UNIX® Future Directions White Paper

John Sontag, Chief Architect
Open Systems Software Division
Hewlett Packard Company
19447 Pruneridge Ave. MS 47LA3
Cupertino, CA 95014

Situation: Growing Demands on Business Critical Computing

Information technology (IT) has become a competitive weapon for companies today. The realization that information is the key to enabling an enterprise to maximize its results is changing the way businesses view their computing resources. Integrating data into information means pulling data from various sources in distributed locations: Not just one format of data but collections of "rich" data, with embedded links, audio and video, and smart applications that have distributed objects.

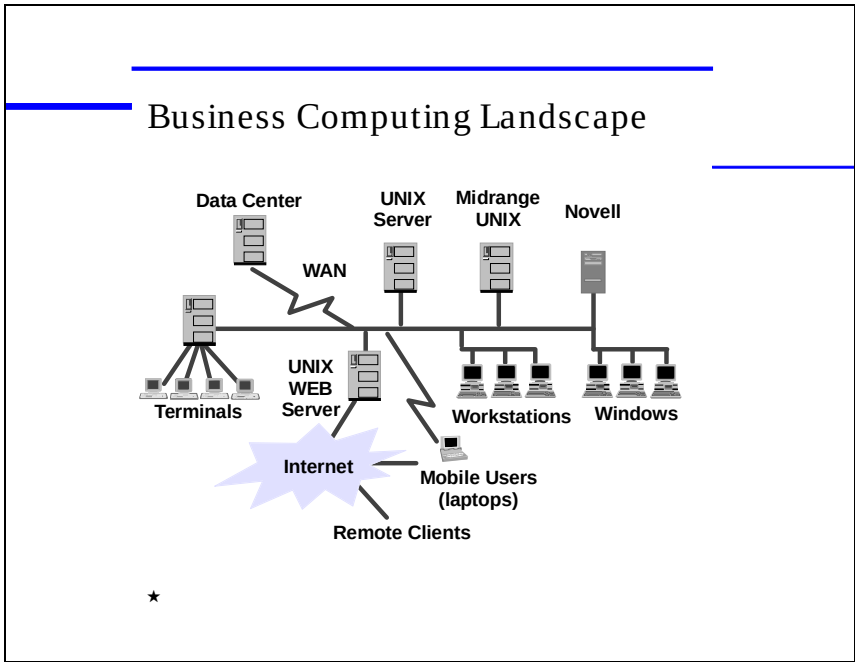

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

Gathering and processing all this data into useful information is what business critical computing is all about. Over the past decade, the computing landscape has gradually changed from centralized processing of everything, to islands of PCs and workstation LANs, and now tiers of computing specializing in certain tasks. Bringing this power to the user is good for productivity and use of resources, but complicated by interoperability issues, compatibility issues, and management issues.

The amount of information, the fact that it's distributed, and the hybrid computing environment makes deployment, administration, and control of this infrastructure a difficult task. The experience level of administrators necessary to manage this complex environment is not available at most user sites...and it shouldn't be.

Add to all of this the rapid introduction of new capabilities and services to the environment and even the strong of heart begin to wonder if it is possible to create an IT infrastructure that will endure and enable their business to compete in a global marketplace. It’s not easy, but there is hope and UNIX plays a key part in the solution.

Understanding this landscape, the needs of business critical computing, and with an eye on the kinds of demands that businesses will make on their IT infrastructure in the future, HP believes that Next-Generation UNIX system technology should have certain inherent capabilities. There should be no question about its robustness and ability to recover from failures. It should enable businesses to take advantage of available technologies today, yet be designed to easily incorporate new breakthroughs in hardware, systems design or software without having to start over. And it should enable the integration and management of a heterogeneous computing environment.

Solution: Enrich the Unix Infrastructure

To meet current and future business critical computing needs, the next-generation UNIX system needs to be more than what is traditionally considered an operating system (OS). It must provide customers with a way to manage and overcome the challenges described earlier:

- Complex data gathering
- Hybrid, diverse environments
- Deploying and managing solutions

The system must provide a basic infrastructure for common frameworks for software development, distribution, and systems management. It must be scaleable to fit needs, from small businesses to heavy duty needs of enterprise computing. It must be easy to incorporate emerging technologies—either by replacing existing functionality or by adding new features.

This set of system technologies must not only meet the requirements of today's computing needs but also enable new opportunities for emerging markets and solutions, including leading network integration for server-based applications.

Next-Generation UNIX Technology

Together, HP and SCO, the leading revenue-generating and volume UNIX solution providers along with our industry partners, will produce the best systems software for meeting business critical networking and computing functions.

This systems software will be deployed on industry leading hardware solutions, both 32- and 64-bit systems, including Intel's iAX86, PA-RISC, the HP/Intel joint architecture, and others.

Key Design Principles

The technology will be made of modular components that fit together, not a monolithic structure. This will enable it to be flexible for special purpose applications and extensible to add new functionality as emerging technologies became commercially available. The components will be extensible and replaceable, allow value-added enhancements for market needs, and will not affect the quality or consistency of an application's platform. This leadership in design will allow market adaptation and added value while proving easy to use and manage.

To help independent software vendors (ISVs), systems builders and customers, the basic framework such as utilities, the APIs, the development tools and environment, software installation licensing and management will all be consistent.

In keeping with the vision of a robust, business critical platform the OS will have Reliability, Availability, Supportability (RAS) features. Performance will scale from small servers to large parallel clusters of systems or mainframe class through massively parallel processors (MPP).

One of the key requirements of our customers is to leverage their existing investments and productivity tools. For example, the new system design allows easy and transparent interoperability with Microsoft Windows' desktops or Novell NetWare LANs.

UNIX has traditionally had the reputation of being hard to use. Designed into this system will be usability and functionality that will provide consistent ways of managing both locally and remotely.

Finally, in keeping with our customer-centered perspective, the system will be able to do all this without forcing users to revamp their existing infrastructure.

Three Dimensional Architecture (3DA)

Apart from making this system reliable, easy to use, and easy to manage, the key is to make it a system that encourages innovation. Today, modifying the OS to adapt it to a new type of hardware, or stripping it down for a specialized application, means a significant amount of work and redesign. In addition to the effort, the eventual result usually "breaks" application compatibility. To balance the need for compatibility with adaptation and innovation, the next-generation UNIX technology will partition the system into replaceable and configurable modules. This structure enables portability, functionality and extensibility with an upgrade path. Each module has a defined interface which must be maintained for compatibility, and any module can be replaced to provide customization for a given processor platform or to enable a broader range of solutions. For instance, application compatibility from a small workstation up to an MPP system can be achieved by changing only the modules necessary to achieve good system scalability.

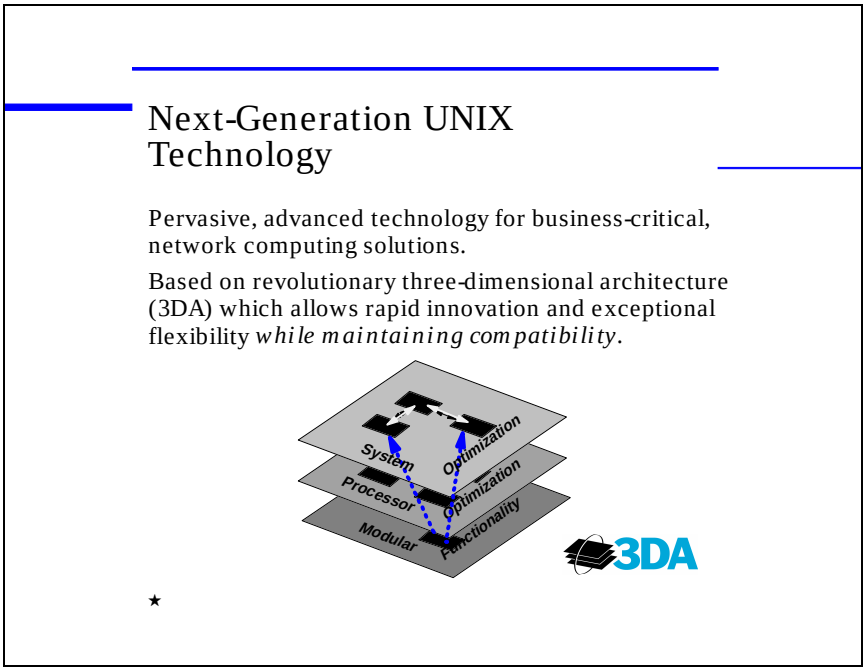

First Dimension—Modular Functionality: Extensible, plug-in functionality for extended markets.

The modularity of the architecture allows for adaptation of one or many components to meet the needs of a particular solution. To maintain compatible solutions and enable convergence between operating system vendors the intention is to refine module boundaries over time. Many of these boundaries are already clear since Unix95 specifies a large part of the API space for applications. File systems and IO subsystems

have been moving to more formal interfaces, as have the networking protocol stacks. This is not magic or wishful thinking, but just good engineering practice.

Second Dimension—Processor Optimization: Efficient processor specific optimization.

The base module layout of the system has been defined. Next, the IT needs to be specialized for different microprocessor architectures to maximize application performance. The intention is to do this without impacting the API that application writers see by replacing a small number of modules in the system with ones that adapt to the hardware or provide some hand-coded optimizations. The difference between this architecture and a microkernel or layered system is that module replacement can happen anywhere in the system necessary to achieve system goals. This is key to realizing the advances in microprocessor architectures, which lend themselves to different algorithmic optimizations.

Third Dimension—System Optimization: Complete system extensibility.

After the system has been adapted to a microprocessor, a broad range of solutions must still be enabled. Once again, compatibility for applications is paramount, and the need to adapt internal algorithms and policies remains critical to delivering value to the customer. At this level, another small number of modules may be replaced to adapt to specific system constructs, such as specialized IO adaptors, clustered architectures, or fault tolerant platforms. This architecture allows for changes which are necessary for proper operation without impacting application compatibility.

Infrastructure & Capabilities

The goal of the system is to use the framework of 3DA to enable delivery of both the core technology and the basic infrastructure for business computing needs for today and in the future. From our customers' perspective, there are two essential areas where an operating system must excel: Business critical support—defined as continuously available systems, and network ready—taking advantage of the Internet and allowing shared information access with other servers, clients and LANs. These must be accomplished with a guarantee that the systems and data are secure from unauthorized access. Finally, these systems must be highly reliable, enabling businesses to access or share information across networks while protecting confidentiality, and easy to install, configure, tune and administer. The management scheme must be flexible enough to allow distributed management or centralized control. The key phrase is EASE-of-USE.

Business Critical Support—Reliability and High Availability

Protecting data and ensuring that it is available when needed is a critical element of an OS for business critical solutions. Journaling and intent logging will be provided to protect data loss or to recreate data from transaction logs. Software RAID and Hardware RAID for higher levels of data integrity will also be supported. Software RAID 0, 1, 5 and combinations of these levels enable reliable solutions using commodity disks. The OS will provide the ability to reconfigure on-line and to be resilient to hardware failures by enabling hot swap, system failovers or automatically shutting off a piece of hardware that has failed rather than panicking the system.

Where possible on supporting hardware architectures, further protection on system failures such as improper memory management can be isolated by hardware protection. These failing services can be brought off-line as needed to prevent complete system failure.

Fast file system recovery and system-error correction will be designed into the system in case of disaster to allow a downed system to come back on line quickly.

Business Critical Support—Performance and Scalability

The 64-bit system technology will be designed to take full advantage of advanced processor architectures, such as the 64-bit Intel Merced processor, enabling efficient use of the hardware from single CPU systems through SMP and clustering solutions. 32-way SMP and support for clustering will enable this operating platform to scale its performance to meet the most demanding requirements of businesses.

Support for 64-bit processors and the use of clustering with high-speed interconnects will allow much larger memory models (>> 4GB) than are possible today. This system will take advantage of systems with huge amounts of RAM to maximize application performance. 64-bit technology will also allow systems to support much larger disks, file systems and files. Support for large, distributed storage units or disk farms will be available. The system will run effectively on medium-sized servers and also on systems with terabytes of disk farms and gigabytes of main memory. Optimization for some very high-end solutions such as MPP systems will be supported by extension through 3DA.

Performance and scalability are not limited to the operating system level. In addition to this, in the case of SMP environments, an application will be able to take advantage of the full processing power of multiple CPUs via kernel and user-level threads. The OS will offer optimized database support via features such as high performance kernel, async I/O, semaphores and timers.

Applications may also be designed to make use of the network for reliability/availability. This will enable a new class of highly available distributed applications.

As businesses move increasingly to client/server environments, the network interface and transport plays a big role in performance. This system will provide support for high speed networks such as 100-Base, and FDDI, and for, fault resilient file, print and directory services.

File and print services will be designed to scale to meet the growing needs of businesses to share resources. One of the strengths of NDS is its ability to scale to accommodate the traffic in large enterprises.

Network Ready—Multi-protocol OS

From the operating system that invented and developed networking services, further integration of PC networking and inter- and intra-enterprise networking environments will be supported. The incredible importance of networking does not need to be explained with today's explosion of the Internet. Next-Generation UNIX technology will include a complete set of Internet-enabled software to provide LAN and WAN connections for doing business on the Net. This includes the ability to do searches for information as well as conducting business on the Internet or on private networks. This is all supported by delivering:

- Leading routing protocols such as OSPF
- POP, POP3 on the client side
- Technology for running Web-enabled applications
- IPversion6 support to deal with the IP address space problem
- Widely used LAN protocols such as IPX, SPX, SPXII, NetBIOS and NetBEUI

Users will be able to connect to the Web, using either packet or switched networks. WAN connectivity means support for technologies such as ISDN, ATM, Frame Relay and protocols such as PPP. It will also provide support for remote client access.

Network Ready—Seamless connectivity

Easy integration of Windows desktops and other PC clients is a key design principle. Several options for integrating clients will be available. These options will cover both resource sharing and management. NetWare file and print services on a UNIX system will allow, for example, a Windows NetWare client to share files and printers off the UNIX system as if it were a NetWare server. DCE services—Executive, Security, RPC—will enable distributed client/server computing. In order to make access to server resources transparent, Novell Network Directory Services (NDS) will be

incorporated to provide global directory services and transparent access to server resources. Users currently using NFS for resource sharing and NIS for managing the name space, can continue to do so.

For those environments that currently use DCE X.500 directory services, the system will provide integration of NDS and DCE directory and security services. From a user's perspective, login validation will occur only on one server. Once a user has successfully logged in, access to allowed resources on other servers will be automatic.

Network Ready—New class of application software

Next-Generation UNIX technology will integrate and deploy a key set of software development and infrastructure tools to make network computing a reality and with this will come the infrastructure for building new classes of networked applications.

The distributed architecture will be CORBA-based and will be able to integrate with Microsoft networked OLE when it ships. Of course, this will be done in a secure fashion. By tying together networking security, naming, authentication and identification, administrators will have an easier time creating and managing their networks and users will have an easier time using distributed systems.

Essential Security—Secure networking

System software to support mission critical security with ease of administration in a networked or standalone environment is critical to business success. Businesses today require security as a core competency for company firewalls and for inter-enterprise business. It must allow Internet commerce and exchange. The networking services and the core OS will be secure but easy to administer with:

- NDS—RSA public key—for authentication and identification of users and resources on the network
- Bundled Internet firewall to protect the system from attacks through the Net
- Private key security—using Kerberos authentication

Essential Security—Secure commercial access

The core system will ensure the data and services provided by the system will not be compromised. The security mechanisms will allow ease of use and administration of the system and ensure access to critical data is guarded with:

- Control individual access to individual objects in the system ACLs
- Share access to individuals, groups, and lists of users for differing privileges
- Backup services will handle security attributes of all objects
- C2 level auditing capabilities are easy to configure and use

- Audit both systems and network access
- Role based administration so there is no "Root see/get all" data
- Removal of completely privileged ROOT access to all data and functions
- Management responsibilities can be partitioned without compromising security

Ease of Management—Software management

Software management is one critical area of administration and management. Next-Generation UNIX technology will provide a comprehensive set of software management capabilities. Some of the features to be included are:

- Standard installation, packaging and distribution formats
- Standard licensing format
- Floating licenses, usage monitoring and metering
- Push/pull software distribution mechanisms. Includes distribution to PC clients
- Version management
- Enables "software on demand"
- Extensible software distribution framework to include new functionality and client support

Together, these services can provide timely software delivery, software-on-demand, effective software inventory and comprehensive on-line software management.

Ease of Management—Frameworks

System based on Next-Generation UNIX technology will be the easiest server-based systems to administer in the industry. The cost of administration will be reduced by providing a common set of tools and infrastructure for applications, networked environments and stand-alone systems with:

- Advances in standard hardware support such as “Plug and Play”
- Self-healing components (ie. disconnected directory services that come on line)
- Self-configuration with reasonable defaults

The best type of management is what is no longer required—make it as self managing as feasible, provide proactive management facilities, alert administrator of imminent problems. This management functionality will cover clients as well and include:

- Management framework with well-defined APIs that can be used to plug into enterprise management solutions
- Common framework and APIs for all management tasks (common repository and admin policies, DMI support for asset management)

- Directory services—single point of administration for all network-wide static and dynamic resources
- Comprehensive network system monitoring capabilities
- Secure SNMP support for network management

The OS will plug in to all industry leading network management frameworks including HP OpenView, IBM NetView, and Computer Associates Unicenter.

Ease of Management—Easy to use utilities and services

Management within many systems is an afterthought resulting in extensive investment to efficiently run computing resources. Systems based on Next-Generation UNIX technology will have the right management framework to support the complex needs for management. Built-in manageability and ease of use through:

- GUI and character-based management features that will allow management across the LAN or WAN environments including the capability to take control of the system (remote control)
- One step installation: Easy installation, minimal questions, remote installation/software distribution, pre configured systems
- On line backup, network backup, remote backup of clients
- Dynamic file management—on-line management (shrink, grow, defragment, etc.)
- Performance monitoring and basic capacity planning tools for sophisticated environments
- Network tracing and logging—powerful network enabled debugging facilities

Ease of Deployment

Each of the above requirements must be accomplished within a framework that is easy to use, install, and manage. One of the biggest improvements will be that what you see and get on any system will be compatible.

The system design allows room for the development of value-added features by ISVs, Systems Integrators, OEMs, etc. but a common environment and framework will be supported on all systems. This common set of services will allow application development and deployment for systems in traditional, as well as emerging markets. Software developers can be guaranteed that applications will run everywhere, reducing the development, deployment, and support costs for all environments.

Customers and partners will benefit from the reduction of total cost of system deployment including installation, administration, training and value-added software. Today, 78% of all computing costs are deployment and administration expenses, and

this system aims to reduce that cost significantly though removal and automation of tasks.

Ease of Use

A GUI framework and environment, based on CDE, will provide a common end user experience. This framework will be extensible to add value in dynamic ways. Additionally, common utilities, management frameworks and tools will all present the same experience to the administrator. Together, these will reduce expenses and time for training. All differing versions will have the same common framework and environment.

Ease of Development

Development tools and environment will be the same on all systems based on Next-Generation UNIX technology across platforms and vendors. This includes different chip architectures.

GUI-based tools for complete product (coding, building, debugging) will be available with the ability to add third party tools.

Performance analysis tools, including analyzing multi-threaded and parallelized applications will be available.

Common packaging format and infrastructure is a part of the core package and will enable ISVs to reduce the number of package formats needed to support software installation and allow distribution over a network. Combined with the licensing framework, this environment can enable a new generation of products and distribution, reducing cost for ISVs, integrators and IS environments.

Tools to integrate PCs into UNIX services (network OLE framework for servers) will also be available.

Investment Protection

As part of this evolution, it is clear that customers cannot upgrade their platform and applications at the same time, or even uniformly across an enterprise. The technology will be designed to make this gradual evolution easier by supporting a broad set of inter-operation features and through binary compatibility support, without displacing their current infrastructure. A customer’s investment in today's systems will not require massive system replacement or re-formatting of data and system software.

Impact

Next-Generation UNIX technology will be the platform which enables a graceful evolution of UNIX into the future. As user needs continue to evolve, the platform will move along with the evolution, without sacrificing installed solutions. It will allow rapid innovation and unparalleled flexibility while maintaining compatibility.

Application developers and ISVs get a simple porting choice

The leading volume and revenue-generating UNIX system companies, together with our industry partners, will produce and deliver a single, compatible operating system solution for leading industry hardware. By defining and delivering a common implementation of the operating system, the business community, ISVs, OEMs and integrators can be assured that their systems and solutions will operate on the leading volume UNIX environments. This reduces the cost of development, deployment and training for the entire computing industry. We will do this in a way that is compatible with existing applications and allows easy migration as users upgrade their systems to those that provide:

- High performance platforms
- Optimization possibilities through 3DA
- Scaleable solutions to tier products from small business through networked enterprises
- Availability of unified, powerful development environment and tools
- Availability of key scalability and enterprise software framework
- Clustering, distributed file and print, Network OLE
- Single, consistent interfaces across multiple platforms enables develop once, deploy many
- Single development effort to target wider platform base/customer base
- Single, consistent platform and tools for system and application deployment that is compatible and can be migrated from existing systems

OEMs can focus on value add

OEMs are a key element of the success of UNIX, with their ability to customize solutions which closely match their customer needs. With Next-Generation UNIX technology they can:

- Free up resources to innovate and focus on value-add.
- Take advantage of emerging technologies more quickly, easily, without having to re-engineer
- Maintain the ability to differentiate via add ons or compatible extensions

- Develop technology incorporated into base operation system through extensibility of the architecture
- Distinguish hardware solutions with optimized operating system hardware
- Increased availability and recruitment of applications will be improved due to write once guarantee
- Level playing field ensured
- Faster time to market with focused investments

Business critical computing customers can rest easy

These companies will have an unprecedented number of features in the core system (such as clustering, reliability, networking, etc.) and entire new solutions capabilities from OEMs including:

- High-end server environment solution: RAS
- New power of networking solutions and integration
- Ease of management
- Consistent system image from desktop servers to large scale enterprise systems
- Enhanced usability via single desktop and management interface
- Higher level of functionality with the system
- Investment protection, with legacy applications compatibility ensured
- Reduced cost of deployment, training and support costs
- Wide choice of hardware platforms to deliver same highly available computing environment
- Lower cost and higher value due to economies of scale
- An extensible framework for future product solutions.

Sign Up!
Information and registration is available at
http://www.hp.com/go/3dunix