Critical Incident Management Communications

Introduction

It is 2am in the morning and you have just received an urgent call from your call centre, informing you that your company's critical business systems have gone down, affecting hundreds of customers. Your mobile is continuously ringing with anxious and panicky staff, asking what they need to do. Angry customers are starting to call, demanding to know what has happened and when the problem will be resolved. What do you do in this situation? How do you solve the problem? Who should you inform? What message should you send out to your customers?

These are just some of the questions likely to be going through a business decision-maker's mind at the time of a major business incident affecting their service. Solving them quickly and efficiently can be tricky, so it worthwhile taking some time to think about how you will tackle such incidents and plan in advance. Irrespective of the size of your business or whether it relies on a real-time or critical service being available to your customers, a disaster recovery plan can be vital to its survival. This plan should include procedures for communicating internally with staff and externally with customers and suppliers.

This article provides some insight into incident management procedures and best practise in determining policy for incident management communications.

My understanding of the issues and difficulties surrounding the incident management process has developed over a period of four years, spent working in a real-time, business-critical environment, processing online payments for the UK's largest online payment service provider. With over 40,000 customers and millions of pounds worth of transactions being processed every day, there was much at risk. When the payment systems go down, customer's income is directly affected, leading to potential claims for liability and compensation, as well as customer dissatisfaction with the service and customer attrition. When money, time and reputation are at stake, the risks are high and it is essential to get the process of handling critical incidents right. Internal staff and customers need to be informed of current or potential issues in the best and most efficient way. A clear policy and procedure for handling critical incidents and communicating with staff and customers is therefore essential.

In December 2003, WorldPay, a subsidiary of the Royal Bank of Scotland Group and the largest online payment service provider in the UK was hit by a sustained Denial of Service (DDOS) attack. The attack lasted several days, during which the payment and administrative systems were unavailable, affecting thousands of merchants relying on the WorldPay payment systems to process their payments. The scale and repercussions of the attack were serious enough to make news headlines around the world.

A DDOS attack is a sustained, potentially criminal attempt to bring down an internet site or service through the use of simultaneous and continuous bombardment of the site by thousands of computers that have been affected by a virus. The result is that the website and servers cannot handle the massive scale of demand for resources and therefore crash. In the WorldPay case, the main website, support site, administrative systems and and payment systems were all taken down. At the time, all customer communications were also being sent through internal systems, so when these went down, there was no quick way of communicating with customers. In the ensuing days, thousands of concerned customers called into the call centre, causing the phone systems to crash. The company had no incident management procedure in place and the result was a panicked response and long delays in providing customers with information.

The incident management policy and procedures guidelines discussed in this article originated out of the trauma of the WorldPay DDOS incident. This has been gradually refined over the course of 4 years, in response to lessons learnt from dealing with hundreds of different types of incidents. The incident management process and team now supports a number of Royal Bank of Scotland companies involved in real-time payment processing. A key element in refining the procedures has been the continual evaluation, through the use of post-event wash-up sessions, to learn lessons and draw conclusions for the future.

As with any business process, what has been set up is not perfect. However, given the constraints and budget provided, the process does take into account the needs of the business, the requirements of customers for information and the realities and costs involved in setting up the systems and personnel to handle business critical incidents. By sharing some of the conclusions we reached in this article, the intent is not to dictate how your own business incident management process should work, but provide insight into some of the considerations, systems and personnel that may be worth considering when setting up your own business communications response.

Here are some of the questions that were regularly debated when setting up the incident management policy:

- What is an incident? In other words, when is an event serious enough to make it an “incident” that requires involving staff and informing customers?
- Who should handle an incident? Should this be a technical person? Someone on the call centre? A business decision-maker? An account manager?
- Who should write the communications that go out to staff and customers? A PR or communications person? Another member of staff?
- What type of training and support should you provide to staff who handle incidents? For example, any special computer or communications equipment? How much do you pay staff for handling an incident? How much should they be paid for being ‘on call?’ How many staff do you need to be on call, at any given time?
- When do you communicate with customers? What is the threshold or criteria for setting this off? How long should you wait, after being informed about an incident, before communicating?

- Which customers do you communicate with? All or a select few? If you have customers who receive their support in another language, do you provide incident communications in that language?
- How do you inform customers of the problem? By phone, by email? By SMS? On your Website? Using automated systems? What systems or methods should you use for doing this?
- What should you say, when informing customers? What should you not say?
- What should you say once the incident is over?

The answers to many of these questions will depend on considerations such as the size, scale and nature of your business, the importance of any incident to your business and customers, how frequently these occur and what budget you have available for this.

Here is what we considered for each of these questions.

What is an incident?

Your business systems may run on several different hardware devices and involve a variety of different software products and processes. You may be dependent on external services providers for some aspects of your service, such as your Internet and telephony systems. Any of these aspects can fail or cause problems.

Some of these problems may be minor enough not to actually affect the service your customers receive, although they may need technical attention to remedy. Other problems might affect your service but be beyond your control, due to a problem with an external service provider.

It is sometimes difficult to identify when an event is serious enough to qualify as an incident.

In our payment service provider environment, the situation was incredibly complicated by the number of factors that could impact on our service. When dealing with a large-scale system that handles millions of pounds worth of transactions daily, from across the globe, backup and redundancy mechanisms are essential, but they cannot cover all scenarios.

Software releases and system updates could potentially impact on services once released live – this could happen on any of several points along the line – sometimes related to other services or systems run by the bank, which had no direct bearing on our service. At other times, new patches put in to fix one problem, caused unanticipated impacts on service when released to a live environment.

In addition to the primary payment systems, at any time there could be a number of other supporting systems and products that could be affected, for example, access to the administrative systems. Other issues related to database problems, downtime that was necessary following the installation of new hardware, or switch-over between primary databases, for maintenance purposes.

Although we had several Internet service provider routes, if any one of these routes went down, it might effect a section of our customer base, in a particular geographical region – e.g., USA or Asia pacific.

At other times, there were issues with the banking network to which our systems were connected – this meant that although our own systems were functioning normally, from the customer's perspective, their transactions were not being authorised.

At other times, a customer might phone in to report a problem, which after investigation turned out to be a problem with their own systems.

What we discovered from experience was that it was not always possible to identify immediately whether there actually was an incident and what was causing a problem. In fact, it usually took time. For example, an on-call technical engineer might receive automated alerts from the systems, indicating a potential issue – however, it then required follow-up and investigation to identify firstly if there were any significant problems and then, what the issue was and which services were affected. At this stage, we might not know enough to be able to tell the customer anything, but on the other hand, not saying anything to the customer if it turned out to be a major incident would be equally damaging.

As a rule of thumb, if an incident did not affect our critical service path (payment and administrative systems) or only lasted for a very minor period, then it could be handled as a routine customer communication, which could then be scheduled if required. If an incident was severe and ongoing (for example, lasting more than 15 minutes), then the incident management process was engaged and an 'emergency' communication would be prepared for customers.

Who should handle an incident?

In any major incident, and even for minor incidents, it is likely that a number of staff and key business decision-makers will need to be involved. Depending on the scale of your business, it may be necessary to appoint one person to be responsible for diagnosis and repair of a problem, another for communication with staff and customers, and another for co-ordination and business decision-making. The reason for this is that an emergency situation requires staff to respond quickly and focus on specific tasks.

In a typical WorldPay incident, we found that identification of a potential problem usually came from a number of sources – Call centre staff and corporate account managers would receive queries from customers and forward these on to the staff involved in incident management. Technical engineers would often receive system alerts, indicating that further investigation was required. Each of these incidents required co-ordination amongst staff in different departments, a decision-making process and a decision to either communicate or not with the customer.

Another important consideration in deciding who should handle an incident was the time when the incident occurred. If the incident occurred during office hours, regular call centre staff and business managers were available to help out. To cater for incidents that occurred on the weekend or after work hours, staff needed to be on call, to handle a potential event. After some initial trial and error, we reached the conclusion that involving regular business staff – pulling them off their current tasks - was not efficient and took them away from everyday business activity, which would then cause delays in other areas of the business.

The solution we put in place was a full incident management team, which would be on call 24 hours a day, to handle any incidents. In addition to their regular duties during office hours, these key staff would be available if required to handle incidents. The same team would be available out of hours. We came up with the following team (after some period of refinement):

- A business incident manager – to co-ordinate with various parts of the business, make sure everyone required was talking to each other and having overall responsibility for ensuring that the incident was appropriately resolved.
- A technical manager on call, to respond to any potential issues and investigate and make any necessary repairs. The technical manager could call on any other technical support staff to assist if required.
- A communications manager – to prepare the communications to the customer and to internal staff, and to send this out as soon as possible, using the agreed upon methods. This person was also responsible for post-incident reports.

The rationale behind this team was as follows: While a technical engineer was investigating a complex issue, he would not have the time or be in a position to talk to key business decision-makers, then draft and send out potentially sensitive information to customers. At the same time, a business decision-maker might be talking to other decision-makers, technical staff, others business divisions or external service providers to resolve the situation and would not be in a position to draft and send out the communications. A communications manager would have the skills and background to draft the communication, ensure the message was appropriately worded for the target audience put it through a review process and send it out in different formats and different languages.

In addition to the core ‘on call’ incident management team, a number of other staff might be involved at any given time, given the nature of the problem. These included senior call centre managers, Public Relations managers and directors, who could make time available for an important crisis. All such staff received ongoing information on any incidents, through internal communications, and where relevant, for major incidents, were involved in any decisions.

In our circumstances, given the critical nature of our services, the complexity and large number of customers and transactions involved, this structure made sense. Obviously, for a smaller organisation, the costs involved of having three or more members of staff ‘on call’ at the same time might not be cost-effective.

We will discuss later the use and limitations of automated alerting and messaging systems, but it is sufficient to say that in our case this did not overcome the need for the team structure we had set up.

What type of training and support should you provide to the staff handling incidents?

Any on call staff need to be involved and trained in the incident management process you have put in place. In addition to policies, this may involve training on different types of systems. For technical staff, this might be on the alerting and diagnostic systems. Communications managers might need training on any tools used to prepare and send information to customers.

On call staff will need access to relevant account usernames and passwords, the telephone numbers of key decision-makers or other on-call staff that need to be contacted. They may need access to secure databases, where information such as customer details is stored.

On call staff will probably need a mobile phone and access to a laptop with a secure Virtual Private Network (VPN) connection to the office, unless you are expecting them to come into the office each time there is an incident. Expecting staff to come into the office each time there is an incident may not be practical if staff have to travel far or a resolution to a problem is needed quickly. You may want to consider wireless connectivity for staff who are on the move, e.g., in an airport or hotel, and do not have access to a phone line or broadband connection at the time.

With regards to pay for on call coverage and out of hour’s support, that would need to be a decision for your business, or down to negotiation with the relevant members of staff. As a rule of thumb, unless specified otherwise in their contract of employment, staff handling an out of hours incident could be expected to receive their contracted overtime or weekend rate for the time during which they are working on an incident. In addition, it is standard practise in most industries for on-call staff to receive a certain amount for making themselves available to respond to calls on weekends or after work hours. They receive this amount irrespective of whether there is an incident during this period – and it does make up for the inconvenience of being woken up at 2am regarding a potential incident, which turns out to be a false alarm.

When do you communicate with customers?

The golden rule is to tell customers the information they need to know, when they need to know it. Customers will often indicate their communication preferences. We found, for example, that some customers requested to be informed each time there was an incident.

If a customer is directly affected by failure of a service you are providing them, then best practise would be to communicate this to the customer. If a business system failure is unrelated to the customer and does not directly affect them, then don't communicate this.

Which customers do you communicate with?

Good practise is to have sufficient information stored about your customer base to be able to segment and target only those customers that are affected, rather than your whole base. For example, if a service failure only affects a specific product, only contact those customers who have purchased this product.

How do you inform customers of the problem?

Using an SMS messaging system is probably the quickest and most efficient way. Our SLA from the report of an incident to sending out an SMS to a customer was 10 minutes. This meant that 10 minutes after we first learnt about a problem on our systems, our customers were aware of it.

We preferred to go down the route of a "managed" communication, as opposed to an automated one. I would recommend that automated alerts be reserved for internal staff, who can then assess the impact and determine what they mean. The more complex your alerting system, the more alerts you will receive, and this may be graded from minor to major. It takes trained staff to determine what these mean.

At the next level, once an internal member of staff has determined that there is an incident, the communication out to customers and other internal staff needs to be "managed".

By managed, I mean that some thought and consideration must go into deciding:

a) whether to communicate out to the customer
b) how to communicate with the customer, and
c) what to say.

It is always best practise, when dealing with customers, to only send out "managed" communications. There are important legal and PR reasons behind this. A poorly managed or unmanaged communication could result in legal claims, disgruntled customers, and loss of business and reputation – which is why it is so important to get your incident communications right.

What should you say, when informing customers?

Keep it short and to the point. Only say what you know for sure, rather than being forced to back-track later.

Our policy for example, was to say "date, time: potential issue with out payment systems. We are investigating".

Provide just enough details to inform the customer of a potential issue, so that they can take remedial action. At this stage you don't need to go into any details about the reasons for the failure. The objective is to inform the customer of a potential issue – so that they can take any remedial steps necessary on their side.

What should you say once the incident is over?

Your business needs to communicate with the customer as soon as possible after an incident, to provide more details about the reasons for the occurrence. We used a standard report format to communicate the incident, which provided details of:

- Date and time
- Nature of incident/who was affected
- Brief description of the cause
- Resolution
- Steps put in place to prevent this happening again

It is important to indicate what steps are being taken to prevent the problem happening again in the future. The objective of the report is to reassure the customer that the incident has been effectively managed and that any lessons that need to be learned have been learnt.

Will you need to pay out compensation?

What should you be communicating out to customers in this regard?

A key question that follows on from any business failure is whether or not your business is now liable to pay out compensation to your customers for loss of service.

A great deal will depend on what type of services have been affected and what the contract between you and your customer states on the matter of service failure. For example, was it a service critical to the customer's business, or something that the customer could resolve in another way? Does your service contract specify this? Are there any penalties specified for failure of service? Is the failure severe enough to amount to a breach of a condition, which could potentially enable the customer to repudiate the contract?

Above all, communicating with the customer in the right way is essential to ensuring that good will is maintained and the impact on the customer's business is minimised.

When dealing with requests for compensation following an incident, your legal team will also need to be aware of the contents of the contract that has been signed between you and your customer. If your contract excludes liability for service failure during an incident and your customer has agreed to these terms, you should be safe (see the ruling in L'estrange vs Graucob 1980). However, you should also be aware of the implications of legislation such as the Unfair Contract Terms Act 1979, which states that any business clauses which attempt to exclude liability for economic loss must be reasonable, especially when this is due to negligence. This is illustrated in cases such as Rambler Motors vs Hollier 1996.

If your contract is directly with a consumer, who has been affected by the loss of service, you may not be able to exclude liability for a critical incident failure, as the use of exclusion clauses is strictly regulated by the Unfair Terms in Consumer Contracts Regulations 1999.

What does the law say about incident management policies?

There are two pieces of relevant UK legislation you should be aware of, which have some impact on the requirements for incident management policies. These are the Health and Safety at Work Act 1979 (HASWA 1979), the Data Protection Act 1999 (DPA 1999).

HASWA 1979 has mainly implications for the internal staff you employ. HASWA (and the resulting six pack regulations outlined in the Health and Safery at Work Regulations 1999) require that your business has conducted appropriate risk assessment, which includes all staff and equipment, as well as training of staff. Therefore, in the event of a major incident, such as a fire on premises, your staff should be trained to handle this. If you employ more than five members of staff, you will need to have a written policy.

The Data Protection Act 1999 has implications for how you store customer details and communicate out to customers in the event of an incident. For example, if you are using a third party supplier or agency to do the communications, have your customers given their consent for their details to be stored or used by this third party? When sending emails or other communications, how secure is the data being sent? For example, you should not include a username and password in the same email.

Recent legislation in the US (e.g., Sarbanes-Oxley Act of 2002), EEU and UK (Civil Contingencies Bill 2005) have made business more aware of the need for effective business continuity planning and disaster recovery measures.

Conclusion

Don't wait for a major incident to happen to your business before setting in place your plans and processes for handling incidents and communicating with customers.

An effective incident management and communications program could save your business thousands of pounds and above all – ensure you maintain a good reputation and the confidence of your customers, enabling your business to continue to grow and prosper.