Operating Systems

As a program that controls and mediates the access of other programs to hardware resources, an operating system (OS) faces the challenges of reconciling ever-changing program requirements and rapidly advancing hardware technologies. At IBM Research, these challenges are being addressed with a variety of approaches.

K42

The K42 group is developing a new operating system for 64-bit cache-coherent multiprocessors, ranging from small-scale to very large-scale multiprocessors. Each virtual resource (for example, file, process or memory region) is implemented by independent objects to achieve a high degree of locality and to avoid global locks and data structures. Applications with special needs, like databases and web servers, can pick the building blocks (objects) that implement the resources they use to match their expected demands. K42 currently runs on 64-bit PowerPC and MIPS® systems.

SawMill

The SawMill project addresses the complexity of developing and maintaining a variety of custom operating systems. By first decomposing an existing operating system into flexible, reusable components, and then defining a suitable architecture, it is possible to compose efficient and robust operating systems. This approach is being applied to Linux® to create SawMill Linux, which consists of Linux-based components to provide both typical system services and general components, such as memory, task, device, and access control managers.

Leeds/Linux

The focus of our Leeds/Linux effort is to explore operating system support for secure embedded devices. Our specific target is the IBM 4758 secure coprocessor, a FIPS-140 level 4 device with hardware cryptographic support and physical tamper protection. One goal is to make Linux suitable as a stand-alone operating system for small-embedded devices. We are also investigating the security aspects of devices, including the handling of tamper-responsive features, encryption of file contents, and handling of trust.

Adaptive Fast Path Architecture (AFPA)

The AFPA architecture complements high-performance network servers. It includes a RAM-based cache for serving frequently accessed static content and a reverse proxy that distributes requests for dynamic content to a set of back-end servers. AFPA has been implemented on several server platforms, including Windows™ 2000, OS/390®, AIX®, and Linux. By conservative estimates, AFPA more than doubles capacity for serving static content compared to the best conventional server architectures.

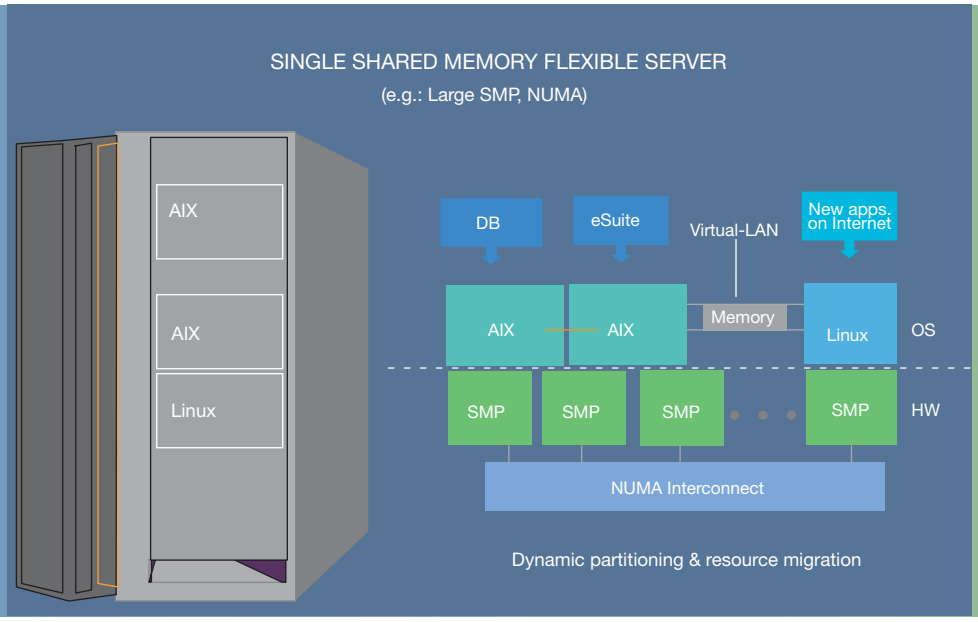

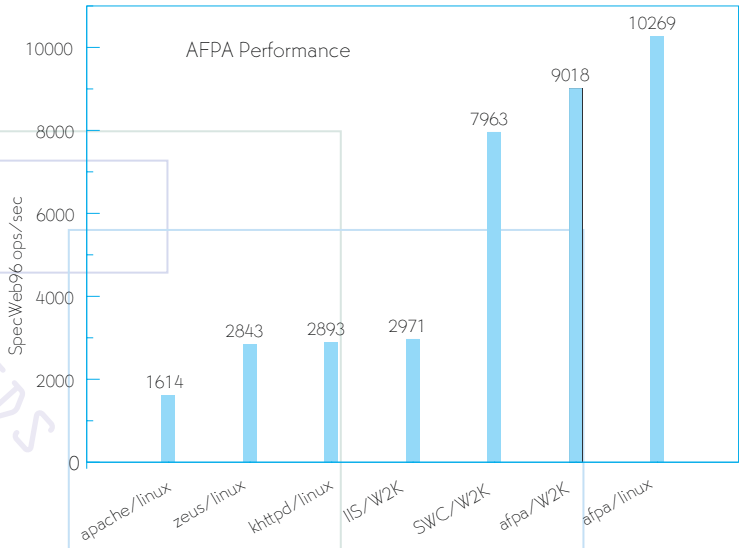

Flexible UNIX Servers

Our flexible UNIX® server group addresses high-end enterprise operating systems with focus on AIX and Linux. We are developing a dynamically partitionable, large-scale shared memory system for SMPs or NUMA systems. Each partition in such a system runs a fault-contained instance of a potentially different operating system. A partitionable system can address scalability, workload isolation, and capacity-on-demand requirements. We are adapting AIX and Linux on PowerPC and Intel® platforms to enable dynamic partitioning through online addition and removal of system resources, such as CPUs, memory , and I/O, without the need to reboot.

Internet QoS

The Web and the Internet together constitute a critical information, entertainment, and commerce infrastructure that is rapidly evolving from a best-effort service model to one in which service differentiation can be provided for users, services, and applications. We are pursuing a number of research directions investigating support for service differentiation on Internet servers and proxies, including flexible and efficient network bandwidth management, dynamic inbound connection rate control for server overload protection, and resource management (CPU, network bandwidth) for web sites and Internet applications colocated on the same server platform, as well as policy-enabled management for server farms.

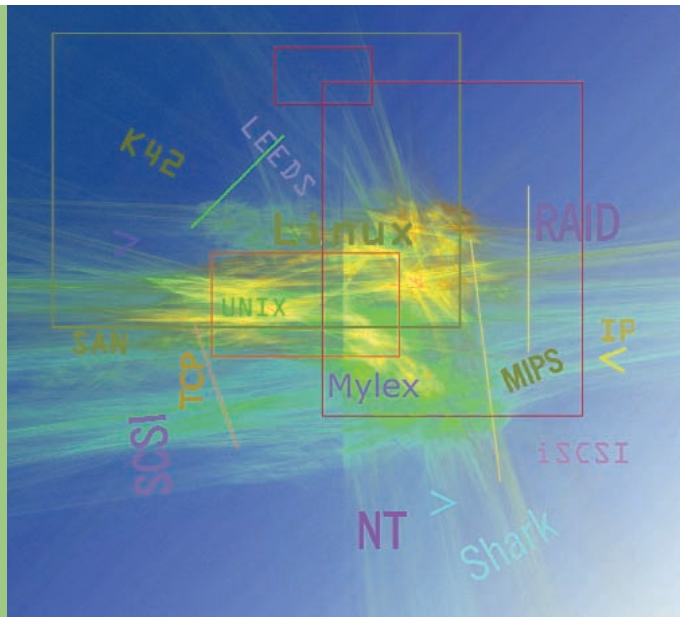

Enterprise Linux

This project focuses on enabling Linux to efficiently support enterprise-level applications. These include business-centric applications in which application availability, vertical and complete solutions, and customer service are important. The work involves both the traditional applications, such as web, file, and print servers, as well as database, e-commerce (WebSphere®), directory services, and Linux kernel scalability. We have implemented different schedulers that increase scalability under heavy load, increased system activity, and increased CPU numbers by reducing lock contention within the kernel.