

# What lies beneath

Changes to enterprise and e-business applications can have hidden consequences for the organisation's technology infrastructure, says Colin Butcher.

BUSINESSES NEED TO be agile and responsive to changes while delivering consistent service to their customers. An enterprise application is a complex tool that can help you achieve this – provided its use is properly supported by the infrastructure of your business.

Changes to business processes in order to fit 'the system' can erode competitive advantages and stifle innovation, so it is imperative to understand how to get the best out of these applications.

The key to making them work is to enable them to extract and deliver useful information from all that data, in a timely manner. This requires a good understanding of the kinds of information you need in order to make good business decisions, which in turn requires a knowledge of the functional elements that make up your business.

The underlying technology infrastructure must allow you to acquire data from many sources, process that data and reliably deliver the resulting information where and when it is needed.

You need to involve the people responsible for the infrastructure of your business at the outset so they can both plan for the future and identify the areas that are likely to require significant changes.

The technology infrastructure consists of the data and storage networks that allow you to

exchange data between systems, as well as the systems themselves (both large 'server' systems and smaller desktop 'client' systems).

Many businesses who implement enterprise applications tend to ignore both the potential impact on, and limitations of, their existing infrastructure. As a result, these businesses often fail to take full advantage of new applications and business processes. At worst, the poorly planned introduction of such applications can be a spectacular failure.

One area that has to be considered very carefully is the 'parallel running' of both old and new systems. This crucial step generally raises a large number of complex issues that need to be successfully resolved.

Once you have moved to an enterprise application that integrates functions across many aspects of your organisation, your business will quickly become utterly reliant on it. It has to work consistently and reliably.

In some cases, you may need to replace your entire network infrastructure to attain the necessary levels of performance and availability. You may also need to invest significantly in storage subsystems and suitable backup and archive mechanisms.

## Framework

All enterprise applications require significant effort to customise and tailor them to your specific business processes. They are not an 'out-of-the-box' solution to a set of business issues.

Without a well-designed technology structure, you will forever be making changes in a piecemeal fashion. You also

have to recognise that every so often your business requirements will change to the extent that you need to revisit the original design assumptions and may need to make some fundamental changes.

The architecture is simply a set of rules and concepts that govern how the different functional elements that you need to run your business will fit together.

In networking terminology, a good example would be the design of a TCP/IP addressing scheme that leaves plenty of spare address space to be filled in later, yet imposes conventions and rules about how you allocate addresses. For example, all 'server' class systems, printers and infrastructure components (eg, network switches) would be allocated fixed IP addresses, and all other devices would use dynamically allocated addresses.

## Availability

It's all too easy for a business to implement new practices, such as internet-based trading, then to wake up one morning to find it's trying to run what is now a 24x7 business on an infrastructure better suited to a 12x5 operation.

In order to achieve uninterrupted 24x7 operation, you need to have a fully resilient infrastructure with automated monitoring and failover mechanisms. This is technically feasible, but will be difficult and expensive to get right.

As far as possible you need to eliminate single points of failure (such as core switches) within the critical components of your infrastructure. You also need to ensure seamless operation of old and new equipment during any transition phases.

In practice, most businesses have operational windows when it is possible to take components out of service in order to make changes or repairs. You need to understand all the critical components of your infrastructure in terms of the probable time to recover from failure. That must include elements such as firmware updates, onsite spares stock levels, maintenance agreement service levels, etc.

You can minimise the risk of disruption by careful planning and testing. This must include testing both the process of implementing changes and the process of reverting to the original state.

Even then, things can go wrong, such as the recent failure of the air traffic systems at West Drayton and Swanwick during an upgrade.

## Disaster recovery

Implementing disaster recovery is another aspect of building a resilient infrastructure.

Disaster recovery is slightly different to high availability in that it tends to be concerned with business continuity following a major event that has caused severe disruption.

The main problem is ensuring that there is no data loss or corruption as a consequence of the event and that business operations can continue within an acceptable period of time, usually from an alternate location.

Losing your systems can be a real problem, but if you irretrievably lose your data, it's just like suffering a major fire – you stand a very good chance of going out of business altogether.

A big part of disaster recovery planning is ensuring data integrity. Backups take

time, which can be a problem in 24x7 operations.

Another crucial area is backup and restore. The primary purpose of a backup is for you to be able to restore normal operations quickly in the event of inadvertent deletion or equipment failure. You have to devise a suitable strategy that minimises the time taken to make a backup, yet which will get you the level of data integrity that you need.

Restores never happen at convenient moments, so the procedures have to be simple, effective and well-rehearsed. How many businesses regularly prove that their backups work by practising restores?

### Mobile working

With the drive to reduce visible direct costs and the consequent increase in home and mobile working, the complexity of the IT infrastructure has increased significantly.

Instead of linking a relatively small number of sites, companies now have a wide variety of locations that need to be connected with secure, fast and reliable networks.

Wireless data networking using wireless Ethernet is now common. It can be extremely convenient to the modern business traveller, for example with 'hotspots' in airports. It can also bring some really interesting security and data integrity issues with it.

Businesses no longer have a technology infrastructure that exists within the corporate boundary. They have networks that need to extend seamlessly beyond the corporation to wherever their workforce and customers are.

To make matters worse, those mobile devices are usually in the hands of non-technical people who require significant training and support to minimise problems and to be capable of recovering from equipment failures. Not unreasonably, these people expect the technology to work – and are frustrated and angry when it

doesn't do what they expect.

At this point, they usually try to fix things for themselves and usually make a much worse mess of it in the process.

Your mobile and remote staff need support and help from the technical infrastructure. They are not working effectively if they spend their time messing around with laptops rather than doing the job they are paid to do.

You need to build and test standard configurations for such machines, and then enforce the use of known configurations to minimise support costs. You must have sufficient resources in place to address these issues and prevent serious problems.

Theft and consequent data losses or breaches of security are another big problem with mobile working. Another area to consider is having laptops at home connected directly to the internet. Those machines become a matter for serious

concern when they cross back into the corporate infrastructure – or someone else's corporate infrastructure.

You must take action to minimise the impact of such problems.

### Access control

The main issue with security is knowing that you've been compromised. That's where intrusion detection and auditing processes come in.

Defences to keep out intruders, or more commonly the enemy within, are important but not as important as knowing you've got a problem.

Security policies and procedures need to be simple, effective and enforceable, otherwise they are unlikely to be followed consistently as part of an everyday routine.

Most security breaches occur through carelessness, such as leaving a cellphone or laptop in a public place. Good physical

security combined with effective communications security techniques should minimise the risks.

### Performance

The two governing factors for performance are latency and bandwidth:

- Latency determines response times. The goal is minimum latency, which is determined by the total end-to-end delay through all the equipment in the path.
- Bandwidth determines throughput. The achievable maximum bandwidth is determined by the capacity of the smallest link in the path.

These two factors are further complicated by the issue of contention, where multiple applications end up competing to use the same physical paths at the same time.

The main performance issue that needs to be addressed is the behaviour of the various infrastructure components under peak load. Intermittent saturation can be very difficult to diagnose and resolve. Quality of Service (QoS) measurement tools can help to both manage and diagnose such issues.

The way an application works can have a significant impact on the need for high-bandwidth and low-latency data networks.

For example, say a desktop client at a remote site needs to perform an SQL query across many fields in a relational database. It is far more efficient to perform the entire query within a 'query server' closely associated with the data and simply pass back the results, than to drag all the data across the intervening network and perform the query locally on the desktop client.

With storage subsystems a similar effect applies – an array of smaller fast disks will generally outperform an array of larger disks, even though the overall capacity is similar.

There is greater parallelism in the array using a larger number of smaller disks so the array



Colin Butcher: most security breaches occur through carelessness

controllers can take advantage of that to increase the overall performance.

With disk arrays, performance is generally related to the number of spindles and the amount of cache memory in the array controllers. Do not make buying decisions based purely on how many disks will give you the capacity you require.

Similar arguments apply to server systems. You can increase parallelism with multiple controllers and processors in order to increase throughput and decrease latency effects.

Understand the workload your applications impose on the systems and networks, then buy equipment accordingly.

### Voice and data

WAN links are the most likely part of a network to fail and they are usually under the control of another supplier organisation.

To minimise risk, it's best to have dual links between sites, with each link taking a

physically separate path. You should be careful to avoid having apparently separate links, only to find they join up with a common service provider infrastructure somewhere between the sites you're connecting together.

Voice over IP is becoming a commonly used technique, usually because it's apparently cheaper to use the data network bandwidth for voice calls as well.

It can also simplify call handling and management. However, be careful that you have sufficient bandwidth to avoid saturating your network links – especially when paths fail.

The first thing that happens when there are problems at a remote site is the people there pick up the phone to call for help.

If the network is saturated due to a virus attack or there is no spare capacity, then the scope of the problem can quickly escalate and trigger other problems.

### Strength in depth

Segmenting the technical infrastructure will help to protect your core business. There are many advantages to segmentation:

- Avoiding a 'monoculture' which is prone to common failure modes – there are risks involved in having everything on a common system and network platform if those systems are not rock-solid and well protected.
- Improved reliability and availability by damage limitation – in the same way that ships are compartmentalised using watertight bulkheads. Don't forget disaster recovery and business continuity planning: that's your lifeboat!
- Improved security through multiple security zones – just like mediaeval castles had layered defences.
- Improved performance through traffic separation – just like the new M6 toll road past Birmingham.
- Improved monitoring and

management through reduced complexity – the scale and scope of each zone is significantly less than that of the entire infrastructure.

● Improved scalability and ease of growth – by building your systems in a modular fashion that fit with your business structure and technical architecture.

All of these benefits come from an underlying architecture that defines a segmentation structure appropriate to your business.

### Conclusion

The rewards of successfully implementing an enterprise application can be significant, but there needs to be a sound justification and realistic assessment of both the business and technological aspects.

You will incur significant costs and risks if you are forced to make major, unplanned and rapid upgrades to your underlying technology infrastructure when you are part-way through the implementation.

The worst kinds of failure are usually performance related where the response times and connectivity become intermittent, typically because various components within the infrastructure have become saturated.

As always, the key to avoiding such problems is to carefully assess what you have, then plan and design for the future evolution of your business (see Table 1).

It is essential to consider the impact of the new systems on your existing infrastructure and business processes. You need to manage the transition and ensure you have sufficient technology infrastructure and staff capacity in place before undertaking any major implementation stages.

● *Colin Butcher is technical director at consultancy firm XDelta Limited.*

*Tel: 0117 904 8209. Email: colin.butcher@xdelta.co.uk.*

TABLE 1: Key areas to consider

ENTERPRISE APPLICATION LAYER	Need	● What do you wish to achieve?; ● What are you likely to want to achieve in the future?; ● Minimum functionality required to meet your requirements
	Usability	● Consistent behaviour across components of overall package; ● Lack of complexity; ● Security
	Deployment	● Backwards compatibility, including mixed version operation during upgrades; ● Ease of installation and automated distribution; ● Co-existence with other software; ● Monitoring and test facilities
	Installation	● Ease of customisation to match business characteristics; ● Ease of data import from previous applications; ● Cost of base system and additional modules; ● Hardware and software prerequisites; ● Hardware and software inter-dependencies
'MIDDLEWARE' LAYER	Provides end-to-end services for the applications	● File access; ● Data access; ● Message passing; ● End-to-end security (access control, authentication, encryption, auditing etc)
NETWORK LAYER	Performance	● Bandwidth; ● Latency; ● Contention; ● Speed of failover
	Segmentation	● Switch configurations (eg, VLANs); ● Protocols; ● Addressing schemes; ● Main data flows
	Availability	● Single points of failure?; ● Monitoring; ● Maintenance and updates; ● Speed of failover
	Physical cabling	● Signal quality: distance and physical environment; ● Colour coding and labelling; ● Requires good physical security
	Security	● Authentication; ● Access control; ● Auditing; ● Encryption; ● Firewalls
	Wireless	● Antenna placement; ● Disable SSID ('stealth mode'); ● Encryption (WEP); ● Authentication; ● Firewall separation from main network; ● Visitor zones
OPERATING SYSTEM LAYER	Availability	● Clustering (shared everything); ● Monitoring; ● Updates; ● Speed of failover; ● Disaster recovery planning; ● Business continuity planning
	Security	● Resistance to attacks; ● Authentication; ● Access control; ● Auditing
	Performance	● Throughput; ● Response times; ● Peak load behaviour
SYSTEM HARDWARE LAYER	Performance	● Parallelism; ● Throughput; ● Response times
	Expansion	● Partially populate a big box; ● Multi-processor systems; ● Partitioning capability
	Availability	● Hot-swap components?; ● Build quality; ● Component redundancy (eg, power supplies)
STORAGE LAYER	Performance	● Bandwidth; ● Latency; ● Contention
	Segmentation	● SAN fabric configurations; ● Device addressing schemes; ● Main data flows
	Availability	● Single points of failure?; ● Monitoring; ● Maintenance
	Data integrity	● Device naming; ● RAID (0+1 is best – shadowed and mirrored)