deals with major accident hazards has had some form of management of change procedure. They typically include a form where the details of change, potential risks, methods of control and formal review and approval are documented. It is recognised that changes to premises, plant, process, substances, procedures and people or organisation need to be managed.

3.9.2 ISSUES RAISED BY ASSESSMENTS

Whilst most companies have had procedures in place to manage change for some time, the application to people and organisation is relatively new to most. I have seen very few examples of documents that demonstrate a human or organisation change has been fully assessed and implemented in a controlled manner.

Human and organisation factors can be affected by all types of change (e.g. changes to plant, process, procedures). Whilst most companies do recognise that factors such as roles and responsibilities, training and competence, and procedures should be considered, evidence suggests that these are often seen as being of secondary importance. This probably reflects that most changes are led by engineers, who are less comfortable dealing with human and organisational factors.

There are quite a few weaknesses in the way companies manage all types of change. There is usually an expectation that every planned change will be successful, and so little thought is given to what could go wrong and the need for a contingency. The focus is on how the changed system will function, and so transition phases are not always identified or considered. And, although nearly every management of change procedure includes a requirement for reviews to take place after changes have been implemented; these happen only very rarely.

3.10 Ladder 10 – Continuous improvement of safety

It is unlikely that we will ever get to stage where we will decide safety is 'good enough.' This is especially the case where the hazards are significant and systems are complex. Even if an organisation were to achieve great performance, the challenge would be to do this in a more efficient manner. Therefore, managing safety in a way that results in continual improvement is an essential aim.

3.10.1 TYPICAL ARRANGEMENTS

Every company I have worked with that deals with major accident hazards has had some form of accident reporting and investigation system. Most carry out audits and genuinely want to learn about events so that they prevent them recurring.

3.10.2 ISSUES RAISED BY ASSESSMENTS

I am satisfied that if an incident occurs where there is a tangible consequence, it will be reported and direct causes would be identified. I am less satisfied that root causes are identified, especially with regard to human and organisational factors. Also, I am not convinced that companies have been very successful at developing systems for reporting and investigating near misses. Overall, the level of effort put into analysing incidents is usually determined by the actual consequences experienced rather than the potential, which can mean learning about low frequency, major accident scenarios is relatively poor.

Many companies only expect to learn from their own experiences. Even if part of a large group, sharing information between sites or divisions can be poor. There are some organisations operating in the UK that provide a good opportunity for companies to learn from others either in their locality or industrial sector. These seem to work well at an informal level but there does not seem to be an appetite for companies to recognise this as a formal part of managing safety.

I am often surprised that companies do not take the time to read published accident reports properly. They may read the executive summary or press reports. But there seems to be a perception that there is not much to be learnt because the company involved may have been working in a different sector or country, or that the hazards, plant and systems used were different. I think part of the reason is that people focus on the direct causes and do not look at the fundamental root causes, particularly related to human and organisational factors that can usually be applied very widely.

3.11 Ladder 11 – Management of safety

The requirements for safety management systems are clearly laid out in various documents including HSE guidance document HSG65 and the health and safety Standard 18001. The general starting point is a defined policy. An organisation and set of procedures are required to implement the policy. It is then necessary to monitor and audit the effectiveness of the system. Human and organisational factors should be an integral part of the system if they are going to be managed effectively.

3.11.1 TYPICAL ARRANGEMENTS

Every company I have worked with that deals with major accident hazards has had a health and safety management system. They are usually compliant with guidance and standards. Some are integrated with quality and/or environmental systems; whilst many are stand-alone.

3.11.2 ISSUES RAISED BY ASSESSMENTS

On the face of it there are rarely concerns with safety management systems. However, as with many things the issues are not with what is documented but how it is implemented in

practice. One of the main concerns is that companies are good at reacting but relatively poor at being proactive. Recent emphasis on the identification of 'leading indicators' of safety performance may help, but this is a new concept for most companies and there is some way to go before it is properly established.

There can be some concerns that safety management systems 'belong' to managers. Shift workers, in particular, feel they have minimal active involvement because they find it difficult to attend meetings or take place in audits. This perception is often reinforced because employees feel they do not get much feedback regarding the issues they raise and the results of audits and incident investigations.

3.12 Ladder 12 – Automated plant and/or equipment

The degree of automation used by companies has steadily increased as new technology has been introduced. This has had many benefits including improved performance and reliability; and it has relieved people of boring, unpleasant and potentially hazardous tasks. However, automation can have negative effects on human factors and so its use needs to be managed.

This ladder was not included in the original methodology. It was added by the EI report.

3.12.1 TYPICAL ARRANGEMENTS

There are many different ways that automation is used in practice. It ranges from simple control loops maintaining parameters at a defined set point through multi-variant optimisation and complex sequential mode changes (e.g. plant start-up). It is used for normal operations and in response to emergency situations (e.g. automated plant shutdown or depressurisation).

3.12.2 ISSUES RAISED BY ASSESSMENTS

It appears to me that most people assume that all automation is beneficial and the aim should always be to use it as much as possible. But it is very rare for any objective assessments to be made after automation has been implemented to confirm the expected benefits have been realised or to identify any negative impacts. Consideration of human and organisational factors is usually limited to the provision of training, but this often just tells people how the automation works rather than explaining how they are supposed to use it.

It is fairly common to find that automated systems are not used in practice, often because people do not understand what it is supposed to do or perceive it makes their job more difficult. People can be particularly reluctant to touch an automated system so that if it is on when they start their shift they will leave it on, but if it is off they will not switch it on.

4 CURRENT BEST PRACTICE

I think it is reasonable to assume that what was considered to be best practice in 1999 is now out of date. I believe the concept of High Reliability Organisations, that has developed over the last decade, gives us a good reference point for what may be considered as best practice now. The table below summarises how a new top rung could be developed for each ladder

Ladder Topic	Defining best practice based on high reliability organisations
Situational awareness	A key feature of high reliability organisations is that they never believe that they have good situational awareness. They spend more time collecting and analysing data than other organisations, looking for indications that things may not be as they should be and that this is an indication that something could go wrong. As a result they are better at avoiding problems, and spend less time responding to them.
Teamworking	Teams within high reliability organisations adapt to suit the situation. When things are normal and under control, teams and individuals work flexibly so that priority actions are performed in the most efficient manner. People rotate through roles so that they do not stagnate. When demands increase the focus changes to making sure the right person is in the right place, with their position within the organisation becoming less important. But in an emergency it is recognised that pre-defined team structures provide resilience, and systems are in place to make sure people have an in-depth knowledge of what they need to do.
Alertness and fatigue	I doubt high reliability organisations have significantly different arrangements to manage the risks of fatigue. However, they are probably more inclined to accept that fatigue is inevitable, and to have a culture where this is recognised and dealt with effectively with minimal fuss. For example, I would expect a high reliability organisation would rather someone took a 'power nap' at work if they felt they needed to, whereas most companies seem to take a pretty dim view of this practice.
Employee health	I would not expect a high reliability organisation to have systems for managing fitness for work that are significantly different to any other company. However, I would expect this issue would be perceived as a higher priority. They would acknowledge the risks of people being at work below their best, and do what they can to minimise the likelihood of this happening. The main difference with a high reliability organisation may be that it could implement arrangements that in other organisations may be perceived as being beyond the scope of the employee/employer relationship (e.g. non-work related health issues). Instead of being seen as delving into people's private lives, it would be perceived as a sign that the employer is genuinely concerned about its employees' health and fitness.
Training and development	High reliability organisations are known to train their employees far more than others. They recognise that the ability to handle all situations that can occur requires an in-depth knowledge of how the system works and the skills to handle unpredictability. Also, they rotate people through roles so that they do not become stale or complacent, which means they need continual training. Anything learnt from an incident, near miss or other event is used to update training. They value expertise, deferring to the person who has the knowledge or skill required at any particular time, no matter what their grade or position in the organisation.
Roles and responsibilities	Roles in high reliability organisations tend to adapt dynamically so that the right people are always in place to address the priority issues at any time. There is a high degree of trust that people will do what needs to be done; but they are also supported so that they can work effectively.
Willingness to initiate recovery actions	Although high reliability organisations put a lot of effort into prevention, they still expect things to go wrong. Other organisations seem to be more satisfied that they can prevent events and see emergencies as being so rare that there is little they can do to prepare for them. This is a fundamental, cultural difference that significantly affects resilience.

Ladder Topic	Defining best practice based on high reliability organisations
Management of procedures	Ironically, whilst one of my main complaints about procedures is that there are too many; high reliability organisations probably use procedures more than others. The main difference is that they make sure they have procedures where they are needed and that they provide the information needed to avoid and deal with problems. As a result their procedures do not just define how to perform a task but also prompt people to be ready for things to go wrong. Procedures are viewed as an integral part of the job so that individuals and teams will not progress with a task if they feel a procedure would be useful and one was not available.
Management of change	High reliability organisations will have a greater appreciation of the fact that change is happening continuously. They recognise that maintaining up to date knowledge of how their system functions is important if they are to prevent problems and respond to them effectively if they do occur. They are aware of the issues of 'change by stealth' where lots of small changes occurring can, over time, take a system a long way from where it started even though no planned intervention has taken place to implement a change.
Continuous improvement of safety	Continual learning is an integral part of high reliability organisations. They do not wait for something to go wrong before identifying opportunities for improvement. They assume there is always weakness in their systems, they look for those weaknesses, investigate how they were able to occur and how resilience can be improved.
Management of safety	High reliability organisations don't need to talk about safety in isolation as it is seen as an integral part of their business. They are happy to discuss any issues raised at all levels in the organisation, whether it is good or bad news. Any safety issue is seen as a serious concern for the business. High reliability organisations treat every event as an important learning opportunity and do what they can to retain the knowledge they have gained so that it can make a positive influence in the future.
Automated plant and/or equipment	High reliability organisations are very wary of complexity, and would see automation as a source of complexity. They would assume that automation could cause them problems and make sure they do not rely on it. They would not avoid using automation but would fully understand its limitations and make sure they always work within them.