C.6 Dependability and Manageability

Principal Authors:

Richard Baseil, Barry R. Lewin and M. Satyanarayanan

Additional Contributors:

Ed Balkovich, Tom Brand, Gary Campbell, Timothy Chou, Harold T. Daugherty, Coyne Gibson, Bradford Glade, Roger Haskin, John Healy, John H. Howard, Ming-Yee Lai, Barbara Liskov, John McCormack, J. Eliot B. Moss, Sushil G. Munshi, Brian Noble, Louis Scerbo, Ian Service, Tom Soller, James Spencer, Arshad Syed, David Vereeke, John Wilkes and Bernard Ziegler

---

1. Introduction

While speed, ubiquity and functionality may dominate the headlines, the success of the National Information Infrastructure will rest ultimately on whether it meets users' expectations of dependability. This section exposes the fundamental threats to dependability and manageability (D&M) in the NII, and details research recommendations for coping with these threats in the NII's architecture, initial deployment and ongoing successful operation. The scope of the section is broader than its title might imply. Our intent is to span all facets of providing high grades of service to users.

The word "infrastructure" connotes dependability to many users. Indeed, society becomes alerted to the need to invest in infrastructure when its dependability is compromised (e.g., national highways, water distribution systems). In general, users will expect high levels of dependability for many services and will expect flawless delivery of those services. Each of the sample application areas (health care, manufacturing, education, commerce and government) presupposes an NII whose dependability is beyond question.

Imagine the impact of NII outages on these and other critical application areas. Can a patient undergo remote treatment if communication with the doctor is uncertain or if the doctor is unable to view critical diagnostic data in an emergency? Will any company entrust a key aspect of its manufacturing process to the NII if it is unreliable? What kind of a curriculum can educators hope to establish if the timeliness or quality of delivery of lesson material via the NII is poor? If the NII is not dependable, how can any government agency use it to deliver services? With millions of dollars of electronic commerce at stake each second, can the nation tolerate significant NII downtime?

The NII will consist of products built by many suppliers and managed by many network operators, all

based on each individual's understanding of standards and requirements. This multiparty interaction lends itself to interoperability issues that can affect users' perceptions of dependable service. These issues can range from the inability to plug-and-play because of multiple interpretations of requirements, to coding errors that cause failures to propagate nationally.

Meeting user expectations of service quality will depend on how the NII is architected, constructed, implemented and maintained on an ongoing basis. To achieve high performance and dependability in the NII, it is necessary to consider all facets of D&M from the start and include them in the basic building blocks of the system. The NII will neither be dependable nor manageable if D&M issues are an afterthought.

2. A Plan for Dependability and Manageability

2.1 The Challenge of Growth and Evolution

As use of the NII grows, so will expectations of its dependability. Unfortunately, so will the threats to its stability. Although we have considerable experience in managing telecommunications and wide-area computer networks, our current solutions in their present forms, based on current technology and scale, are not going to work for long in the NII. Each quantum increase in scale and service complexity will stress old solutions and render some of them inadequate.

For the NII to remain viable, continuous and ongoing research is needed in scaling up of current solutions, in major improvements to them and in exploring new solutions. It is important to emphasize that refinements to solutions will be as important as their initial development and incorporation into the NII. A one-time, up-front research investment will not sustain the NII as a dependable and manageable entity forever.

A systemwide perspective is essential when addressing these issues. The NII will be a sophisticated and interdependent combination of hardware, software and storage media spanning many levels of the system. Problems can arise due to individual failures at any of these levels or because of unanticipated interactions across levels.

2.2 Balancing Progress With Stability

Clearly, the NII should be built with maximum flexibility to allow for deployment of as many new services as possible. But it is also important to recognize that precautions need to be taken to assure that these services can successfully coexist. In effect, the strategy for coping with growth and evolution must walk a thin line between two kinds of biases. One extreme is being excessively liberal, imposing no controls at all on increases in scale or introduction of new experimental services. The other extreme is being excessively conservative, opposing changes unless their safety and efficacy has been proven beyond doubt. A liberal bias will encourage innovation, lower entry barriers, and encourage competition

and market forces to play their legitimate role in realizing the full potential of the NII. On the other hand, a conservative bias is more likely to yield an NII that is dependable and well managed.

One approach to coping with this dilemma is to classify services and to offer different levels of confidence in the dependability and manageability of those classes. In its simplest form this would involve two classes: "core" and "peripheral." Core services are those considered essential to the NII and for which centralized resources (such as management attention) may have to be committed for meeting service guarantees. Examples of core services might include those needed for a connection-oriented service, access to and use of selected servers of critical national importance, and services related to the NII's management. Adding a core service would be non-trivial and would require "certification." That certification process will need to be defined, but should include important aspects of interoperability. In contrast, peripheral services are those regarded as valuable but less critical and for which availability is offered on a best-effort basis. Adding a peripheral service involves minimal bureaucratic overhead, and no centrally funded resources are committed to its sustenance. During emergencies, peripheral services may be dropped or degraded in favor of core services.

What is core and what is peripheral will change over time. Since today's luxury often becomes tomorrow's necessity, some peripheral services will be promoted to core. Other core services may be created in response to new perceived needs. There will have to be a process of achieving consensus between users, service providers and NII administrators to determine what services are core and what are peripheral. Because the mechanism for achieving this consensus involves public policy rather than research, we do not address it further in this brief. But it will be an important aspect of the overall design and operation of the NII.

2.3 Architecture and Initial Deployment

D&M must be designed into the NII. The architecture needs to be sufficiently robust to help reduce the effects of failures, and it should act as an aid toward fault recovery, not a hindrance. Today, products and networks often have extensive delays in their recoveries from failures because D&M support is built on top of these systems instead of being built in as integral parts of the systems. The areas that should include D&M considerations from the beginning are:

- *Architecture:* The NII's architecture must be of a robust design to assure that reliability requirements will be met. It must include management systems that receive accurate and timely information to monitor and maintain components of the NII. Management information must not be relegated to such a low priority that it is unavailable at the very times when it is critically needed.

- *Requirements and standards:* Incomplete or ambiguous requirements and standards are major threats to NII's reliability. It is vitally important for requirements and standards to be built with complete and unambiguous standards that stress the importance of reliability.

- *Distributed management:* The NII will clearly be managed and updated by many parties. We must avoid having each party act to optimize their own operation at the possible penalty of others or at the expense of the NII as a whole. Automated and manual operational safeguards need to be constructed to ensure distributed manageability can be accomplished safely and effectively.

- *Service characterization:* The NII's services need to be built with dependability in mind. Parameters for characterizing reliability as a part of the service definition are needed. Not all services may require the same level of dependability, but we must avoid designing less dependable services initially and deluding ourselves into thinking that improved D&M can easily be accomplished later.

- *Postmortem capability:* Because failures will happen, the ability to understand those failures and make corresponding improvements needs to be part of the NII's capabilities. We must recognize the eventuality of NII failure and anticipate it by designing in mechanisms to trap key failure data and establishing processes to recreate NII faults in a controlled setting for further study. The airliner "black-box" is an example of where this concept has been adopted and used effectively in another industry.

- *Measurements:* The manageability of the NII depends heavily on what characteristics of the services (and networks) need to be measured, how they are measured and how those data are used. Too much data is almost as useless as too little data, if the uses for the data are not clearly understood from the outset.

- *Service deployment:* The NII's success will depend on integrating new, reliable services into the existing infrastructure with minimal service disruption. The trade-off between speed and reliability needs to be better understood.

2.4 Perennial Problems

While sound architecture is a good first step for the NII, it will not remain problem-free forever. We have identified a number of problems that we characterize as perennial problems for the NII. Although not exhaustive, they typify the threats to the D&M of the NII. These problems are never going to go away; they will always be in the wings, waiting to strike. Constant expenditure of resources and attention will be needed to hold them at bay. A program of continuous research and development will be needed to address these problems as the NII grows and evolves.

- *Failures of networks, servers or environment:* Every hardware element of the NII is subject to mechanical or electrical failure. Further, there is already considerable evidence that software failures overshadow hardware failures; this trend is likely to continue as software in the NII grows more complex. Other failures such as power, earthquakes, floods, operator errors and fiber breaks due to accidental dig ups will also pose a threat to elements of the NII. Unless masked by sufficient redundancy, these failures will result in unacceptable service outages.

- Failure containment is critical because a failed subsystem that is not rapidly isolated may easily bring down other parts of the network. An especially challenging task is establishing that the recovery mechanisms of the NII are indeed capable of handling anticipated failures. This requires proper simulation of the full range of abnormal operational conditions to ensure that recovery actions are triggered and stressed.

- *Overloading of the NII:* For an application that is time critical, delivering information late may be as bad as not delivering it at all. In electronic commerce, for example, a market opportunity may be lost in milliseconds. For remote medical consultation, on the other hand, images of an exploration need to be delivered jitter-free to the specialist's screen. Overloading of NII components will be a major threat to timely delivery of information. Such overloading can be steady-state or transient, and may arise from extraordinary traffic characteristics, unplanned growth of the system, new applications or unexpected service interactions.

- *Security violations including denial of service:* As the NII assumes an increasingly vital role in society, it will also become an attractive target for criminal activity. Unauthorized disclosure and modification of information are obvious threats, but denial of service caused by unauthorized system load and circumvention of accounting mechanisms will be an equally serious threat. Denial of service is especially difficult to deal with because it is often indistinguishable from accidental overloading of the system.

- *Incompatibility and interoperability:* Even at the initial scale of the NII, it will be virtually impossible to pause the system for upgrades. Increased scale will only worsen the problem. As the NII grows, it is likely that upgrades will occur when the network is under stress, leading to even more reliability problems.

- This implies that all changes in the system will have to be introduced gradually rather than atomically. While difficult enough with routine upgrades, this becomes a particularly challenging problem when the motivation for the upgrade is an emergency fix for security or reliability reasons.

- *Clashes and inconsistencies between global and local policies:* Because the NII will be used by a diversity of individuals and organizations, it will be necessary to decentralize policies as far as possible. For example, one organization may centrally coordinate the attachment of new machines to the NII. Others may take a more laissez faire approach and leave this task to individuals or groups. As another example, one organization may bear the entire cost of traffic originating from that organization. Another organization may require individuals or internal cost centers to bear that cost. Unfortunately, a proliferation of policies renders system management more complex. It is inevitable that at least some systemwide policies be adopted for the NII to be manageable. Balancing this tension between centralization and localization will be a constant challenge as the system evolves.

- *Scale-related increase in management complexity:* Operator error is already a major source of failures in high-availability systems. As the NII increases in scale, it will become increasingly difficult to administer and manage. A key problem will be information overload. It is easy to present NII operators with so much information that they are overwhelmed. But excessive data reduction is also harmful: Problems may be masked and fault isolation will be difficult. Visualizing this huge and complex system in a manner that allows timely and effective decisions to be made by operational staff will be a major challenge.

2.5 Research Goals

The D&M research agenda for the NII should stimulate and nurture any activity that will improve our ability to cope with the perennial problems listed in the previous section. Some of the detailed recommendations in this brief, such as research on replication, caching and load balancing, follow directly from this broad goal. The value of such research activities is already recognized today; the creation of the NII will undoubtedly increase their importance.

But our discussions also identified a number of critical research areas for the NII where there is a dearth of current activity. These areas are best described and understood in terms of the goals they support. We list these goals below:

- Make it easier to characterize dependable infrastructure and services. This includes quantifying currently intangible factors such as quality of service, ease of management and usability.

- Support pre-implementation design, analysis and verification of dependability in hardware and software building blocks of the NII and in compositions of them. Especially important is the ability to model applications that adapt to changing NII conditions.

- Guide the deployment and evolution of the NII by modeling of infrastructure and services, and developing techniques for risk and cost-benefit analyses.

- Substantially automate management of security, resource optimization and configuration control, and better understand the human role in complex NII systems and services. Such automation should include continuous monitoring of system health as well as anticipatory actions to forestall problems.

- Enable validation of metrics, models and architectures through prototype construction.

- Allow non-disruptive introduction and reconfiguration of services and provisioning of service databases.

We wish to emphasize the importance of continuous improvement as well as radical innovation. Specifically, we recommend a balanced portfolio of research activities that 1) scale up and bullet-proof

deployed mechanisms and subsystems, 2) extend and refine existing technologies and 3) develop and validate new enabling technologies. D&M are characteristics that will often require in situ study of implementations as well as of system usage and behavior. Hence, the research plan for the NII should recognize that the traditional distinctions between "research," "development" and "deployment" will be fuzzy in the context of D&M.

3. Research and Development Recommendations

The research required to meet the D&M goals of the NII can be grouped along three distinct dimensions. All three dimensions are important, and research on them will be required throughout the life of the NII.

- Characterization and validation of service quality.
- Continuous system operation.
- Orderly growth and evolution.

The next three sections list specific topics pertinent to each of these three research dimensions. For brevity, we list each topic only once even though it may be relevant to more than one research dimension. These topics are not intended to be exhaustive. Rather, they are meant to be examples of the kind of research that must be done to preserve and enhance the D&M of the NII.

3.1 Characterization and Validation of Service Quality

Unless we can crisply specify and quantify the resource requirements and performance of a service, we will have to rely solely on anecdotal evidence to decide if that service is being delivered satisfactorily. Without such characterization, it will be impossible to assess the impact of a new service on the NII. Developing the specifications is not enough; efficient runtime techniques that can confirm that the specifications are being met must also be developed.

1) Developing and Validating Metrics to Describe Service Quality:

- Performance specification.
- Reliability and availability specification.
- Quantifying manageability.
- Characterization of other service parameters (such as jitters, bandwidth, availability, reliability).
- Evaluating effectiveness/appropriateness of metrics.

2) Measuring Service Quality:

- Development of efficient and accurate measurement techniques for service metrics.
- Design and standardization of benchmark suites for service metrics.
- Development of techniques to efficiently monitor service quality in the running system.

3) Incorporating Service Quality into Interface Specifications:

- Development of specification techniques for service quality.
- Design methodology for incorporating and validating appropriate service quality metrics into interfaces.
- Techniques for empirical substantiation of specifications for individual services.
- Research on balancing transparency with user-awareness in services specifications, especially fault tolerance.
- Development of cost/benefit models for different levels of service quality.
- Interoperability and standardization activities across suppliers and operators.
- Compositional techniques to assess service quality from component qualities.

3.2 Continuous System Operation

Techniques to improve the reliability and availability of hardware and software components of the system are clearly needed. To complement this effort, research is also needed on techniques to offer viable fallback options for services. The Titanic mentality ("It can never happen.") and the mentality that "there is no escape anyway" must be avoided. An overall approach that combines failure avoidance with contingency handling is likely to be more robust. Research on techniques to simplify routine system management as well as to help in troubleshooting and crash recovery are also important.

1) Replication Strategies for Masking Failures:

- Service replication techniques for availability and performance.
- Hardware and software redundancy techniques for environmental failures.

- Transactional techniques.
- Fault-tolerant replication protocols.

2) Fallback Mechanisms and Graceful Degradation:

- Exploitation of caching.
- Adaptive techniques for coping with changing conditions.
- Exploration of trade-offs between effort expended to sustain a given quality of service and cost of resorting to fallbacks.
- Validation techniques to ensure degraded service expectations are indeed being met.
- Techniques for failure containment.

3) Software "Black-Box" Technology:

- Efficient techniques to record detailed event histories in compact form.
- Postmortem analysis techniques to determine causes of failure.
- Effective feedback into design and implementation phases.

4) Configuration Management, Resource Optimization and Security Administration:

- Techniques for non-disruptive service introduction and reconfiguration.
- Reliable, fast and non-disruptive database provisioning techniques.
- Load balancing techniques.

5) Resource Control and Accounting:

- Strategies for efficient billing and quota enforcement.
- Dynamic inquiry and negotiation of service costs by applications.
- Anonymous electronic payment strategies.

- Price-based congestion control strategies.

6) Reduction and Visualization of System Management Data:

- Graphical presentation techniques to avoid operator overload.
- Ability to visualize effects of proposed changes.
- Modeling of traffic patterns to distinguish normal and abnormal situations.

7) Management Tools and Techniques:

- Failure detection and isolation techniques in hardware and software.
- Network and service monitoring tools and systems, including those supporting real-time tracing and diagnosis.
- Better understanding of human role in administering complex NII systems/services
- Self-management techniques to reduce the number of highly trained system administration personnel.
- Early-warning techniques to predict service disruptions.
- Intelligent techniques and expert systems to assist and partly automate management.

3.3 Orderly Growth and Evolution

Avoiding problems before they arise will be an essential component of the NII's overall strategy for dependability and manageability. Toward this end, research in tools and techniques to simplify development and stress testing of robust services will be valuable. Research to develop mechanisms for certifying services will also be important. Empirical research on the NII to identify imminent bottlenecks and predict future traffic patterns will also be required.

1) Design and Development Methodologies:

- Incremental construction techniques to reduce cost and enhance reliability.
- Abstraction techniques to reduce apparent complexity of highly available services.

- Techniques for reducing and surviving Byzantine failures.
- Development methodologies supporting change and extensibility of the NII.
- Simulation and emulation methodologies to understand vulnerability of NII to specific types of failures.
- Validation of metrics, models and architectures through prototypes.
- Research supporting systems reliability and analysis: risk and cost-benefit analysis.

2) Development and Validation Tools for Robust Services:

- Interoperability and regression testing tools.
- Tools and analytical techniques to identify and correct undesirable service interactions.
- Reliability analysis tools.
- Service quality assurance tools.
- Frameworks and tools to simplify future implementations of new and existing protocols.
- In situ and standalone stress testing techniques for hardware and software.

3) Modeling Based on Analytical Techniques or Simulation:

- Techniques to evaluate long-term cost-effectiveness of alternative resource allocation strategies and to compare design choices before investment.
- Infrastructure and service modeling to guide the design, deployment and evolution of the NII.
- Modeling and analysis of adaptive applications.
- Techniques for assessing software reliability.
- Tools and techniques for reliability, performance and quality-of-service modeling.
- Analytical approaches to network and server traffic analysis, demand modeling, and capacity measurement and management.

- Measurement, modeling and prediction of hardware and software failures.

4) Long-Term Empirical Studies:

- Workload and traffic characterization.
- Postmortem analysis and understanding of system failures.
- Historical analysis of data and prediction of future evolution.
- Empirical approaches to demand modeling, and capacity measurement and management.
- Validation of analytic and simulation models.