



White Paper

metagroup.com • 800-945-META [6382]

May 2005

David Willis

The Impact of Emerging Applications on Network and Security Infrastructure

A META Group White Paper

"Infrastructure is being stretched in new ways, requiring more intelligence and more efficiency in the network. Corporate networks will evolve well beyond basic connectivity, to be smart, secure, and survivable."



METAGROUP

Contents

Executive Summary	2
The Impact of Emerging Applications on Network and Security Infrastructure	2
The Distributed-Network Challenge	3
Emerging Enterprise Applications and Initiatives	4
<i>Typical Critical Enterprise Applications</i>	4
<i>How ERP Affects Network Design</i>	5
Anticipating Application Growth	7
Implications for Security Infrastructure	8
The Need for Security Intelligence	9
Overall Impact on Secure Network Design	10
No Performance Compromises	12
Bottom Line	13

Executive Summary

The corporate IP network has become the platform on which most businesses function. As a company becomes more dependent on a reliable and secure communication infrastructure, that infrastructure must evolve to keep pace. Security threats are also evolving rapidly. The network will evolve well beyond basic connectivity to be smart, secure, and survivable.

The following are key components of smart network infrastructure:

- A security services layer with centralized policy management
- The ability to evolve the network over time and customize installations to the local environment
- Support for a wide variety of applications
- Applicability to private networks and the Internet
- Remote management
- Suitability to large and small deployments
- Cost effectiveness
- Rich services with no compromise on performance

Here, we describe the challenges of new network applications and the implications for advanced networks.

The Impact of Emerging Applications on Network and Security Infrastructure

Business thrives on communication. It is indisputable that companies are dependent on their data networks to provide customers with applications and data. But new requirements are stretching infrastructure in new ways, requiring more intelligence *and* more efficiency in the network.

IT systems exist only to further the objectives of the business. Leading organizations are attempting to align IT capabilities, costs, and goals directly with those of the business — starting with the corporate network. The effort begins by executing more efficiently. As an example, companies have found that supporting multiple networks is highly inefficient and seek a converged infrastructure to

handle all traffic types. By the end of the decade, the migration to a common voice and data network will be underway in almost all companies and completed in 40% large and midsized enterprises.

The new intelligent IP network operates in a cost-effective and secure fashion, providing convenient, simple, and reliable services to its users. Security will be integrated into the overall infrastructure, not simply bolted on by adding a device here or there. It is necessary to create a logical security service layer using enforcement points (e.g., routers, firewalls, intrusion detection) along with base infrastructure, but under the control of a central policy. *To adapt to current and future needs, network and security infrastructure must evolve into a unified, multiservice platform.*

Why now? Let's start with overall IT trends. Infrastructure is consolidating as organizations attempt to improve efficiencies. Data center services, along with applications, are frequently centralized. But at the same time that applications are being hosted centrally, users are becoming more distributed; instead of working only in large facilities with hundreds or thousands of workers, they are also working from smaller offices, from home, and wherever they may find themselves — and they need network-based applications to get their jobs done. Network survivability has also become important, due to the impact of world events and infrastructure failures (e.g., the US power grid failures of 2003). This has led corporations to be more concerned about business continuity in the face of the unexpected.

Putting these themes into the context of the network, it is clear there will be more users, more devices, more applications, and more ways of connecting than ever before. Applications and application components will interact via the network in new, often unpredictable, ways. Network systems must be adaptable in various deployment scenarios.

The Distributed-Network Challenge

Network management must adapt to remote computing, with many services offered via the wide-area network (WAN). Branch offices are a special challenge for network administrators. Although the desire to cut cost has led to centralization, solutions must match local needs. There may be some degree of local specialization, but the solution must still enforce a central policy. Further, systems must support a wide variety of underlying technologies and services. The optimal solution for a large office is different from that of a large headquarters location, and different still from the small or home office. Geographical variations in WAN services can further complicate things. The carrier service of choice in Europe could be substantially different from the best selection for the US, Japan,

or Asia; some sites might be attached over the Internet, others over frame relay, and others using carrier MPLS services. Access technologies also vary by location (e.g., T1/E1, DSL, Ethernet, wireless). The trend toward personal mobility only increases the variation of devices and access mechanisms. Whatever the connection and the local configuration, remote devices must be securely manageable over the network with a minimum of local expertise.

Emerging Enterprise Applications and Initiatives

As networks converge, they are becoming more aligned to application needs. Leading businesses understand just how critical communications are to their business — not simply a cost to be controlled, but also an enabler of new business. In fact, 62% of Global 2000 companies view their networks strategic assets. Without a strong core network, users have no applications; without applications, the businesses cannot function. Devices must be able to recognize which applications are most critical and handle them according to business policies.

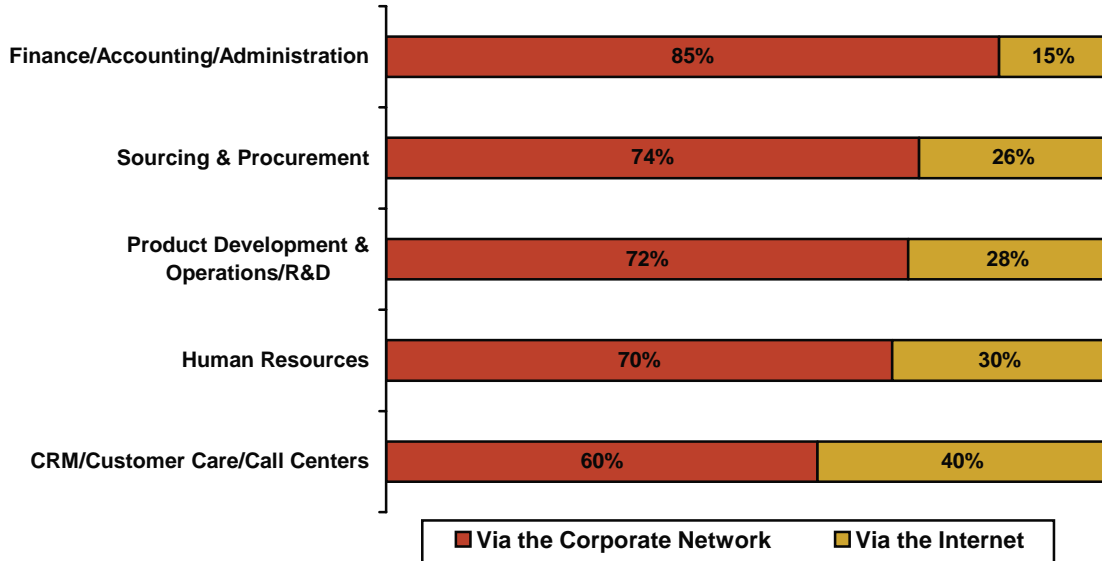
An understanding of the most essential business applications will help illustrate the challenges.

Typical Critical Enterprise Applications

Enterprise resource planning (ERP) systems span a wide range of functions that are mission-critical to companies — from manufacturing and logistics, to distribution and inventory, to back-office functions like accounting and human-resources management.

These applications are no longer confined to the traditional intranet. Increasingly, these applications are exposed both inside and outside the firewall (see Figure 1).

Figure 1 — Network Use Inside vs. Outside the Firewall



How ERP Affects Network Design

Network and Security Implications of Expanded ERP Systems	
Activity	Impact
Increased reliance on ERP applications	High-availability design No single point of failure User performance measurement
New ERP users outside the firewall	DMZ designs must allow secured access to critical systems while protecting data <ul style="list-style-type: none"> • Better user authentication (<i>who the user is</i>) • Better user authorization (<i>what the user can do</i>) • User activity auditing • User life-cycle management

Customer relationship management (CRM) systems enable companies to more effectively capture and understand their customers' and competitors' behavior. Not only do CRM solutions span sales force automation systems and marketing functions, but also they increasingly integrate to customer services and support roles as well — where a large number of users exist outside the firewall, and infrastructure must be secured appropriately.



The Impact of Emerging Applications on Network and Security Infrastructure

Network and Security Implications of CRM Systems	
Activity	Impact
Sales force automation	User is frequently disconnected <ul style="list-style-type: none"> • Remotely managed policy controls • Desktop hygiene controls • Distributed access; variable-access scenarios • Access network may not be trusted (e.g., Internet WLAN hotspot) • Device may not be trusted (Internet kiosk, home PC)
Field force automation	User is frequently disconnected <ul style="list-style-type: none"> • Remotely managed policy controls • Desktop hygiene controls Support for advanced file types (images, video)
Customer service	DMZ design to protect access to critical systems Web-based portals with user personalization Management of thousands of users at multiple trust levels User activity tracking Privacy controls

Knowledge worker infrastructure is also evolving dramatically. New platforms dramatically improve productivity, supporting virtual teams and other models. This approach often incorporates real-time collaboration (RTC), which provides synchronous communication services between groups of users and between users and real-time-enabled applications (e.g., alerting, online meetings, virtual classrooms). RTC uses interaction models that integrate audio, video, data, and other natural interfaces (e.g., voice, speech, gestures).

Although the benefits of these systems are becoming clearer, the combined impact on the network is not as well understood. Many new collaborative systems such as those from Groove, Kontiki, and Skype use a peer-to-peer model — presenting challenges for both network provisioning and security. Real-time voice and video also deserve special attention. In particular, latency must be bounded and predictable (in technical terms, jitter should be low). Voice and video also create a large number of very small packets, which can stress legacy equipment such as routers, firewalls, and switches. Video can require substantial bandwidth, and both voice and video require a method of arbitrating access to limited network capacity, typically in the form of quality of service (QoS).

Network and Security Implications of Knowledge Worker Infrastructure	
Activity	Impact
Collaboration platforms	Peer-to-peer traffic flows Must be easy to add incremental network capacity
Rich media Real-time media	Support for voice and video streams <ul style="list-style-type: none"> • Latency controls • Jitter controls • QoS • Bandwidth for video Support for advanced file types (images, video) Policy controls (by user, time of day, etc.)

Anticipating Application Growth

Although these examples illustrate the most typical of applications, each network is different. Whatever the application mix, the pace of internally developed applications is accelerating. New approaches to application delivery, such as service-oriented architecture and grid computing, radically change the flow of traffic between points on the network. Network technicians are in constant catch-up mode and have barely enough time to build what is needed, let alone understand application flows and optimize them.

The Impact of New Applications on Network Infrastructure

As a result, infrastructure must become more aware of applications, to allow administrators to do the following:

- **Define a policy** that may be understood and controlled by business priorities and implemented in the infrastructure
- **Identify application traffic**, including applications not already known to be in use
- **Determine the appropriate policy** to handle the application
- **Map policies into traffic classes** supporting the specific behavior
- **Monitor and report performance levels**

Network and Security Implications of Application Growth	
Activity	Impact
Rapid application development	Less pre-deployment visibility into the network <ul style="list-style-type: none"> • Network must discover applications • Network should optimize applications
Service-oriented architecture	Non-deterministic traffic flows
Data center growth (consolidation, with proliferation of Intel-based servers)	Non-blocking switch architectures Isolation (firewalling) between business units Security domain controls within the data center

Implications for Security Infrastructure

From a security perspective, the most significant implications of emerging applications are that they are new and that they often diverge from standard, well-behaved communications patterns.

- ***New applications.*** The implication is that existing network and security devices lack the specific, detailed knowledge to properly control and protect the application itself as well as the environment in which it resides.
- ***New communication patterns.*** Peer-to-peer and Web services technologies put an end to there being a consistent two- or three-tier communication structure between clients and servers. Wireless networking enables numerous network attachment points — some inevitably unprotected — versus a handful of well-defined ones that are outfitted with requisite security controls (e.g., authentication, malware filtering). In both cases, the point is that the luxury of having a relatively small number of natural chokepoints is disappearing.

However, it is not just the proliferation of new/different applications that is having an impact on information security solutions. Another major driver forcing change is the rapidly evolving threat landscape. Because much has been written about this issue in recent years, we will summarize only a handful of the most relevant observations:

- Outside target-specific hacking, the most serious threats are **self-propagating malware** that incorporate multiple attack mechanisms (i.e., so-called "blended threats").
- Despite the best efforts of security vendors and users alike, **malware and other forms of attack continue to be alarmingly effective at eluding currently deployed safeguards.** More than 75% of organizations were affected in 2004.

- **The vulnerability-threat window is continuing to shrink.** The trend is quite obvious — with MSBLAST trailing notification of the RPC DCOM vulnerability by only 26 days and then the Sasser worm trailing the LSASS vulnerability by just 17 — and as a result, true day-zero exploits appear inevitable.
- **Threat propagation times are accelerating, and the impact is broader.** For example, the Slammer worm doubled its infection count every 8.5 seconds and achieved 90% coverage within 10 minutes. Broadband adoption rates and the increasing degree of global interconnectivity, coupled with persistent innovations in threat design, will ensure this trend continues.
- **The "bad guys" behind the threats are increasingly being motivated by money,** not merely status among their peers. Better funding, focus, and organization behind their efforts will only bolster them.

The Need for Security Intelligence

The result of all this is the need for a security system that is more intelligent (i.e., capable of accounting for new applications as well as new, sophisticated threats), more comprehensively deployed (i.e., leaving no open paths), and more efficient (i.e., automated, to both reduce administrative effort and keep pace with fast-moving threats). In other words, these conditions result in the need for a security infrastructure that exhibits the following characteristics:

- It must perform two distinct, primary functions. First, it should **govern which traffic is allowed to flow** on the basis of pre-established policies that define acceptable usage of the computing resources. Second, it must also examine this traffic that is "allowed by policy" to ensure that it does not contain attacks of any sort — and if it does, then it must ensure this otherwise-allowed traffic is also barred from the network.
- To be effective at both its usage control and threat-prevention functions, it should have at its disposal **knowledge at multiple layers**. Operating solely based on information at the network layer, as has been the situation in the past, is no longer sufficient. Going forward, the security infrastructure should have visibility into and **control at the application and user layers**, in addition to the network layer. It is important to recognize that the user layer includes not only the user's identity (and how that relates to specific roles and entitlements), but also the security state of the client device the user is utilizing.
- It must include components that conduct the primary functions and that are **suitable for deployment throughout the entire computing environment** —

not just at a handful of critical chokepoints. This means supporting various form factors, capacities, and price points to ensure pervasive coverage — from remote client stations, to Internet boundaries, to subnet/business unit boundaries, to data center boundaries, etc. — is achievable.

- Information and responsibility for responsive actions must ultimately be shared among the individual components comprising the infrastructure. In other words, it must begin to **operate as a cohesive system**. This is fundamental to achieving an even higher degree of effectiveness, as well as significant operational efficiencies.

Overall Impact on Secure Network Design

The intelligent enterprise network will have the following characteristics:

- **Evolvable.** The common theme among all the application scenarios and security threats is change. Infrastructure must be designed so that the impact of change is minimized; modularity is critical. Software updates must be easily applied. Requirements can change quickly; the network manager may choose to change to a different deployment scenario after only a few years into the useful life of the equipment. Rather than deploying new equipment, systems should be expandable and remotely reconfigurable. New security policies should be applied automatically as the deployment scenario changes.
- **Integrative.** Systems must be minimally disruptive. New platforms must interoperate with existing infrastructure and new platforms must conform to industry standards that enable easy integration.
- **Securable.** Systems should be secure by default, enabling the administrator to gradually open up access as necessary. Security policies should be easy to create, monitor, and audit.
- **Manageable.** Networks devices should be administrable on a systemwide basis, rather than as a collection of discrete parts. Local installers should be able to place a minimally configured unit to a remote site, from which point a remote administrator may complete the installation to fit corporate policies and local requirements.
- **Best fit.** Finally, network infrastructure must be able to match the appropriate level of service required. IT organizations must balance any potential loss to the business against the cost of building and maintaining a communications network built for high availability. High-availability features such as redundant equipment, dual connectivity, hot swappability, and failover capabilities at

higher layers may be essential for critical services. Conversely, the cost of high availability should not be levied on installations that do not need these features.

Network devices for emerging applications should feature the following:

- **Various form factors:** The ability to size devices appropriately based on the needs of the local office.
- **Integrated security:** Network functions and security must work together, not independently. A centralized, policy-driven infrastructure should span the network and security layers.
- **Support for multiple WANs,** such as Internet VPN, MPLS, frame relay, and private line.
- **Remote installation:** Devices should be configured over the network. Secure configuration via Web and telnet interfaces, with strong authentication options. Support for polling and alerting via SNMP.
- **QoS mechanisms** that enable voice, data, video, and multiple data type priorities. QoS policies should enable applications to share bandwidth and limit (or block) undesired applications. Devices should enable the mapping of QoS policies across Layer 2 and Layer 3, supporting multiple QoS signaling mechanisms (VLAN tags, Diffserv, and RSVP).
- **Policy-routing mechanisms** that enable various traffic types to be selectively directed over multiple links as available.
- **Reliability mechanisms** that eliminate single points of failure. Failover options such as redundant connections and dial on demand.
- **Administrative flexibility,** supporting both centralized management models where all traffic is brought into a common location, as well as distributed models with remote management.
- **Network addressing options** allowing administrators the freedom of using various IP addressing techniques, such as local Net10 addresses and network address translation. Ability to handle multiple conflicting IP addresses systemwide.

- **Support for multiple VPN tunnels:** Systems should allow voice, video, and other traffic streams to be independently tunneled.
- **Real-time traffic management:** Support for secure admission and transport of voice and video traffic (H.323 and SIP) without compromising data security.
- **Flexible routing protocols,** using industry standards such as RIP and OSPF.

Bringing Networks and Security Together

Given the economic constraints most organizations face, and the growing aversion to adding "yet another box" (or piece of software) to the environment, the ideal situation would be for all these security capabilities to be coincident with the network infrastructure. Although it is unrealistic to expect a thorough merging of network and security infrastructure to occur overnight, the evolution is inevitable and has already begun.

Organizations must seek products that exhibit characteristics indicative of an architected, system-based approach to information security.

An often-overlooked aspect of security management is the need for instrumentation with the security administrator in mind. Security administrators should be able to define policies and verify that they are being followed, using the "separation of duties" principle. Another compelling approach is in role-based administration, allowing for the partitioning and delegation of management under a common overarching policy.

No Performance Compromises

Network solutions must be able to provide security without affecting performance. They must be designed to allow a gradual evolution of services over time without upgrade, rather than simply being "dumb pipe" infrastructure. Solutions must also be secure right out of the box and configurable over the network on an individual unit or systemwide basis.

Best-Fit Capabilities for Network Security	
Enhanced security	<ul style="list-style-type: none"> • Application awareness and control (antivirus, URL filtering) in addition to network-layer control (encryption, authentication) • Multilayer attack protection • Segmentation and compartmentalization (VLAN, VPNs) • Admission control and quarantining
Enhanced performance	<ul style="list-style-type: none"> • Efficient inspection models and mechanisms • Specialized hardware (e.g., accelerators)
Enhanced management	<ul style="list-style-type: none"> • Scalable and flexible policy models • Delegated administration • Integration with perimeter controls • Simple, automated deployment (real-time and offline configuration, GUIs, and command lines)

Bottom Line

Network infrastructure must become more aware of users, devices, applications, and activity and be evolvable, integrative, securable, and manageable. Network and security infrastructure must become more integrated; the days of "bolting on" a security solution are ending. Infrastructure should be designed with change in mind, readily adapting to new requirements without high-impact upgrades to software or hardware and radical redesigns.