Oracle Comments Respecting the Draft Report "Enhancing the Resilience of the Internet and Communications Ecosystem Against Botnets and Other Automated, Distributed Threats"

Oracle Corporation provides products and services that address all aspects of corporate information technology (IT) environments—applications, platform and infrastructure. Our products are delivered to over 400,000 worldwide customers through a variety of flexible and interoperable IT deployment models, including on-premise, cloud-based or hybrid, that enable customer choice and best meet customer IT needs. Our Oracle Cloud offerings provide a comprehensive and fully integrated stack of application, platform, compute, storage and networking services in all three primary layers of the cloud: Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). Our on-premise IT offerings include: Oracle Applications, Oracle Database and Oracle Fusion Middleware software, among others; hardware products including Oracle Engineered Systems, servers, storage and industry-specific products, among others; and related support and services. We provide our cloud and on-premise offerings worldwide to businesses of many sizes, government agencies, educational institutions and resellers with a sales force positioned to offer the combinations that best suit customer needs

We are pleased to have the opportunity to comment on the report.

1. NTIA questions and overview

We would first of all like to note that the draft is overall quite good. Throughout the report there are individual insights that, even when widely acknowledged, are often not stated as clearly or forthrightly as they are here. For instance, the point (p 12) that enterprise networks are both "simultaneously a victim and a source of risk" brings home the point that enterprises cannot just "hunker down" as a defensive strategy. And the suggestion (p 19) that we ought to respond to changing facts on the Internet in part by developing new vocations is a welcome. Similarly, the frank admission that most network technologies are general-purpose -- as useful for corporate networks as video games (p 20) -- nicely shows why special-purpose responses are futile. It is refreshing to have such frank acknowledgement of some challenges before us all.

In its request for comments[1], the National Telecommunications and Information Administration asks some questions. We answer them here in short form, with numbers and forward references to discussion below. The numbers are those appearing in the Federal Register; the entire question is not quoted in the interests of brevity.

1. Ecosystem

The description of the ecosystem is broadly correct, but might benefit from some consideration of the way the ecosystem is changing. Some of the relevant ways are outlined in section 2 and

[1] 83 Fed. Reg. 1342 (January 11, 2018) (Docket No. 180103005–8005–01, RIN 0660–XC040)

section 3 below. In particular, the report appears to make a fairly clear distinction between enterprise networks and more informally-managed networks, but does not talk in any detail about how the distinctions are collapsing; see section 3.1 in particular.

The report makes clear that automated, distributed threats include more than just threats from denial of service. The evolution of botnets into effective multi-purpose tools is an important development that network operators are just beginning to face, and the observation in this report is important.

2. Goals

We largely support the goals in the report, and believe the achievement of these goals would reduce threats to the Internet. We are pleased with the report's emphasis that the nature of the Internet means that no single actor can address the threats independently. We especially agree that some changes in both network operations and the marketplace of devices will be needed.

We are of two minds about actions 1.1 and 4.2, where there is an emphasis both on bilateral and multilateral action, and international standardization. Many of the Internet's benefits have been realized through voluntary standards rather than through treaties or formal international standards. It is not clear to us whether action 4.2 is intended to encourage international embrace of those kinds of voluntary standards, or whether the idea is to widen the international regulatory regime for the Internet. We believe the latter would be undesirable, and is inconsistent with a theme in the report. For some discussion see section 2.

We suggest an alternative approach to address the recommendations under Action 1.2 regarding software. We strongly agree with the draft report's problem statement that the use of software components, libraries and modules can potentially create security issues. However, rather than a narrow focus on one possible solution -- software transparency -- we recommend that NTIA take a broader view of this challenge and its possible solutions. We would strongly support NTIA in calling on stakeholders to discuss wide variety of solutions to address security challenges pertaining to software components. For instance, best practices that relate to vendors' internal processes, practices and policies governing their use, and vulnerability management of these libraries need to be further developed and promoted.

There may indeed be some cases in which some form of software transparency for sophisticated enterprise customers could be valuable; but even in those cases, transparency itself would be only one option of a broader suite of security improvements that stakeholders in the software ecosystem should seek to advance. In our view, transparency alone would usually not provide meaningful security improvements.[2] For these reasons, we suggest that

[2] First, such disclosures would in almost all cases give non-actionable information to the customer: the latter generally cannot fix the software, only the vendor can. Second, this information would often be misleading: the presence of a vulnerable library in a software product does not mean it is exploitable in that product – often it actually is not. That is the case when the implementation of the vulnerable library does not rely on the exploitable part of the

NTIA seek a broader inquiry of the improvements that are possible with regard to software components, libraries, and modules. Oracle would eagerly participate in and promote the recommendations of such an initiative.

3. Stakeholder roles

A core strength of the report is its strong emphasis on how it is not possible for any one actor to address all the challenges, and we like how the report outlines different actors' responsibilities. One thing the report does not make clear, and where it would benefit from some additional background, is that these different roles are not a contingent matter of policy. Instead they are a feature of the way the Internet actually works. The report would be stronger for making this point clearer. One way to do that would be to modify the passage on p 9 to read, "This response underscores both the *designed-in* interdependence of the infrastructure, and the ability of individuals and organizations to quickly learn and adopt." (Emphasis added to illustrate the addition.) Similarly, theme 1 (p 7) could be adjusted to read, "The majority of the compromised devices in recent botnets have been geographically located outside the United States, *which is not surprising considering the number of devices on the Internet and their international distribution*." These suggestions are for small alterations; for discussion and a more comprehensive suggestion see section 2.

4. Road map

The report's emphasis on envisioning a future for different parts of the ecosystem, and its approach of outlining goals and some actions in the service of those goals, seems to us to be better than a more detailed road map. Road maps have the advantage of being able to rely on a mostly static environment. But both the current security environment on the Internet, and the nature of the Internet itself, are changing so fast that detailed plans would be overtaken by events before the plans were even completed.

In our view, the most urgent actions are listed under goal 2, with action 2.1 of particular importance. It ought to be noted that industry and SDOs are already acting in this area; so what is needed is not the setting up of additional fora for collaboration or standards activities, but instead, additional and enthusiastic participation in efforts already underway, to make sure that everyone's needs and use cases are addressed. See section 4 below for more discussion.

5. Incentives

---

library. Third, such communication would actually often weaken rather than strengthen cybersecurity by diverting resources: customers would flood the vendor with inquiries about an announced library vulnerability, forcing the vendor to prioritize this patch (even though it may be a low CVSS score) to the detriment of patching other vulnerabilities that may higher CVSS scores. This problem may compound the fact that the vulnerability may not even be exploitable in the product: the vendor may have to prioritize patching a low CVSS vulnerability that is not even exploitable over a higher CVSS that is exploitable. (We note that vendors do patch vulnerable libraries embedded in still supported products even if they are not actually exploitable, but that non-exploitability is an important factor in prioritization and thus in allocation of resources in their security teams.)

The report is clear (correctly in our view) that effective responses to the threats will come from voluntary and co-operative action, but also states boldly (p 8) that market incentives are misaligned. One incentive that sometimes causes industries to adjust their behavior is a clear understanding of the consequences of regulation. Promoting the study of consequences from comprehensive regulation for the burgeoning device market may be a goal worth consideration. This point parallels the observations in the last paragraph in the text box at the top of p 23: liability and outright regulation are related issues, and it will not do to leave these topics unstudied. This does not mean that the report is incomplete without such study, but the text box might directly suggest that industry and the academic sector create models to understand what the consequences of various regulatory regimes would be.

Many of the anti-abuse mechanisms on the Internet to date, and certainly the most successful ones, have depended on reputation systems. The consumer market especially is almost completely unaware of how their own purchase decisions might affect their online reputation, even though consumers are highly sensitive to reputation and fashion in many purchases. This is another angle for action in the "convening power" theme in action 5.1.

The report does not emphasize end-consumer incentives as much as it might. In particular, the link between botnets and personal information compromise is touched on (p 19), but linked only casually. Yet many botnets are used specifically to harvest users' data and to interrupt services. While it is true that defending a home network against botnets will not be sufficient keep one's personal data safe, *not* defending against botnets is an excellent way to lose control of such data. Making that clear to the general public might improve incentives in favor of device security.

6. Metrics

Progress in foiling abuse is itself impossible to measure, obviously: one cannot tell whether an attack that did not come is because of good defense or because the attacker was not ready. But several of the goals amount to readiness of industry, products, or consumers. So, there are things to be measured, including the following:

a) Penetration rates into networks by network sector and by product type.
b) Product patch frequency, number of 0-day exploits, and frequency of patch before attacks are "in the wild" as opposed to discovered theoretically.
c) Size of the network-security and product-security employment sector as compared to the size of the internetworking sectors overall.
d) Comprehensiveness of security considerations in standards and product design, as measured by subsequent revisions to address weaknesses.
e) Size and effectiveness of information-sharing networks as compared to botnet growth.

In what follows, we expand on some of the themes we raise above.

2. The Internet and international boundaries

The report is extremely clear that Internet infrastructure is interdependent across organizations and across countries. The report is also clear that that interdependence should continue, but it

does not really explain why that is the case. Already we see various countries attempting to arrange access to the Internet along national lines, and without an explanation of why this is neither feasible nor desirable it is possible such efforts will grow.

The report appears to be built on the basic understanding of the Internet as a network of networks. But that usual formulation, while entirely correct, appears not to make plain to everyone a key consequence. The Internet requires interdependence and poor correlation with national boundaries not as a political fact, but because those features are fundamental to the very functioning of the technology. Each operator of a network operates it to serve the operator's own needs, and in large organizations those needs can easily cross international boundaries. Each such network is an autonomous system. When one connects those autonomous systems to create an internet, the resulting network of networks *automatically* has an international character. Similarly, the global Internet is efficient because it depends on protocol conformance to provide easy and cheap inter-network packet routing without every participant in the network needing to have a contractual relationship (either directly or indirectly) with every other participant. This ease of communication and low contractual burden is part of what has made internetworking technology so successful, and what is now delivering new efficiencies and models for business and wealth creation.

Therefore, it is not merely the case that "solutions specific to particular countries or jurisdictions put at risk the global nature" (p 20) of the Internet; but that such solutions will not work. Because the Internet is *essentially* interconnected, those kinds of solutions will impede conforming activities and hinder legitimate business practices, without doing anything to those attacking the network. For the same reason, it is not some political stance on the part of technical and business experts when they say ingress filtering at the level of international backbones is a bad idea (p 10). Instead, it is a bad idea because it would attempt to divide up internetworking along lines that are not consistent with the nature of internetworking: one cannot save the Internet by destroying it. Finally, the need for security standards that are "flexible, appropriately timed, open, voluntary, industry-driven, and global in nature" (p 16) comes not from some political program or preference for open networks, but from the very fact of voluntary interoperation that makes up the Internet.

Only that kind of voluntary interconnection has ever delivered us a technology that grows and scales the way the Internet has done. The Internet provides opportunities for US citizens, for US workers, and for US businesses unlike any previous technology, precisely because of its open nature. We must not permit the fear of attacks on the Internet to lead us to undermine the voluntary interconnection that makes it strong. That would hamper the innovation and commercial opportunities that companies like Oracle need.

Recommendation: The final report should make these points unmistakably clear as a fundamental precept of cybersecurity policymaking. Background asides of a similar nature are found in the report in a text box. A similar approach would work for this background point. Alternatively, we provide two small adjustments to make the point explicit, without emphasizing it, in our response to question 3, above. Oracle is eager to work with you to flesh out these essential issues in the lead-up to the final report and beyond.

3. The nature of kinds of networks

The report quite correctly differentiates among different kinds of networks (p 9). Plainly, there are differences among enterprise networks on the one hand, and home and small business networks on the other. But there are two critical items to observe that may impinge on the future network environment.

3.1. Cloud computing and the erosion of network-type distinctions

The first is one the report notes: the distinctions among networks decline as the consumer-grade devices gain capability. This tendency is what, for instance, has caused the explosion of "BYOD" policies in enterprise networks that, in the past, would never have allowed employee-supplied devices to connect to the corporate network. The report correctly observes that changes in the enterprise environment are partly due to the sophistication of consumer-directed devices. It also notes, again correctly, that better enterprise security practices around network devices paves the way for better consumer-grade security. Yet the report should clarify the additional leap: the distinctions among these kinds of networks was always somewhat artificial; and the extent to which the distinction was useful has already eroded, and will erode further as devices get more capable.

This observation is, indeed, precisely why the report is correct in noting that enterprise networks are at once a source of new vulnerabilities, even as they are themselves vulnerable. Yet the report, especially in its "Vision for the Future of Enterprise Networks", does not completely embrace the central fact: the "internal/external" way of thinking about networks is now largely obsolete.

Cloud computing is frequently described as though a traditional segregated-network model of computing were still in effect, and that the only difference lies in moving the boundary between "inside network" and "outside network". But the real gains from cloud computing come from realizing the plastic boundaries: "my networks", "our networks", and "all networks" in the cloud can be cast as questions about credentials and access, rather than network location. The cloud is about service availability and delivery, not about network nodes.

But the cloud model of computing means that "inside" and "outside" cannot be first-class distinctions. Imagine two services, A and B, delivered in the cloud. Each of them may change IP addresses, service region, or deployment strategies at any time. Yet they must interact with one another, which means that neither can have fixed firewall rules of block lists unless it is willing to endure an outage with the other. Scale this to the scope of the Internet, and it is obvious that the "internal/external" distinction must give in.

If this is correct, then the fairly hard distinction in the report among enterprise, SMB, and home networks begins to collapse. More importantly, perhaps, the distinctions among edge nodes in any of these kinds of networks starts to fall apart, unless there are automatic ways of determining what an "edge" is at any given time. And perhaps most important, enterprise networks and networks delivered by companies such as Oracle cease to be different things.

Recommendation: Without adding a significant section on cloud computing to the report, it might be possible to address this point by adding some text to the discussion of Enterprise

Networks, starting on p 12. To begin with, a paragraph along the following lines might set the stage:

> A clear distinction between the inside of a corporate network and the outside has always faced practical difficulties. While the network architecture might be clear about the boundaries, many networks have made inevitable compromises as originally-internal services were exposed to devices outside the network, and as merger and acquisition activities caused formerly-external networks to become internal. The advent of cloud services puts even more pressure on the distinction, because cloud computing often depends on flexible networking arrangements in order to deliver on a promise of increased reliability at reduced cost. Even if "locking down" an internal network was once possible, demands for both high reliability at low cost and very fast deployment means that tight boundary controls will not be realistic for many enterprise networks in the foreseeable future.

Recommendation: Another paragraph added to Vision for the Future of Enterprise Networks (starting on p 13), before the final paragraph along the lines of the following would be valuable:

> The policies also must be applied with the knowledge that the distinction between the enterprise network and other kinds of networks will begin to blur. Consequently, policies cannot concentrate on network nodes alone, but must concentrate on data flows and protecting enterprise data even when it leaves the enterprise network.

The collapsing distinction we note is implicit in the report's embrace of the proposed Manufacturer's Usage Description (MUD). There is no practical way to scope MUD to different kinds of networks except by constraining capabilities either of the device, or the network. Such constraints may be desirable, but they are equally desirable regardless of the kind of network in which the device appears. Given that the devices can appear in any network, then, efforts should concentrate primarily on those that deliver results no matter which kind of network is involved.

Recommendation: With this in mind, Action 3.2 should not be scoped only to home IT and IoT products, but to *every* IT and IoT product. Secure-by-default should be the stance of every product, even ones aimed at experts. If the experts want to turn such features off or engage functionality in a more specialized way, they can do so by overriding the defaults.

3.2. Will the home network be a full part of the Internet?

As the report notes (e.g. p18), many "home networks" (really, those of residences and also small and medium sized businesses) live behind a single IPv4 address, with Network Address Translation (NAT) used to map the (increasing number of) devices, all connected using RFC 1918 address space, into the single IPv4 address. The report suggests (see e.g. note 39, p 18) that NAT contains a kind of benefit, because it often ends up working as an accidental firewall. The report goes on to note that IPv6 may change some of these facts, and Action 3.4 urges investigation of the consequences for attack and defense.

Recommendation: Action 3.4 is laudable, but we would revise it to go slightly further and in order to investigate the potential for truly self-administering networks that operate at a fairly high level of sophistication. In effect, home networks need to become first-class participants in

the Internet, just as enterprise networks participate today. But because (as the report correctly notes) most such small networks will not have professional administration available, they will need to operate safely all on their own.

Home networks have relied on NAT as they have partly because of history. Early access to the Internet was usually over dial-up connections. This meant that the access point of the customer was always one phone line, which necessarily corresponded with one connection. That model was maintained when broadband connections became common. Even if ISPs had wanted to provide a different model, the shortage of IPv4 addresses by that point meant that they could not. But consumers, also, did not want to pay extra for additional IP addresses, and so "home networks" came to be things that were not really first-class participants in the Internet. The "edge" of the Internet conceptually stopped at the customer premise equipment (CPE). NAT traversal technologies such as STUN, TURN, and ICE have made that edge somewhat less sharp, but in general home networks have been like large clients of the Internet.

It is not plain that this state of affairs will persist. Already, various kinds of smart devices represent a kind of service delivered from inside the premises: smart meters, appliances that deliver maintenance and usage information to manufacturers and utility companies, and even IoT devices all act as kinds of servers, providing the relevant service to outside consuming devices. As a matter of implementation today, the data is often "served" over an HTTP client connection to a web service, often operated by the equipment vendor. But while such arrangements are good for the vendor, because they encourage lock-in, we see two problems. First, from a consumer's perspective the real value of Internet of Things is likely to come from the connected things acting in concert with one another. Even if it were a practical possibility, it is unlikely consumers will commit to a single vendor for all the things in their houses and cars. So, interoperation among the connected things seems like an eventual outcome. Second, a connected thing that fails every time the vendor's service is unavailable will not be very useful. So, we can expect the things to interact with one another, and possibly with arbitrary services on the Internet.

Making all of that work reliably will require some services to be inside the home network, and such services will work more reliably over IPv6 than they will with IPv4, NAT, and various NAT traversal technologies. So, we may suppose that the increasing availability of IPv6 encourages the use of IoT devices, which will then promote more IPv6 adoption, and so on in a virtuous circle.

If that is true, however, then the need for automatic network management that builds on something like MUD becomes even more urgent.

Recommendation: Pursuant to this urgent need, the report should call for work from industry and standards bodies to make safe automatic network management practical. Oracle will eagerly engage in these efforts and welcomes additional stakeholder and government activities. To achieve the ends outlined above, this will need to be more sophisticated than simply putting a stateful firewall at the CPE device for IPv6, to mimic the role of stateful NATs in use with IPv4 today. There have been some initial moves by standards bodies in this direction, such as some

Internet Engineering Task Force working groups[3], which suggests that industry and standards bodies are already interested in this topic. Perhaps such efforts can be encouraged to reach for more ambitious goals.

At the same time, everyone involved in networking should remember that the "end to end" argument[4] applies actually to *applications*, not network nodes[5]. Some nodes in the network are "dumb" on purpose, and we will not get the benefit we want from the network if we attempt to treat every node the same. For use in networks without professional administration, this makes automatic network management again more important.

4. Development of standards and practices in industries

The report does an excellent job of recognizing the global nature of the problem and the marketplace and consequently emphasizing the role of international standards to establish globally recognized technical frameworks, methods, and processes. We offer the following feedback and proposed changes to strengthen this view and Actions such as 1.1.

A. Recommendation: Strongly urge stakeholders, particularly those from markets not traditionally oriented around global networks, to participate in standards processes. Standards and best practices are most useful when they are developed with the broadest possible participation, so encouraging stakeholders to get involved in the efforts already underway would be a desirable outcome.

   Such encouragement of participation in existing bodies is especially important for new arrivals to the Internet of Things marketplace. Many of them come from traditional industrial- or consumer-goods environments. In those cases, these participants may not be used to thinking in global-network terms, and may find the approach for handling IPR and interoperation questions foreign. They may be tempted to attempt to move the standards development to sector-specific bodies. As the report notes, however, most of the technologies in question will not honor such industry boundaries, so the standards development cannot either.

B. Recommendation: NIST, NTIA, and other relevant U.S. government agencies should increase their already prominent activities in international standards development. As the goals and actions in this report are carried out, we believe it is vital for government to encourage, support, and participate in standards development. Government agencies

---

[3] See in particular the Autonomic Networking Integrated Model and Approach (ANIMA, https://datatracker.ietf.org/wg/anima/documents/) and Home Networking (HOMENET, https://datatracker.ietf.org/wg/homenet/documents/) working groups.

[4] Saltzer, J.H., D.P Reed, and DD. Clark, 1984. "End-to-end Arguments in System Design." *ACM Transactions on Computer Systems* 2 (4): 227-288. November 1984. doi:10.1145/357401.357402.

[5] For more on this point, see Sullivan, A, 2017. "Avoiding lamentation: to build a future Internet." *Journal of Cyber Policy* 2 (3): 323-327. 2017. doi:10.1080/23738871.2017.1400083.

such as NIST and NTIA play a valuable role when they act as conveners to coordinate, identify, and communicate federal government views regarding emerging technologies. And for developing requirements or best practices, government agencies are invaluable partners in market-driven standards committees.

Government should not, however, initiate activities that compete with private sector standards development organizations. In addition to straining industry resources to participate, the output of such activities is regarded as a US government document and therefore not usually recognized as widely as international standards. Further, it provides a justification for other governments to undertake the same activities, fracturing the marketplace and creating a burden on producers and consumers alike. Some actions that the USG chooses to take - with the intent of making positive progress - may be perceived in other countries as putting stakes in the ground and motivate similar acitivies in their countries. The USG should keep this potential burden on industry in mind when it implements actions such 1.2 and 1.4 calling for government and industry to collaborate.

When government efforts do produce frameworks or requirements, the ultimate goal of the effort should be to rely on or produce international standards. To this end, the USG should consider if/when it would be appropriate to start discussions about new work in international fora in parallel to USG efforts. For example, an international activity that set out to agree on higher level layers of a model or high-level concepts in a framework could help align efforts across countries, while also allowing each country to tailor their model or framework to their own needs.

C. Recommendation: Strengthen actions in the report by noting activities already being undertaken in the standards and operations communities. In particular, Actions 2.1 and 4.5, and possibly Action 4.4, might be strengthened by calling attention to the DDoS Open Threat Signaling (DOTS – see https://datatracker.ietf.org/wg/dots/documents/) work that has been undertaken at the Internet Engineering Task Force (IETF). This kind of effort can provide an open standard for automatic interoperation among different vendors.

D. The report spends some time on ingress and egress filtering: BCP 38 (p 10 ff). It could benefit from noting, first of all, that there are already efforts to promote good routing practices (such as the Internet Society's Mutually Agreed Norms for Routing Security or MANRS – see https://www.manrs.org/). In addition, the report should note that filtering alone will not solve the problem. Historically, attackers used address spoofing as one of their primary tools, and so BCP 38 adoption was a critical item. But the sheer number of compromised devices means that attackers do not really need to spoof any more to be effective: they can simply use the addresses of the many compromised devices, without hiding them. Recommendation: The report would therefore benefit

from making an explicit link between the large number of compromised devices and the degree to which filtering alone is ineffective.

E. Actions 5.1 and 5.2 suggest standards and labelling regimes to help consumers. Any regime must ensure both that the products bearing the label actually undertake practices that will be helpful for the consumer, and that consumers can understand what the label indicates. Keep labeling simple by making it digestable by consumers. Keep it actionable with categories that have an effect on the market both through purchasing choices consumers can make and pricing decisions producers can make. One example of this would be an indication of who manages the device and its security: is the device managed by the consumer directly, or does the device come with support for some duration? Recommendation: NTIA should convene industry and consumer experts in a process to determine effective and meaningfully informative informational tools for consumers. This should build on the consensus recommendations regarding patching/upgradability that stakeholders developed through the most recent NTIA multistakeholder process. Oracle is eager to engage with other stakeholders to take concrete steps to advance the market in this regard.

5. Vendors of mitigation also create some risks

It is certainly true, as the report notes (pp 10-11) that a robust DDoS protection market can help network operators (especially enterprise network operators) face the threats against them. It is also good to see the recommendation of hybrid approaches. The discussion might be improved by noting especially a danger to the Internet from overreliance on a large degree of filtering in case a concentrated market emerges. There is a danger that the network may come to be too centralized if everyone uses the same techniques and vendors. This risk is especially acute as cloud operations become more common. Operational concentration presents risks, and enterprise networks in particular will need to understand those risks when making their decisions about network service vendors.

Recommendation: Articulate the affirmative risks of overreliance on filtering and network centralization.

6. Conclusion

Oracle is committed to a stable, useful, open Internet – one that continues to offer its many opportunities based on voluntary interoperation, regulation primarily restricted to areas already regulated, and interconnection without a large degree of contractual overhead. We are therefore pleased that the report makes recommendations that continue to support such an Internet. We congratulate the authors on an excellent draft, and hope that our comments prove useful during final revisions.