

Tracking Assets and Vulnerabilities in Corporate Networks

by Dieter Gantenbein

In today's dynamic information society, organizations critically depend on the underlying computing infrastructure. Mobile users and computing devices add to the challenge. Business operations, intellectual property, and corporate value are quickly at risk. The IBM Zurich Research Laboratory has developed a tool called Intelligent Device Discovery (IDD) that collects information from all possible sources, computes an aggregate picture of assets, and categorizes their usage and vulnerabilities.

Tracking computing devices as assets with usage, health, and vulnerability information facilitates the provision and maintenance of an efficient, optimized service. Recent incidents with viruses and worms appear to support this. In general, a precise understanding of the operational infrastructure is a key element of many corporate decisions. Examples also include the negotiation of outsourcing contracts, the planning of mergers and acquisitions, server consolidation and business optimization.

Building an accurate inventory of computing assets in heterogeneous dynamic systems and networking environments is difficult, especially when only limited privileges are available and no prior device instrumentation has taken place.

Classical methods for inventory and asset management quickly reach their limit in today's dynamic environments: Periodic physical inventories ("wall-to-wall") have the distinct advantage of identifying the actual location of devices but require costly human visits ("sneaker net") and cannot detect mobile, currently out-of-office equipment nor the existence and use of logical assets. Financial asset tracking, although an accepted process in its own right, cannot detect additional equipment brought into or accessing the resources of an organization. Periodic self-assessment questionnaires to be filled out by individual end users or their cost-center managers are another, often complementary approach. Apart from the human effort they require and the inaccurate, incomplete data they may contain, most forms pose questions the answers to which could easily be obtained from the infrastructure itself.

Well-managed computing infrastructures typically equip servers and end-user devices with software daemons to track the system's resources and health. There are many situations, however, in which these daemons cannot be relied upon. In many organizations, there are a fair number of devices that are brought in ad-hoc without appropriate instrumentation, for which instrumentation is not available, or on which instrumentation has been disabled.

It therefore appears advisable to complement process and policy-based asset and security management with "automatic sensors" to recalibrate the dynamic and heterogeneous environment. This is where IDD fits in.

Scenarios

Your organization may already have detailed inventories, but what about the IT infrastructure of the company you interact with? How can you quickly substitute in any missing data?

How many old Windows systems with less than 512 MB of memory do we need to upgrade this geography? Where does this device with a strange MAC address come from? Which machines are unsafe, i.e. without the latest patches and not protected by antivirus and security software?

Computer A is a safe machine used by a secretarial employee on the mornings of 4 out of 7 days a week. Computer B - perhaps a student-lab workstation - has many server ports open, is not completely safe, and is shared by users X, Y and Z. Some server has local unprotected databases and no screen password set. A portable device currently in office C has no hardware password set.

And those rogue servers should either be turned off or turned into a policy-based operation. We may want to send emails to those employees whose machine configurations harbor discovered or accepted risks. Discovery provides the means to detect early potential exposures and track down the exact location of assets in violation of current policies.

Information Sources

Corporate networks of large companies span geographies, the individual locations of which typically consist of separate segments for user access, servers, as well as intranet and extranet connection zones. The operation of wide-area and/or campus networking, servers, and/or user devices may be handled by different organizations. In general, there are neither universal administrators nor "one size fits all" credentials. In order, nevertheless, to be able to derive an overall picture, we propose that data be harvested from whatever source possible, and then combined into an integrated information model. The range of information sources considered encompasses network-based autodiscovery, privilege-based additional access to on-line network and application subsystems, and the lean processing of manually provided ledger tables containing financial and physical inventories and configurations. Bothering the end user is considered only as a last resort.




Figure 1: Networks, communication paths, and sources of information.

Techniques

The sources of information are manifold, and the techniques to access them are even more so. The communication stack and the device ports can be actively mapped by scanning from administrative machines, including IP ping sweeps, UDP/TCP port scans, and remote Windows SMB/Registry/WMI fingerprinting. Should established policies limit the yield of this technique for some of the target devices, analysis of subnet traffic can also passively map otherwise stealthy devices. Although this approach assumes probes on appropriate segments and VLANs, the trend to corporate netflow architectures that collect aggregated traffic accounting records from strategic subnets provides for a more strategic access-point for network flow information. The communication traffic among devices can also be observed in the the logs of networking services, including (DNS) name servers containing directory information about registered resources, DHCP address servers with lease information, WLAN wireless access servers, and PIX firewalls. Some of the related events may also be readily available at syslog servers. Whereas managed end-user systems may be instrumented with agents that interact periodically with corporate management servers, e.g., to keep software up-to-date or to download new virus and worm defense

policies, there is also the entire realm of other application servers for mail, group communication, Web applications, mobile-user access portals, etc., that may provide detailed information about usage patterns.

Conclusions

IDD is a network-based IT asset discovery and categorization tool. It combines various device discovery, network and security scanning programs with enhanced data collection, mining techniques, and distributed automation to form a single application built upon the IBM e-business platform with WebSphere and DB2 on Windows and Linux. It is used today by consulting, outsourcing, and security teams on internal and customer networks.

Link:

www.zurich.ibm.com/idd

Please contact:

Dieter Gantenbein, IBM Zurich Research Lab
Switzerland
IP phone: +41 1 724 88 53
e-mail: dga@zurich.ibm.com