ghaff@
illuminata.com

Gordon R Haff
Illuminata, Inc.

The End of Cobalt and the Appliance Era that Never Was

Quick Note

Gordon Haff
14 January 2004

Like Zeppelins filling the skies in some futuristic Gernsbackian vision from 1930's pulp fiction, server appliances were supposed to crowd the datacenters and remote offices of a future in which the Internet boom never ended and the need for specialized net servers never slackened. Alas, neither prediction ever came true.

It's not that hardwired, function-specific server appliances didn't help solve some very vexing problems; it's just that they proved to be too specialized for functions that were changing so rapidly that even general-purpose servers could hardly be modified quickly enough to keep up. The first generation of server appliances took many forms and ran many different types of applications. All were focused on a particular task (or two) and were designed to be ready to install with minimum muss, fuss, or skill. This assembly line approach to server farms was to be the secret sauce that made possible infinite growth without infinite IT staff.

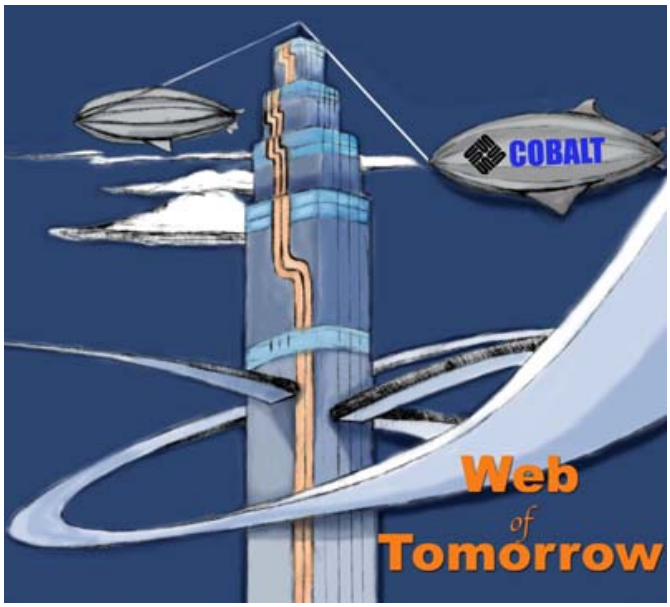

Cobalt Networks was perhaps the best known and most sophisticated of the companies offering these types of products. Sun Microsystems eventually bought Cobalt and tried to integrate it into its product line—but those fitful efforts are now over.

As Sun winds down what remains of Cobalt, we're hearing the predictable emotional rants that Sun mishandled the acquisition and failed to nurture or develop Cobalt's products or people. Sun certainly didn't get its money's worth out of Cobalt, even if the $2.1 billion it spent was in its then-high-flying stock, not cash. At the same time, however, it wasn't just Sun that failed. Server appliances as a class went down in flames. Though it's tempting to place all the blame squarely on the economy, that's not the whole story, either. Rather, the very specialization and streamlining of function that was their *raison d' tre* proved to be less applicable an approach than the adaptability and flexibility of general purpose systems.

As Cobalt and the server-appliance dream fade away, it's worth considering why the idea didn't succeed in the mainstream, why it survives in certain niches, and how the server appliance is being reinvented in a new form today.

Sic Transit Cobalt

Cobalt was the most visible of the early appliance vendors whose pitch was appealing because they focused specifically on serving up Web pages at a time when the World Wide Web was still new and more than a little bit exotic. Cobalt soon surpassed other early entrants like Data General's THiinLine division and Network Engines. Certainly its servers were stylish and well-designed, whether for the datacenter (RaQ) or the small office (Qube). But it was Cobalt's sophisticated software that underpinned its success. Although Cobalt was one of the earliest vendors to embrace Linux as a fundamental platform OS, Cobalt never ran around trumpeting Linux as the source of unassailable virtue in its own right—unlike many Linux-oriented companies that did so during the Internet boom. Cobalt also clearly understood that, even within an appliance model, adaptability and customizability had to exist alongside simplicity.

Sun, viewing itself as the company that "put the dot in .com" and the platform underpinning of the Internet build-out, probably saw Cobalt's simplified appliance approach as just what was needed to let the build-out proceed apace. After all, this was a time when companies were citing the dearth of skilled workers as an impediment to growth, and appliances addressed that lack by requiring less attention than most servers. But Sun seemed never to take advantage of its pricey acquisition, even before server deployments slowed to the point that appliance simplicity became a less pressing concern.

Perhaps sensing that they and the Cobalt technology were sideshows at Sun, top Cobalt executives like CEO Steven DeWitt soon left.[1] One problem was that Sun's resistance to, and suspicion of, Linux—even in the under-the-covers incarnation that powered Cobalt's servers—kept Sun from making much use of Cobalt's Linux platform expertise. Sun steadily de-emphasized both Cobalt's various products and then the brand name itself. It's been clear for about the past year that the Cobalt Qube and RaQ product lines had little future, partly due to Sun's lack of enthusiasm, and partly due to the market's. Sun now plans to formally discontinue them.

The First Generation

Adjectives like non-programmable, pre-configured, and sealed were among those commonly used to define early server appliances. The term "appliance" was itself a nod to home appliances designed to serve one and only one function—a tradeoff between flexibility on the one hand and ease of operation and lower costs on the other. However, most vendors moved away from such absolutist views of what makes a server appliance in favor of a broader definition that allowed for many approaches to simplifying installation and use.[2]

But even the least appliance-like units shared certain characteristics of hardware, software, and the means by which software was delivered or packaged, that were significantly different from general-purpose servers. Choices within each of these dimensions represented potential tradeoffs between often-conflicting goals, including ease of use, flexibility, physical density, and cost. For example, a server appliance designed to run a specific application that was accessed through a simplified browser interface might have been a good match for users with few computer skills, but it tended to be less amenable to upgrades or changes by more sophisticated users with more sophisticated needs.

First among those to jump on the appliance bandwagon were small vendors founded to exploit the apparent market opportunity, including e-Server, InfoBlox, LinuxLabs, Mirapoint, Network Engines, PictureIQ, Rebel.com, StarBak, Vbrick, Vividon, and WireX. Only a few still exist. The large system companies also joined in—whether by tossing appliances in the product mix of their in-house systems business, relying on partners, or—as in the case of Sun—making a major acquisition.

1. See Illuminata report "The Changing of McNealy's Guard" (May 2002).

2. See Illuminata report "Appliances—They're Not About Toasters" (August 2001).

Though it's not usually the case, a slow approach did more good than harm for Dell, HP, and IBM, who were late and hesitant with their own appliances. Compaq and Sun, on the other hand, invested more heavily and pursued an appliance strategy with more vigor, and therefore wasted a lot more money and energy. By about late 2002, all the majors had sidelined or shuttered their appliance-specific divisions.

Broad Ambitions, Narrow Niches

The first generation of server appliances was intended for three general locales: the network edge (whether in enterprises or smaller businesses), enterprise server farms, and SMBs that couldn't afford the staff to maintain more general-purpose hardware.[3]

Vendors designed a broad range of appliances, trying to fit one to every need and every niche in each of those market segments. The level of specialization and variety led the appliances' more ardent promoters to predict that datacenters would soon be built up, in large part, using specialized-function server appliances.

But if the appliance concept promoted during the boom wasn't totally discredited, it didn't capture much mindshare in the mainstream either. A few smallish companies still sell appliances for mid-tier datacenter applications (e.g. Mirapoint's Message Server for email), but most appliances have been relegated to the network edge as firewalls, load

3. Storage devices like file servers can also be thought of as a type of appliance.

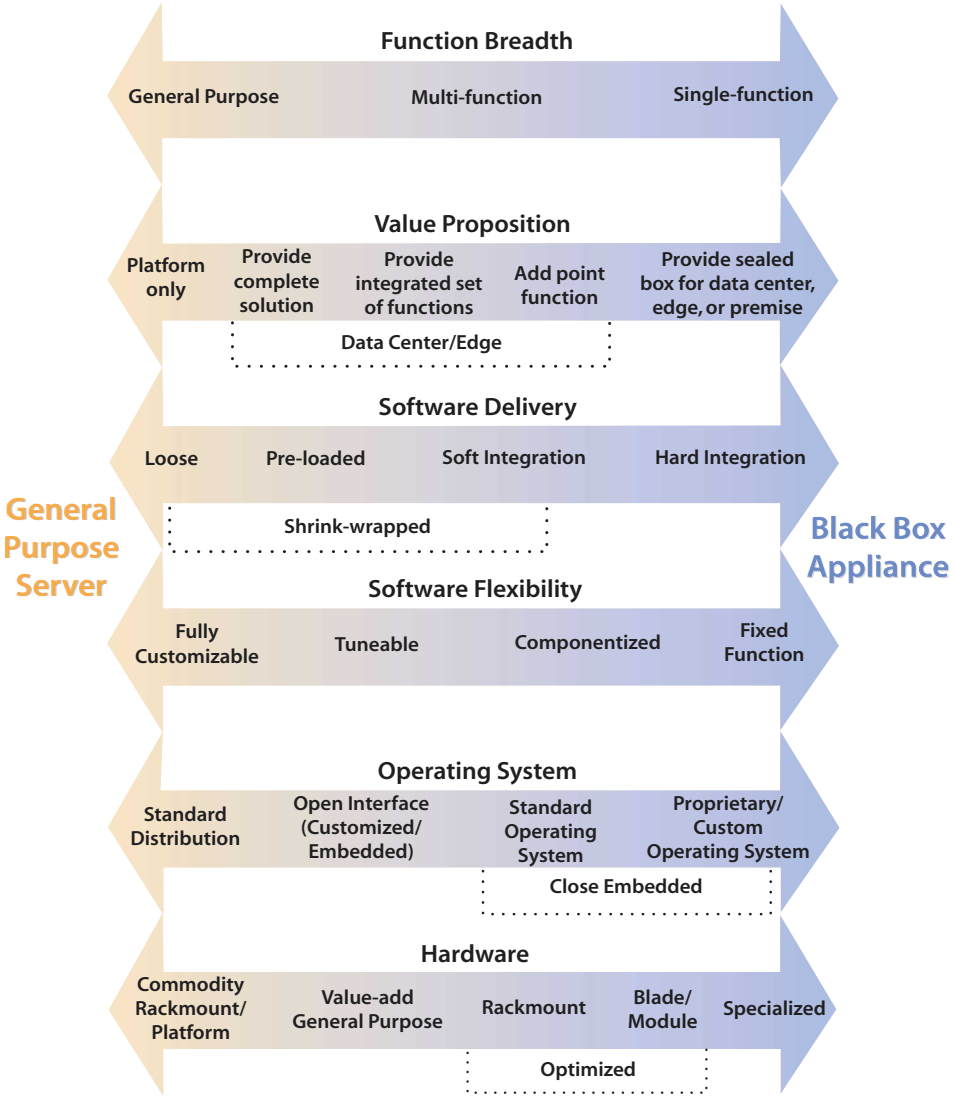

Gordon R Haff
Illuminata, Inc.

balancers, and other extensions of the network infrastructure. Customization tends to be less important in this type of device. Indeed, it can be a negative insofar as edge-of-the-network functions tend to be fairly narrow and well-defined, and twiddling can easily end up compromising both correct operation and security.

In addition, many network edge devices—such as those used for content acceleration—tend to reside in distributed locations where in-depth IT knowledge may not be available.

However, the appliance-on-the-network-edge is a pale shadow of the grand appliance dream. General purpose hardware has largely won—as it so often does. Even relatively flexible appliances are inherently less malleable than a general purpose server. If they weren't, they wouldn't be appliances—at least as the term has traditionally been used. And this comparative rigidity makes it more difficult to augment a Web server with additional software that a site may want to use internally or externally. General-purpose servers running general-purpose software are more receptive to a modification such as adding a set of scripts to manage bulletin board-style discussions which, in turn, require a new version of PHP or Apache that, for its part, requires an update to Perl.

Furthermore, fully buying into the appliance concept could mean purchasing many types of specialized servers—many running on unique hardware platforms with their own spare parts, service needs, and management interfaces—from multiple companies. The problem of how to bring together this hodge-podge of pre-integrated and largely immutable systems into one maintainable whole was never solved or even really acknowledged by appliance zealots. Having hardware and software systems that were so simple they didn't

	Purpose	*Typical functions*	*Appliance acceptance*
Edge Devices	Extension to the network that regulates the interface between core server and storage devices and external networks.	Proxy servers Directory services Caching servers Caching servers for streaming media Firewalls Load balancers	Fair to good.
Server Farm	Primarily delivery, transformation, and management of data delivered from the middle business logic/ application tier and the back-end database tier.	Traditional Web serving Streaming media delivery Image preparation and delivery XML acceleration SSL acceleration Collaborative Application specific (e.g., e-Commerce) Database Management	Poor. Limited examples remain in security and other areas, but lack of customizability effectively smothered appliances in this role.
SMB	A targeted set of functional requirements (e.g. e-mail, intranet) and a greater need for simplified deployment that minimizes or eliminates the need for IT expertise.	A multi-function device that provides key functions such as e-mail, Web serving, and file serving-in addition to a consolidated and secured internet access point.	Poor. Use low-cost network edge appliances (e.g. routers from Linksys, D-Link et al) but not multi-function appliances à la Cobalt Qube.

need to be configurable was, after all, the whole reason to go the appliance route in the first place.

Yet many of the factors that gave birth to the appliance concept—the cost to install new servers, difficult configuration, and the cost of errors—remain. The appliances just didn't provide the right tradeoffs.

The Coming Soft Appliance

Appliance makers tried to achieve simplicity by opting for specialization that limited the ability to adapt. They used expensive low-volume hardware, or made customers dependent on some highly customized software stack, for example. For most buyers, for most purposes, those tradeoffs against general-purpose hardware running general-purpose software were too great for the potential benefits.

The best tradeoff is the one you don't need to make. This led datacenters—especially the largest ones (think Google and its ilk)—to use a variety of homegrown and third-party tools to create and manage "appliance-like" software loads on low cost, flexible, standard servers. They could burn images onto new servers and reconfigure them as required. In some cases, the server vendors themselves even delivered pre-imaged systems to a customer as part of a custom services engagement.

That need, and the eminently sensible approach Google and others have taken to address it, is driving server vendors toward a more comprehensive and systematized approach under the general rubric of "provisioning." Rather than defining it as an important function on its own, vendors like HP, IBM, and Sun treat provisioning as an important underpinning of their grand virtualization and automation schemes. But while those companies' enterprise-spanning visions for computing still have many years of evolution ahead of them, the provisioning component is proceeding apace through a combination of internal development, acquisitions, and partnerships.[4]

Provisioning does more than just help manage large server farms or speed rollouts, however. It effectively provides the promised simplicity of appliances while maintaining the advantages of general-purpose servers.

The details vary from product to product but, fundamentally, provisioning software can create an image such as a "basic Web server" that includes not only the Linux OS and Apache Web server, say, but also any desired site-specific data files, tweaks, or settings. This image is then laid down on a general-purpose server (or blade) automatically, creating a Web server without the need to load and hand-configure each system individually. This Web server image is effectively an appliance in that it's standardized and can be easily replicated as often as needed. But it's a "soft" form of appliance rather than the harder, less mutable form that earlier appliances took. The IT department is in control of the image contents, not the appliance vendor, and can customize the software load as much as needed—but only needs to do so once. Hardware selection isn't tied to the software either.

What's more, once configured, the "basic Web server" doesn't need to remain a Web server forever. It can be replaced with some other software image—and therefore a completely new function—as the needs of the enterprise and its datacenter operations change over time. For example, a Web site that has to add servers to support spikes in video traffic could shift some machines from serving static pages to serving video. In the ultimately automated, policy-based visions of the major system vendors, server roles will shift dynamically in response to the changing demands of the moment—something that couldn't be done at a site using one vendor's Web server and another's video streaming server. This is a shift away from the decades-old production IT philosophy in which servers, once set up, weren't touched unless a touch was absolutely required, and toward the much more

4. See Illuminata reports "IBM Thinks Dynamically" (May 2003) and "Sun Accelerates N1 Development with Terraspring Acquisition" (November 2002).

fluid configurations associated with the telecommunications industry—which effectively invented the provisioning concepts that the computer companies are now co-opting and adapting for their own purposes.

Conclusion

The first server appliances were point products. For some narrow applications such as firewalls that fundamentally sit apart from the rest of the computing infrastructure, point solutions can be appropriate. But, in general, IT managers not only prefer to avoid special-purpose hardware that is tied to special-purpose software, they are entirely right to do so. This is particularly true as we move towards a model of IT infrastructure that favors flexible and dynamic apportioning of resources while simultaneously reducing the amount of *manual* configurations and manipulations required.

In one sense, the Server Appliance Era is one that never was. Especially with Sun bowing out of the Cobalt business, appliance-like devices now survive only as boutique, specialty products, or in certain narrow niches on the network edge. But, in another sense, the new server provisioning future will make every server an appliance, but a soft one, that brings together the best of general-purpose servers and more specialized ones.