

Managing Enhanced Network Services: A Pragmatic View of Policy-Based Management

John Vicente, Information Technology, Intel Corporation
Harold Cartmill, Information Technology, Intel Corporation
Glen Maxson, Information Technology, Intel Corporation
Shelby Siegel, Information Technology, Intel Corporation
Russ Fenger, Intel Architecture Labs, Intel Corporation

Index words: policy, policy-based management, PBM, QoS

ABSTRACT

The convergence of public and private networks and the rise of the Internet as a global medium for information exchange and economics are prompting corporations to augment existing business computing models. New information technology initiatives for collaboration and business exchange are essential to gain competitive advantage. These initiatives are being realized through emerging applications (e.g., e-Commerce, groupware applications, multimedia) over converging private and public boundaries. Enterprise strategies are requiring timely evolution of network infrastructure and management to support delivery and management of end-user and network services. Within this context, quality of service, security, productivity, and infrastructure efficiency are critical to achieving bottom-line business results. In this paper, we take a closer look at policy-based management as an enabling technology and paradigm shift for Intel's Information Technology organization. We explore the dimensions of policy-based management, and we provide a pragmatic review of the technology, discussing the deployment challenges, roadmap considerations, and practical usage scenarios.

INTRODUCTION

Current trends in corporate and Internet networks are shifting from best-effort, vertical network architecture towards a more intelligent, end-to-end, service-aware network paradigm. As evidenced over recent years, the need for enhanced network services such as virtual private networks (VPN), quality of service (QoS), security, collaboration, and directory technologies

demonstrates that customers are demanding more from the core infrastructure for enabling productivity, flexibility, service differentiation, isolation, privacy, and manageability. Moreover, critical network resources must be aligned with business objectives where networks are i) more content or application-aware; ii) provide dynamic features for service creation; iii) observe and enforce network-wide policies; and finally, iv) enable control from the network provider to the administrator to the end-user. The migration to a richer network infrastructure allows corporations to be more agile and optimize infrastructure costs, while meeting the diverse requirements of emerging application demands. The following requirements are driving innovations in current network infrastructure technology.

Mission Criticality

IT organizations recognize that the availability and reliability of network infrastructure are essential to critical applications and services that rely on them. These same applications can compete for network resources (e.g., bandwidth) with other less critical and diverse applications for successful transport delivery and performance. Core network capabilities are required to ensure delivery priority, security, access control, and fair performance allocation. However, certain users (e.g., company presidents) or groups in an organization may have a more critical need for access to resources, and thus, that person or group may get priority or have a different authorization from the rest of the organization. Mission criticality is raised as organizations move towards extranets or public services for corporate business computing or service transport.

Network Architectural Agility

Customer demands and applications are growing rapidly while networks continue to be gridlocked by standards and proprietary implementations. In general, enterprise networks have been implemented with multiple vendor device solutions, while heterogeneity of bandwidth resources and legacy device inequalities are evident over geographically dispersed regions. While such infrastructure complexity exists and diverse application-network demands continue, network architectures and management frameworks must still support more rapid evolution, enabling faster service creation and deployment while maintaining legacy integration.

Service Flexibility

The application and network infrastructure need to be integrated to meet the ever-changing needs of current computing and distributed application demands. Providing dynamic network reconfiguration, timely access control, or dynamic, class-based bandwidth allocations can provide greater flexibility to the end-user or network administrator. An increase in network features allows administrators to manage service levels more effectively, while allowing application developers more control features for end-user services.

Efficiency

The increased complexity of managing network systems along with the need for service-level management requires current network management solutions to be more sophisticated. While corporations continue to grow, optimal planning and management of the distributed infrastructure are essential. Better tools for Just-in-Time (JIT) bandwidth and service provisioning, effective use of resources, and automation of management tasks can reduce the total cost of ownership while improving service-level management.

In the first section of this paper, we present the background on policy-based management including our view of the requirements for this technology. We follow this with a discussion on the evolving IT business and operational models. Next, we discuss the technology implications and challenges that IT organizations must address to realize policy-based management. Following this, we present our current progress within Intel IT, providing a transitional discussion under the e-Business model. We close with a summary and a brief look at how we can move forward with this technology.

BACKGROUND ON POLICY-BASED MANAGEMENT

As defined in [1], policy-based management is "the combination of rules and services where rules define the criteria for resource access and usage." Alternatively in [2], the authors define it as a "unified regulation of access to network resources and services based on administrative criteria." We view policy-based management as a viable technology to provide greater control and management of underlying networks via the creation and distribution of high-level policies (business rules), integrated with the enabling mechanisms of the network infrastructure. By way of automated and rapid configuration and the integration of business policies with the network infrastructure, new opportunities for managing both infrastructure and network services are introduced. To support these new initiatives, we define the following general requirements for policy-based management technology:

Service differentiation. This is the ability to control or manage the quality of the service or service delivery mechanisms in order to meet some predefined network-based performance delivery/metrics. This may extend to enabling Service Level Agreement (SLA) management for service-level validation.

Network provisioning and bandwidth management. These provide proactive bandwidth management by facilitating control and allocation of bandwidth through device configuration management: that is, facilitating manual, multi-device network configuration and performing admission control or traffic segmentation.

Integration with network management systems and legacy devices. Policy-based management must be integrated with current paradigms for managing IT organizational structures. These include existing operational models (e.g., centralized control or change management), security requirements, and business computing models. The requirement to support or address legacy systems and device limitations is also mandatory.

Scalability. The PBM technology infrastructure should architecturally scale to Intel's enterprise environment and private/public models for e-Business computing. These include the policy server environment to business computing hierarchy, the directory and database infrastructure, and scalable policy overhead in terms of administration and protocol communications.

Industry standardization. This means that policy-based management must conform to industry standards and use best practices to support network device and policy management interoperability. There must be standards for such things as policy terminology [1] and protocols (e.g., COPS [3] / LDAP[4]). Moreover, an open framework for policy management schema [5] and directory integration must exist.

- *Security.* The policy-based management tools should facilitate resource access control and authorization, and they should provide integration support for authentication and accounting.

Policy-Based Management Framework

Functional Overview

The components of a policy-based management system (PBM) include the policy console, policy server, policy database, and policy clients, all of which are shown in Figure 1 below. The policy console provides administrative and operational access to the PBM system.

The Policy Decision Point (PDP), or policy server, embodies the decision-making functionality of policy-based management. One or more such policy servers exist in a control domain, with each server configured to support policy management for some defined group of policy clients or Policy Enforcement Points (PEP) in the domain. The policy server provides each policy client with policy information; the policy client in turn carries out (enforces) the policies to the best of its abilities. The policy server's inner structure is shown in Figure 1 below.

Policy Server Structure

The policy client communication component handles all exchanges between policy clients and the policy server. The PBM client-server communications protocol uses the (proposed) IETF standard COPS protocol.

The message processing component is responsible for policy protocol message decomposition and composition, and for interpretation of wire-format objects for use in the policy server.

The core-processing component embodies the logic needed to support the policy server's policy decision making—rule processing and housekeeping. This component also logs all relevant usage by policy clients in support of accounting and billing applications. These logs may also be used for further tuning of resource control policies.

The policy console communication component gives support for communicating with the policy console. This allows multiple PDPs to be maintained from one policy console. It also allows multiple consoles to access the same PDP.

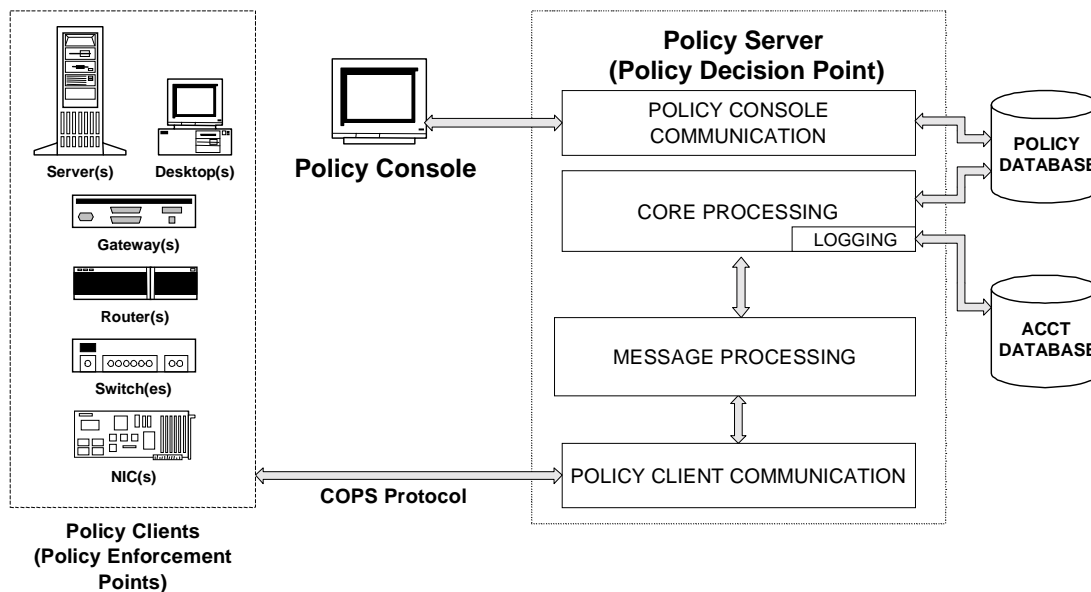


Figure 1: Policy server structure

Implementation Alternatives

There are alternatives to using a policy-based network management server as described above. For example, one can manually configure policy clients to perform policy actions on packet flows, such as Type of Services (TOS) bit settings in the packet header that tag packets for priority Quality of Service (QoS) or client-initiated

RSVP (ReSerVation resource Protocol) [6] sessions. However, one of the key functions of the PBM server is the centralized management of network resources and the allocation of resources to clients based on organizational rules (the organization owning and managing the network resources.) This ensures that only authorized clients have appropriate access to and

use of network resources. This function is extremely difficult to manage solely at the client.

There are other approaches that use directories and directory protocols, such as Lightweight Directory Access Protocol (LDAP) to set up and administer policy. These approaches work well when configuring network devices and setting static rules. However, when dynamic control of network resources is required, directories are ill-equipped to deal with the complex decision-making process, based on current resources in use, priorities of requests, and administration of usage. Also, the LDAP protocol does not currently provide a robust level of information or a dynamic framework to manage network resources in real-time.

Intel Architecture Labs (IAL) is helping drive the networking industry to fully use IETF's Common Open Policy Services by making available the COPS toolkit, which enables clients to talk to policy servers. IAL has developed and is deploying a COPS Local Policy Module that enables Windows 2000* clients to talk directly to policy servers.

EVOLVING IT BUSINESS MODELS

Current trends in business computing are blurring the boundaries between the Internet and Intranet, while application demands are becoming more content rich and diverse in quality of service requirements. Network vendors are building infrastructure components (e.g., VPN) with more enabling network features (e.g., QoS) and management systems to support these new users or applications. Nevertheless, it is unclear how IT organizations should structure or restructure their business processes to map business units, application types, and transactional priorities to critical infrastructure resources. If this new paradigm is to be realized, policy-based management must emphasize and speed up the development of new or the revision of existing business processes. For IT organizations to move towards a more service-oriented and evolutionary computing model, processes will have to be developed to manage the integration of business objectives with network and distributed infrastructure through policy. Figure 2 depicts this integration with a proposed hierarchy (see also [2]) with which policy is introduced, translated, and propagated within the network infrastructure. At the highest level, business rules should dictate global (i.e., domain-specific) directives on the effective use and priority of resources to support

the business objectives. These objectives are translated into network-wide policies within the administrative domain on the necessary bandwidth, communication resources, and topology requirements. Network-wide policy rules are translated to node policy rules that are specific to the required behavior of the node to manage its local resources. Finally, local policy rules are applied and enforced through one or more specific command instructions on device-level functions (e.g., scheduling and queuing parameterization).




Figure 2: Integrating business rules with network policy

We propose the following methodology from which policy-based systems can be implemented as an aspect of the management process within a corporate enterprise environment or managed network domain:

1. Create a network baseline and track network usage against key applications, network services, and business unit users or usergroups.
2. Establish network domains, policy groups, group identities, and hierarchies that map to core business activities.
3. Establish organizational directives and create corresponding policy rules and service-level requirements over various applications supporting policy groups/users on the allocation or priority use of critical infrastructure resources.
4. Administer and deploy policies across the network infrastructure, typically on a domain or inter-domain basis.
5. Audit and validate network policies against service requirements.
6. Refine business directives and network policies based on policy-enforced behaviors.
7. Repeat steps above.

Operational Models—Shifting to Policy

The operational management models of today's IT environment are based on traditional models of systems and network management. These evolved from a

* Other brands and names are the property of their respective owners.

central, mainframe-orientation to a client-server and distributed systems management paradigm. A current challenge, both for IT management and solution providers, is managing both the end-to-end service model and the vertically disjointed infrastructure elements. With recent developments in network QoS capabilities and policy-based management, this challenge may be met. Operationally, however, service-level management and change management processes will need to evolve or adjust to support a policy framework. Furthermore, training of operational personnel on business-oriented policy management, network QoS/CoS, and the supporting management tools will be a requirement. These are discussed in more depth in the next section.

Change Management

A critical component of the IT operational environment is understanding infrastructure changes and managing operational schedules. For example, avoid changing a network the same day software is updated for a large population of servers or clients. Generally, the change process tends to have a high overhead in terms of administration and personnel, which is complex to manage and difficult to administer over a large enterprise environment. As we move to a policy-managed environment, using existing IT change management processes will require greater interdependency and be more complex and costly. Thus, we anticipate an evolution in the administration or change management process.

The existing change management system is built on IT being the direct provider of all IT services. The current model lacks operational flexibility or agility. Policy-based management, on the other hand, will allow organizations to map the network to the business requirements, thus changing the role of the IT organization to that of a coordinator for service provisioning. This is an evolutionary step for traditional IT organizations as the change process will directly engage customers in the change processes. Furthermore, IT project managers are generally focused on managing the customer and the expectations for their business unit, application, or area of focus. This approach works but has unforeseen consequences for the enterprise network and end user when a customer's application is fully turned on. The PBM environment should motivate business groups' project planners to drive changes in the IT project office to consolidate planning activities. The evolutionary process will require project leaders to coordinate their activities to engage stakeholders and to ensure that their application/service is properly prioritized and business rules are communicated and tested. Engaged parties

responsible for sustaining business processes must understand the implications of specific agreed policies on the business environment.

As a policy is implemented, the project manager and operations change bodies must tightly integrate business requirements and the policies deployed within the operational environment. This requires cross-training personnel on the integration of business and policy management so that they understand the effects or ramifications of change policies within the infrastructure. We envision a consolidation of existing processes and tools through the integration of change and problem management with PBM systems. By streamlining processes for change and problem management groups, we will enable data sharing and true interdependence. Thus, we believe the generation of Customer Resource Management databases and tools should speed up and personalize the change/problem management system to provide customers and stakeholders a view of the relevant data that supports their activities.

Service-Level Management

Service-Level Management agreements (SLAs) are contracts between the provider delivering a service and the recipient of the service. The SLA codifies the understanding between the parties to ensure delivery of services and value for payment. A network/application-oriented SLA places value on service delivery of the key components of *availability, delay, throughput, customer service, and affordability*. A service-level policy is a method for controlling and regulating service differentiation. Together, service differentiation and service-level policy form an integral part of the service SLA function that has become increasingly important as TCP/IP networks evolve.

The shift to PBM will require IT organizations to have the tools and motivation to manage to tighter service levels on service performance, security, reliability, and customized agreements. This will require a more dynamic and granular auditing and reporting function for the underlying services and the managed environment. Additionally, the reporting function must be designed to refer to the negotiated SLA and provide a meaningful report that supplies the customer with validation that services are being delivered as agreed. Moreover, the customers must have objective proof of reliability and transaction responsiveness, service availability, as well as be able to rely on the operations supporting their core business.

TECHNOLOGY IMPLICATIONS

Policy-based management technology can be applied both to areas with obvious Returns on Investment (ROIs), such as bandwidth management (e.g., savings on WAN circuit costs) and to areas where the benefits are harder to measure, such as productivity for certain users. Since it potentially can require a large capital outlay to implement (i.e., if a large amount of legacy equipment needs to be replaced), policy-based management initially will be implemented as point solutions targeted to specific conditions with definable ROIs. In this section, we discuss the key barriers IT organizations have to clear and the key technology areas that they must address in formalizing an ROI and roadmap strategy.

Network Technology

Multi-Vendor Interoperability

The infrastructure of an IT organization is made up of technology from a host of different vendors. The type of management policies desired will determine which equipment will be involved in receiving and enforcing policy information. For example, if a company wants to give high priority to certain SAP R/3 users, it may have to configure the network composed of mixed vendor equipment, the SAP servers, the SAP application itself, the users' client systems, and possibly some auxiliary servers (e.g., database and directory servers) in order to make this happen. As mentioned above, because of the nascent level of policy-based management technology and products and the embedded base, multi-vendor issues will persist for some time. This issue is being exacerbated by the move to the Internet and e-Commerce, where policy-based management promises some of its biggest rewards but the multi-vendor issues abound.

Quality of Service (QoS) Technology

Policy-based management technology solutions and tools presently focus primarily on the enabling mechanisms for delivery of QoS and resource (bandwidth) management. QoS policy can be defined [2] based on some criteria including the endpoints of communication, route or communication path, community of interest or usergroup, application types, network or traffic characteristics, or specific time period. There are primarily two methods to support end-user QoS: signaled or provisioned. Within the IP communication model, RSVP with Int-Serv (Integrated Services) [7] performs signaling to ensure QoS on a per flow basis by using dynamic resource reservations. RSVP leverages policy-based controls to support admission control and flow regulation. Alternatively, Differentiated Services [8] or DiffServ operates on a slower time-scale. It is essentially a provisioning model

to reserve or establish service classes. The DiffServ model operates on a traffic aggregate basis where flows (one or more) are bundled together according to a set policy and treated based on a negotiated class of service. The administrative provisioning model is based on SLAs translated to Traffic Conditioning Agreements (TCA) and enforced through underlying mechanisms (e.g., classification, scheduling) within a router or IP device.

Multi-protocol Label Switching (MPLS) [9] and 802.1p [10] are both layer-2 protocols that can be used in isolation or paired with DiffServ or RSVP to deliver QoS on a flow QoS or provisioning basis. MPLS label switch routers (LSRs) are positioned (among other capabilities) to support the integration between ATM and IP by using Label Switching Path (LSP) protocols to map layer-2 ATM VCI/VPI identifiers with a layer-3 IP device to deliver (not exclusively) QoS within an ATM WAN core. Within the LAN environment, the 802.1p protocol enables priority or service classes within a LAN environment, where once again, mapping of layer-2 and layer-3 service classes or traffic precedence can be supported. The original scope of the 802.1p protocol was to support the integrated services model within the local LAN.

In current vendor solutions, early QoS/CoS capabilities to support these protocols/mechanisms exist, but there are many problems with their adoption. Int-Serv/RSVP has been plagued by the scalability issue where the claim is that maintaining per flow state within the Internet or a large enterprise network breaks the fundamental IP architectural model, which is based on end-system state maintenance and control. On the other hand, Diff-Serv has not reached standardization or industry maturity. This latter point is especially true when it comes to deployment, where the need to support inter-domain SLAs has not been fully hammered out. However, the concept of the Bandwidth Broker (BB) [11] has been proposed by the DiffServ community to support provisioning within intra-network domain boundaries and across inter-domain boundaries to deliver end-to-end quality of service. This work is still in the early stages and requires more investigation. Nevertheless, the motivation is there to establish standard building blocks towards enabling QoS within the Internet/IP model.

A broader analysis of some of the aforementioned services and protocols is presented in [12] including alternative QoS-supported models for Constraint Based Routing and traffic engineering. Alternative proposals leveraging the best features or motivations behind these protocols or services have also been suggested to deliver scalable, signaled, and provisioned QoS [12, 13, 14].

Within current IT environments, proprietary network device features that support service differentiation as well as legacy devices are variables in the ROI decision process. Support for legacy systems and proprietary devices is an essential aspect of our early QoS and policy investigation. This includes looking at traffic segmentation (e.g. VLAN switching) or access control (e.g., multicast filtering) to manage traffic propagation. Moreover, over-provisioning, especially within a LAN

environment, is a practical alternative in the short-term. Nevertheless, with increasing bandwidth requirements, QoS-dependent applications, and the move towards the Internet, the industry QoS models described previously may clearly be the favorable long-term choices. To deliver or manage QoS, the consensus is that policy-based controls must be integrated with device mechanisms (proprietary or otherwise) to support the provisioning, admission control, and regulation of traffic.

Network Management

As the demands for distributed and global computing increase, and Internet-based electronic businesses continue to grow, network growth (e.g., traffic volume, traffic types) and complexity (e.g. devices types and counts, network events, interoperability) increase along with it. With demand for innovative services (e.g. desktop video collaboration, knowledge-based management) extending current user communication models, the requirements for these new services will continue to increase with a corresponding increase in network growth and complexity. While this growth is essential for business and the evolution of information technology, it places rigorous demands on our infrastructure. IT managers are faced with the challenge of managing the infrastructure for optimal service delivery while at the same time reducing operational costs (e.g., manpower, operational tools). Finally, traditional network management tools are device centric and require manual configuration, which leads to duplication and task redundancy.

Because of these constraints and limitations, policy-based management is justified. PBM will abstract physical and virtual elements of the network that facilitate the automation of traditional network management tasks across multiple objects of the network. The automation feature lessens the need for human administration, thus speeding the change management process. Furthermore, abstracting network devices or components raises the issue of implementation or proprietary differences across network devices. This allows (under certain conditions) management interoperability and reduces the complexity associated with managing implementation-specific devices over alternative interfaces. However, management capabilities available in current policy-based management solutions are not broad enough to deal with the overwhelming activities associated with network management. Nonetheless, the focus on bandwidth management and QoS is certainly the right choice in the move to this new paradigm for managing networks.

Managing Network Services

A complementary view of the enabling capabilities of policy-based management is the notion of regulating or controlling distributed and abstract network resource objects in concert to deliver or manage end-user network services. Such enhanced network services include video-based distance learning, voice over IP (VoIP), quality of service, multicast services, and security. The service orientation of network management is a major shift from the traditional approach of managing networks. This has been enabled through the introduction of middleware [15] services and the corresponding abstraction or raising of the lower-level physical communication infrastructure. The notion of policy is introduced here as "the ability to administer, manage, and control access to the network resources of network elements in order to provide a set of services to clients of the network" [15].

Security Services

Interoperability

Corporations need to protect their business assets from competitors. This traditionally has been done with physical barriers and firewalls for information technology assets. E-Commerce and the use of the Internet to provide IT services (such as Web and application hosting) is blurring the line between inside and out, requiring that security be implemented on a more information-specific level. The IPSec standard, for example, allows information to be encrypted between sender and receiver, thus satisfying a privacy requirement. However, the use of this encryption may also hide pertinent information from any policy-enforcing infrastructure in the path between the sender and the receiver. Thus, some implementations of security services could eliminate the possibility of establishing other policies. We recognize that both standards groups and vendors are working to reduce or eliminate this conflict.

Policy-Driven Security

While the current focus of network-based policy management tools is on quality of service, the motivation behind policy work extends beyond the functional area within the network transport. Clearly, privacy, complexity management, dynamics, resource access, and authorization, etc. will require similar facilities or capabilities so that growth and traffic propagation in the context of a security policy can be managed. This is especially true for the e-Business model of computing.

However, just as the network vendors have been moving toward policy-based management so has the security

industry. Abstractions away from managing individual access controls on objects, management of heterogeneous environments, and training of personnel on multiple administrative products all support policy-based management initiatives in the security arena. We therefore see a parallel to the network PBM toolset in the applications and server space for policy-driven security. E-Business may provide synergy for merging future policy-based management tools in these traditionally separate areas.

There is an interesting perspective that arises primarily in the application/server space that divides the policy products in half. This comes in the positioning of management of policy versus execution (implementation) of policy. In the network PBM space, typically one expects that one vendor might provide the management tools, while another vendor provides the implementation, or run-time enforcement of policy. Given the underlying infrastructure of network management protocol convergence, it is reasonable to expect a network management toolset to have the capability to manage across a heterogeneous environment.

In the application and server PBM space, on the other hand, there is no standard for administration across a heterogeneous environment. As a result, administrative tools must either be built with interfaces to many proprietary targets, or they must provide a single centralized security server implementation, and expect the secured resources to query the security server in a common language. This split delineates two significantly different architectural solutions and bifurcates the product solution space. Intel's current architecture selection is the former one: a common policy-based management tool manages multiple heterogeneous targets, and the execution-time security is native to the proprietary target environments.

As we move forward with PBM solutions in the application and network space, it is likely we will merge and simplify this overlap of administrative functions in the security space. Administrators will associate policy with people, and that policy will apply to a variety of target objects. Whether these targets were classically "servers" or "network devices" will be irrelevant. Users and administrators will both have abstracted views of the physical computing environments, defined by business intent and not by topology or by a quirk of a vendor's security implementation.

Directory Services

Historically, directory services have been designed with specific requirements in mind. The relatively static

nature of the directory information update model was accepted, and instead the design optimized the high-speed access and replication characteristics. However, the advent of directory-enabled networking, Dynamic Host Control Protocol (DHCP) leases, data flow-rate, network state, and routing statistics have made it necessary to store and distribute low-latency, transient, and dynamic data. When a directory service is capable of supporting 'dynamic' data types, it becomes useful to policy-based network operations, supporting service provisioning or delivery of network state information to applications. An active association between a user or application context stored in the directory and the network can be discovered and used. New possibilities enabled by this technology follow:

- provide secure management of information from a variety of sources, including applications and network devices
- define, register, and provide publish/subscribe features for network events
- process network events for applications, devices, and users
- expose APIs to applications to take advantage of directory-based services
- maintain state information for devices, users, and applications to support policy-based management

The core directory service acts as the single point of administration for all resources, including users, files, peripheral devices, databases, Web access, and other objects. Extended functionality such as that provided by CNS/AD [16] provides access to vendor-specific network elements and services, and securely and efficiently propagated dynamic data. High-speed replication services securely propagate cached data among all directory servers, enabling them to manage dynamic network information such as IP lease or user password. Core directory functionality relates relatively static information about users and applications to dynamic information describing a given service request and the context in which the request has been issued. In summary, directory services will play a significant part in the establishment and the long-term success of policy-based management.

TECHNOLOGY EXPERIENCE AND USAGE

Within Intel IT, our current agenda is to qualify the suitability of policy-based management for the

corporate enterprise environment as well as to define the transitional models to support e-Business. We believe this approach matches current industry and product roadmaps. Our initial technology evaluation objectives will primarily focus on QoS and bandwidth management to enable service management and resource management, although we hope to gain insight into other usage areas including security and network management in a wider context. To facilitate these objectives, as illustrated in Figure 3, we have developed a Quality of Service Network for Emerging Technologies ("QoSNET"). QoSNET is a production-level and policy-managed network environment to support proof-of-concept of QoS and policy-based management technologies. QoSNET will bring together emerging technologies including multimedia collaboration technologies and other next-generation applications (e.g., VoIP) to investigate intelligent bandwidth management capabilities to support scalable and manageable deployment of these emerging capabilities. Our current activities are focused on technology evaluation to support the integration of policy and network QoS features, while recording the operational procedures necessary to support policy-based management. One of our goals is to validate the technology capability as it merges abstract policy rules with device QoS capabilities, namely 802.1p and IP precedence/TOS control mechanisms within layer-2 (switches) and layer-3 (routers) enabled devices. Using the most recent developments in network equipment, we are also investigating key QoS capabilities including rate control and application-based recognition managed through policy invocations. Secondly, through hands-on experience using policy-based management tools, we will assess potential process or operational management improvements in bandwidth management and network and service management. Interoperability between alternative policy-based management tools and multi-vendor devices is critical to our success. We are also investigating opportunities to policy manage or control standard QoS-oriented technologies being defined through the IETF to support QoS/CoS, specifically, RSVP, Differentiated Services, and MPLS. Other planned areas for technology evaluation will include policy-based routing and evaluation of directory technologies to support the integration of system or application-based policy information. This latter item will include looking at scalability issues on policy within a large enterprise network (i.e., Intel's corporate network) or across administrative domains to support Internet-based service provisioning and policy management.

QoSNET Trials Network

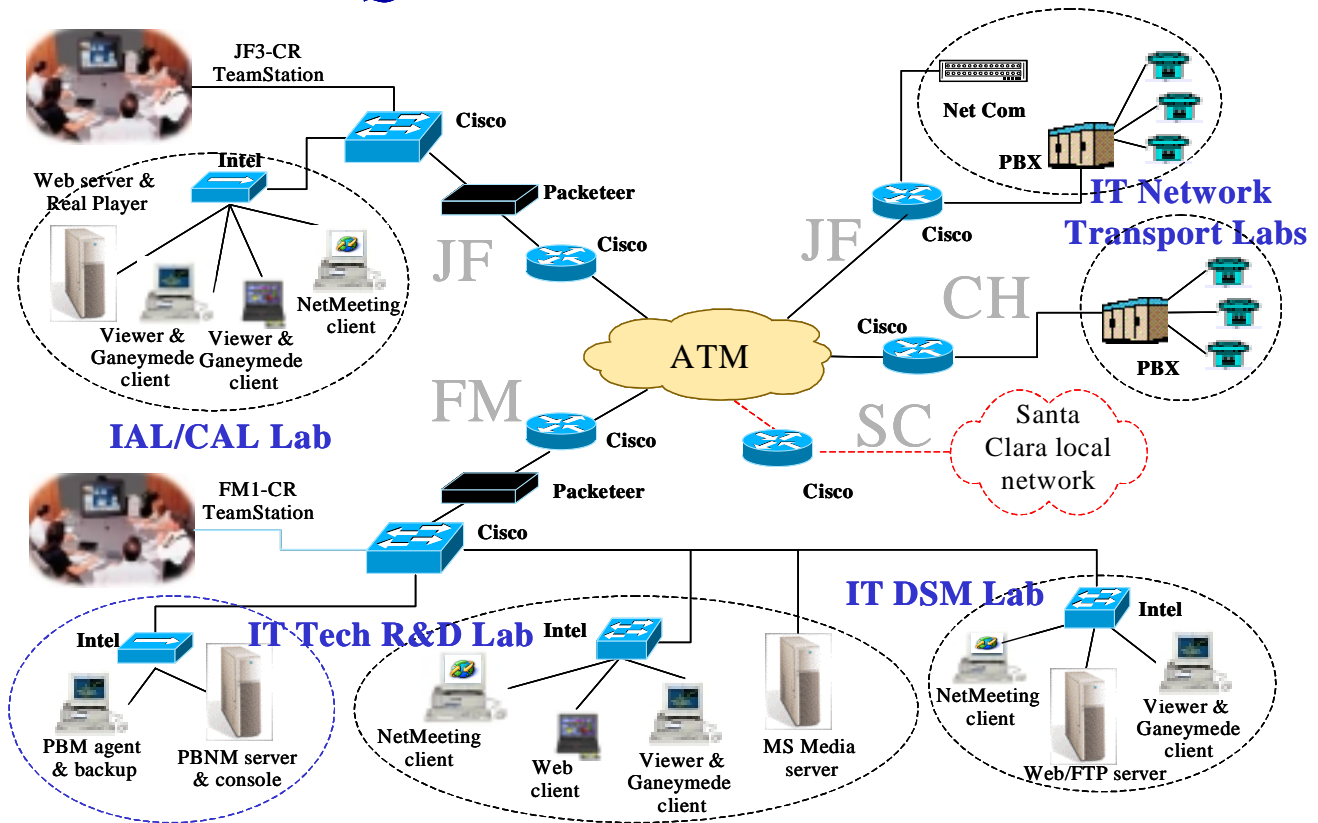


Figure 3: Information Technology QoSNET Trials Network

Usage Scenarios: e-Business

With the introduction of e-Commerce and the move by corporations towards Internet-based businesses, the information technology model will need to support a wide range of applications from a diverse IT customer base including existing internal employees, new employees from acquisitions, corporate suppliers, or individual consumers. Further, the user communication model may source from the corporate private network, the telecommuter accessing corporate resources via a remote access connection (e.g., DSL or satellite), or perhaps a corporate partner operating over an extranet via an ISP virtual private network. Information Technology must meet the diverse connectivity requirements as well as specific user requirements for productivity, network flexibility, service differentiation, isolation, privacy, and manageability. The complexity of the network along with a diverse user base is greatly increased under this new computing model. Furthermore, there will be a shift to delivery and management of services. Under this new paradigm,

motivated by e-Business computing, we identify the following applied areas for policy-based management.

Service Management

As QoS technology and policy become more commonplace, the Internet will support increased business-to-business communications and inter-domain negotiations to ensure resource preservation and SLA's. This will require SLA specifications to be translated into traffic management policies. Moreover, by way of static or dynamic provisioning the SLAs would be enforced through traffic conditioning and device control mechanisms. The introduction of policy management can facilitate these changes by providing the means to translate high-level business policies into device-specific mechanisms to support SLA management. Such a model is proposed [17], including *customer policy*, *service policy*, and *flow policy*. This model closely follows the policy and provisioning administration models.

Dynamic Provisioning

There is a general industry trend towards more cost-effective models for bandwidth provisioning (e.g., VPNs). Static models, based on leased lines through network carriers, are substituted for dynamic virtual pipes or remote access connections available through either service providers or network carriers. The expectation of support for both short-lived and long-lived networks based on changing capacity and geographical domains, supporting either unscheduled or negotiated schedules, is a reasonable one for ISP customers or IT organizations. In addition to traditional office and business applications, provisioning models will require support for alternative QoS traffic types enabling emerging computing applications: VoIP, video-based technologies supporting real-time streaming applications for online training, or real-time applications to support video conferencing or desktop collaboration. The provisioning or signaling support for intra-domain and inter-domain SLAs, in addition to resource management features for load balancing or rate control, will require tools that can facilitate network control and administration automation over complex networks and alternative user models. Policy-based management tools will provide obvious value here.

Internet Pricing and Billing

Models for Internet-based e-Business will involve some form of accounting for subscribed services and use of shared resources. Although several proposals [18, 19] have been published in the area of Internet pricing, it is possible that the right model is somewhere between the shared flat rate model of Internet ISP's and the network carrier telephony model of usage-based pricing. Moreover, with alternative levels of service (i.e., QoS) becoming available to Internet users or organizational subscribers, the pricing model may be extended to support subscribed service levels, in addition to resource use. Policy-based data stores will support the availability of such information to support billing, SLA auditing, and validation. For example, a billing model based on the following policies is proposed [17]:

- Billing policy: customer type and credit associated with a request for a given service
- Charging policy: charging tables describing service, resources, time, and cost for a domain
- Accounting policy: accounting information about resources used and the cost for each customer

Security

The requirements for policy-based security focus on the integration of policies across administratively separate

or heterogeneous domain boundaries. Security has to allow individual consumers, corporate suppliers, and acquisitions/mergers as well as corporate business partners to operate under a shared, and what is perceived to be, border-less infrastructure. End-end security will require a tighter integration of policy across the separate security realms, which are traditionally disjointed across networks, applications, and servers. Finally, dynamic policies will be a requirement for the administrator, providing him/her the means to perform on-the-fly control or automation of short-lived policies to secure content and communications (e.g., inter-company video conference or online supplier training sessions). Enabling such a paradigm will require a higher level of abstraction, automation, and integration across infrastructure elements. We feel that policy-based management solutions would work here by aiding in the definition of allowable security associations, integrating policy abstractions across administrative or heterogeneous security boundaries, facilitating encryption parameterization, and by rapid and dynamic configuration of boundary devices (e.g., VPN, firewall, and proxy services).

SUMMARY AND FUTURE WORK

In this paper, we introduce the key drivers for current IT infrastructure evolution: mission criticality, network architectural agility, service flexibility, and efficiency. To enable these drivers, we presented the background on policy-based management beginning with the necessary set of requirements to realize the technology. We then presented a framework for policy-based management and discussed implementation alternatives to support the proposed framework. We believe that IT business and operational models have to be augmented to transition to a policy-driven management approach; yet, we argue that the changes should help IT organizations align business objectives with more effective use of infrastructure resources. We proposed a simple methodology, which could be implemented by enterprise managers and we suggested that change management and service-level management evolve to align with policy-driven business and operational processes. The ROI decision criteria for policy-based management and technology challenges that IT organizations must address will include the selection and integration of QoS technology, vendor technology interoperability, enabling policy-based security, network management, and directory service integration.

Finally, there are still pending issues that will not be resolved until policy-based management matures industry-wide. Industry standards (primarily IETF and

DMTF) in the areas of policy-based directory schemas, QoS technologies (e.g., DiffServ, RSVP, MPLS), and policy and directory communication protocols (e.g., COPS, LDAP) are still under development and may delay full vendor adoption. Intel is very active in driving technology in the direction of these standards. Additionally, policy scalability, QoS and security conflict resolution, and interoperability will further influence IT strategies and the adoption of PBM technologies.

Over the next year, Intel IT will investigate these issues while gaining experience with policy and QoS technologies. We anticipate continued convergence in the directory arena, as this technology should serve as the foundation for the success of PBM. A widely deployed solution will depend on the eventual integration of alternative technology. The move to e-Business and Internet-based computing will force organizations as well as ISP's to focus on and speed the delivery of a policy-driven approach to managing Internet-based IT infrastructure and enhanced network services.

ACKNOWLEDGMENTS

The authors thank David Mills for his insights and his contributions to the security sections presented in this paper, as well as folks from the Intel Architecture Labs and various groups from within Information Technology including Finance, Strategy & Technology, and Infrastructure Products Engineering for their contributions.

We also thank NSD for their technology support and our external associates, including Hewlett Packard, Cisco Systems, Packeteer, and Net Com Systems for their technology support and collaboration.

REFERENCES

- [1] J. Strassner, E. Ellesson, "Terminology for Describing Network Policy and Services," *Internet Draft draft-strassner-policy-terms-01.txt*, Feb. 1999.
- [2] R. Rajan, D. Verma, S. Kamat, E. Felstaine, S. Herzog, "A Policy Framework for Integrated and Differentiated Services in the Internet," *IEEE Network Magazine*, Sept./Oct. 1999.
- [3] J. Boyle, R. Cohen, D. Durham, S. Herzog, R. Rajan, A. Sastry, "The COPS (Common Open Policy Service) Protocol," *Internet Draft draft-ietf-rap-cops-07.txt*. Aug. 1999.
- [4] W. Yeong, T. Howes, S. Kille, "Lightweight Directory Access Protocol," RFC 1777, March 1995.
- [5] B. Moore, E. Ellesson, J. Strassner, "Policy Framework Core Information Model," *Internet Draft draft-ietf-policy-core-infomodel-00.txt*. June 1999.
- [6] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource ReSerVation Protocol (RSVP)," Version 1 Functional Specification, *RFC 2205*, Sept. 1997.
- [7] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture: an Overview," *RFC 1633*, June 1994.
- [8] Y. Bernet, S. Blake, J. Binder, M. Carlson, S. Keshav, E. Davies, B. Ohlman, D. Verma, Z. Wang, W. Weiss, "A Framework for Differentiated Services," *Internet-Draft draft-ietf-diffserv-framework-01.txt*, work in progress, Feb. 1999.
- [9] R. Callon, P. Doolan, N. Feldman, A. Fredette, G. Swallow, A. Viswanathan, "A Framework for Multiprotocol Label Switching," *Internet Draft, draft-ietf-mpls-framework-01.txt*, May 1998.
- [10] M. Seaman, A. Smith, E. Crawley, J. Wroclawski, "Integrated Service Mappings on IEEE 802 Networks," *Internet Draft, draft-ietf-issll-is802-svc-mapping-03.txt*, Nov. 1998.
- [11] K. Nichols, V. Jacobson, and L. Zhang, "A Two-bit Differentiated Services Architecture for the Internet," *Internet Draft, draft-nichols-diff-svc-arch-00.txt*. Nov. 1997.
- [12] Z. Xiao, L.M. Ni, "Internet QoS: Big Picture," *Dept of CS, Michigan State University*.
- [13] I. Andrikopoulos, G. Pavlou, "Supporting Differentiated Services in MPLS Networks," *Proc. 7th International Workshop on Quality of Service (IWQOS'99)*, London, May 1999.
- [14] R. Yavatkar et al., "A Framework for use of RSVP with Diff-serv Networks," *Internet Draft draft-ietf-diffserv-rsvp-00.txt*, June 1998.
- [15] B. Aiken, J. Strassner, B. Carpenter, I. Foster, C. Lynch, J. Mambrette, R. Moore, B. Teitelbaum, "Terminology for Describing Middleware for Network Policy and Services," *Internet Draft, draft-aiken-middleware-reqndef-01.txt*, May 1999.
- [16] Cisco White Paper: *Cisco Networking Services for Active Directory*, available at

http://www.cisco.com/warp/public/cc/cisco/mkt/serv/prod/sms/common/tech/cnsad_wp.htm

- [17] T. Ebata, M. Takihiro, S. Miyake, M. Koizumi, F. Hartanto, G. Carle, "Interdomain QoS Provisioning and Accounting," *Internet Draft, draft-ebata-interdomain-qos-acct-00.txt.*, Oct. 1999.
- [18] N. Semret, R. F. Liao, A. Campbell, A. Lazar, "Market Pricing of Differentiated Services," *Proc. 7th International Workshop on Quality of Service (IWQOS'99)*, London, May 1999.
- [19] S. Shenker, D. Clark, D. Estrin, S. Herzog, "Pricing in Computer Networks: Reshaping the Research Agenda," *ACM Computer Communications Review*, pp. 19-43, 1996.

AUTHORS' BIOGRAPHIES

John Vicente is a member of Intel's IT organization where he is involved with strategy and technology in the areas of Internet-QoS, policy-based networking, multimedia, and programmable networks. John is also working with Dr. Andrew Campbell as a Ph.D. candidate at the Center for Telecommunications Research at Columbia University, New York. He received his M.S. in Electrical Engineering from the University of Southern California, Los Angeles, CA in 1991 and his B.S. in Computer Engineering from Northeastern University, Boston, MA in 1986. His e-mail is john.vicente@intel.com.

Harold Cartmill, since joining Intel in 1990, has applied his extensive LAN, WAN, server, mainframe, and system management expertise to solving administration and support problems in Intel's mainframe and Windows NT server environments. Harry has served as Technical Lead for Intel's Remote LAN, cc:Mail, Global Remote Server, and Event Management projects. His most recent accomplishments have been in the design, implementation, and support of systems management methodologies. Concurrently, Harry is actively pursuing his B.S. degree in electrical engineering at California State University, Sacramento, CA. His estimated completed date is the Spring of 2001. His email is harold.l.cartmill@intel.com.

Glen Maxson worked for the Boeing Company in Seattle, WA for 17 years before coming to Intel in 1994. While at Boeing, Glen spent the majority of his time solving the 'Enterprise Directory' problem, a quest that continues to challenge him as Intel's Directory Service architect. He completed his undergraduate studies at Pennsylvania State University in 1977. Glen holds a

Certificate in Data Resource Management (1991). His e-mail is glen.maxson@intel.com.

Shelby Siegel is a network architect in Intel's IT Strategy and Technology organization where he does strategic planning for Intel's worldwide network. He received his B.S. degree in mathematical sciences and his M.S. degree in computer science: computer engineering from Stanford University in 1975. His e-mail is shelby.siegel@intel.com.

Russ Fenger is a senior architect and engineering manager in the Intel Architecture Labs where he leads the research and development of Policy Based Network Management architectures. Russ's other research interests are quality of service in the Internet and in network security. Russ has been with Intel since 1983 after graduating from Iowa State University. His e-mail is russell.j.fenger@intel.com.

Copyright © Intel Corporation 2000. Legal notices at <http://www.intel.com/tradmarx.htm>.