Approaching a Breakthrough p. 7 Being a Good Executive Sponsor p. 32

Winter 2011 Edition • Volume 30, Issue 4

SOFTWARE MAGAZINE

The Software Decision Journal

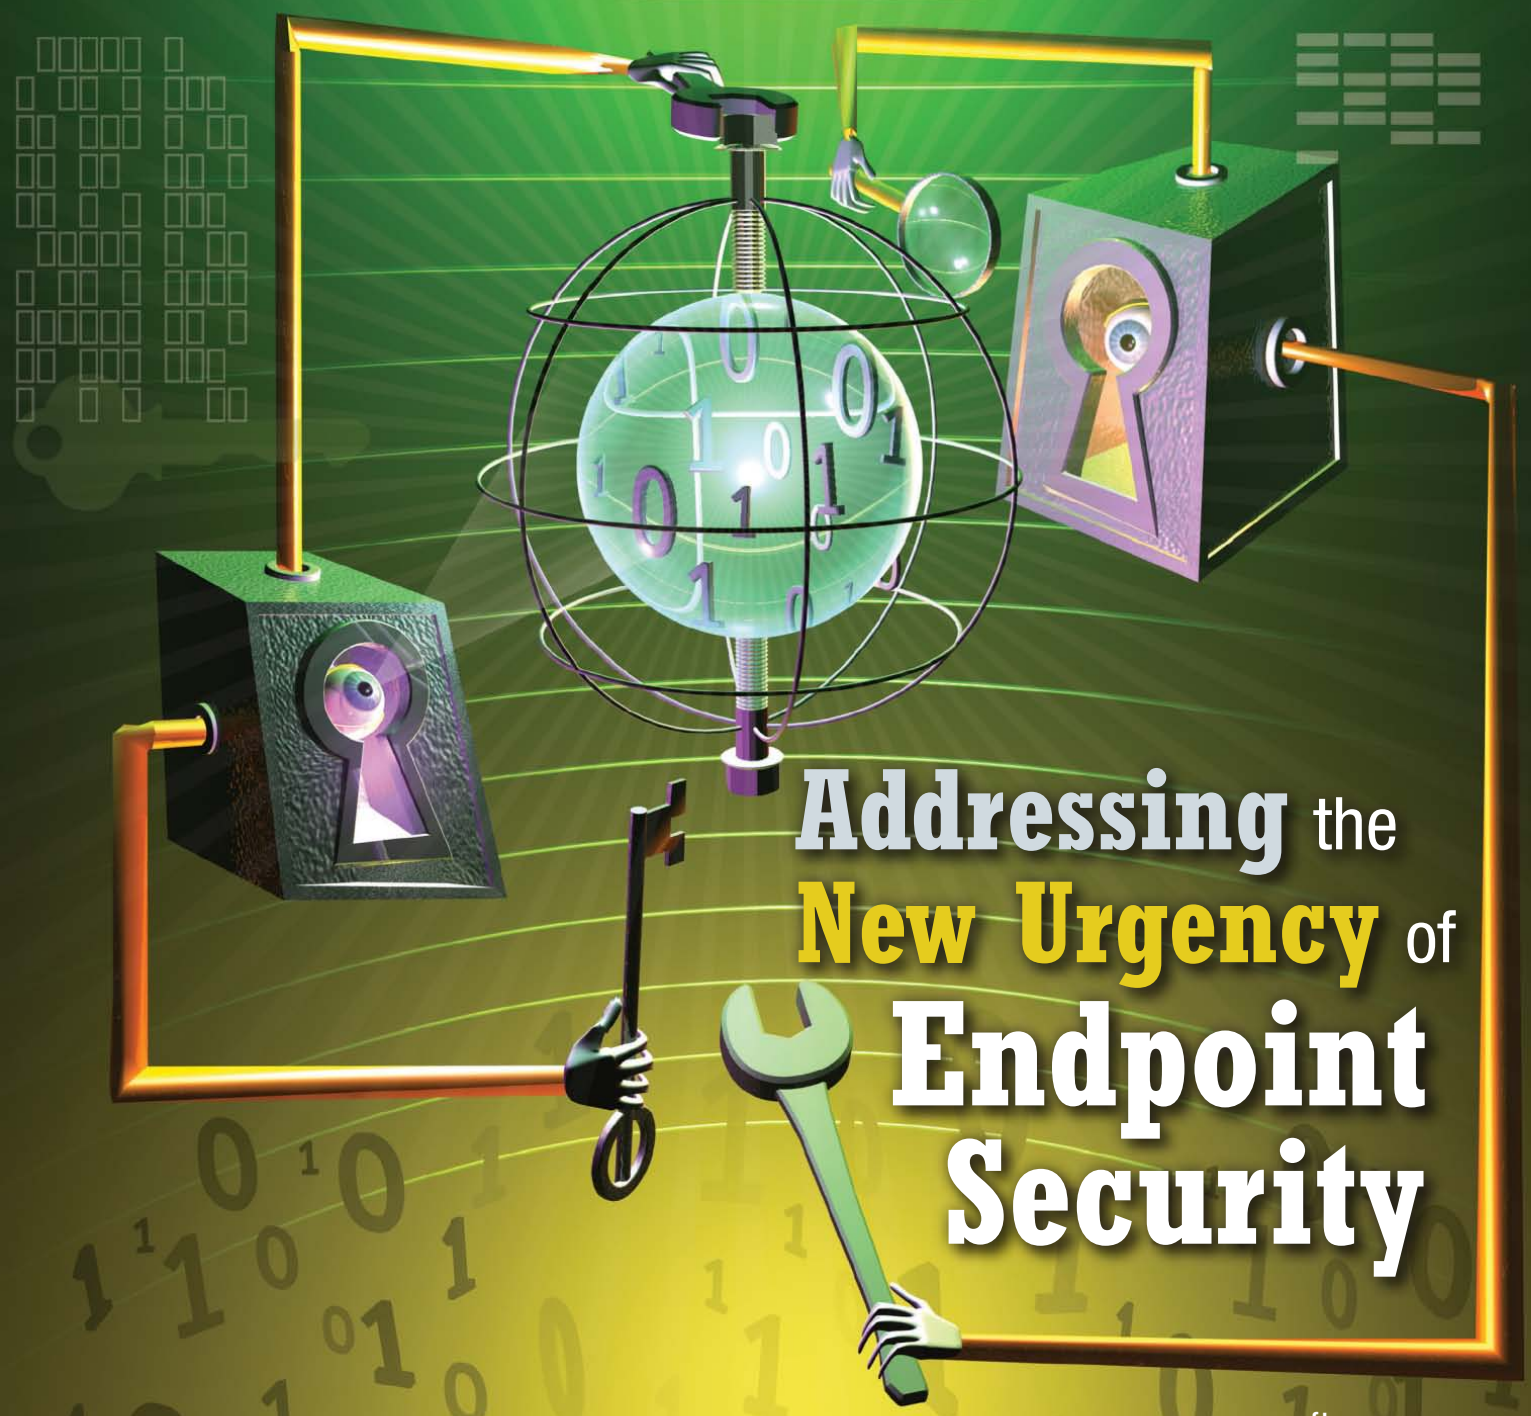

Addressing the New Urgency of Endpoint Security

A KING CONTENT PUBLICATION

www.softwaremag.com

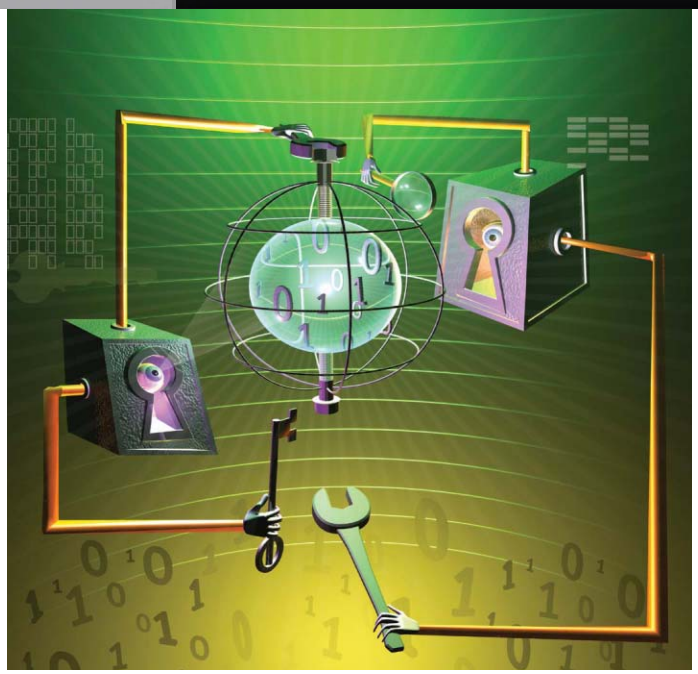

COVER AND ILLUSTRATIONS BY CHUCK MACKEY

COVER STORY: Security

24 Addressing the New Urgency of Endpoint Security

Unless enterprise security architecture addresses endpoint security comprehensively (i.e., taking into consideration things like virtualization, mobility, and social networking), cyber-security will remain elusive

By Sreedhar Kajeepeta

>>COLUMNS

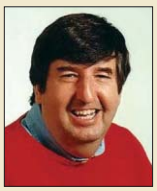

>>DEPARTMENTS

DIGITAL SOFTWARE MAGAZINE
The Software Decision Journal
VOLUME 30, ISSUE 4 • WINTER 2011 EDITION

>>FEATURES

Application Development

11 Estimating the Duration of Project Tasks

Insufficient time spent on schedule development is a key risk to project success; you can eliminate the target-date tango by building a schedule defense that manages the risks

By Kim Pries and Jon Quigley

Business Intelligence

18 Discovering the Big Deal About Big Data

Today, with a significant decrease in cost of storage and processing—along with a slew of new tools for capturing, storing, organizing, and analyzing data—big data has become big opportunity

By Umesh Jain

Digital Software Magazine is distributed quarterly in 2011 by Qmags, the most established platform in the digital magazine delivery business, now in its 13th year, to the following titles from our database of opt-in registrants:

CIO/CTO/VP of IS/MIS	3,000	Product Managers	3,050
VP/DIR/MGR/IS/IT/MIS	3,100	Engineers	7,000
VP/Dir/MGR of Devlpmnt, Netwrks	4,500	Sales	10,000
Software Specialist	5,100	Marketing	14,650
Other IT Management	5,600		
CEO/President/Chairman/Partner	25,000	**Total**	**81,000**

editor's letter

WINTER 2011 EDITION

Urgency of Endpoint Security and the Big-Data Talent Gap

Two pieces in this issue provide timely insights into some major computing trends, that of "big data"—datasets beyond the ability of typical database software tools to store, capture, and analyze—and the security of endpoints, any device with an IP address and a port that is formally attached to the corporate network.

Writer Sreedhar Kajeepeta, global VP and CTO of technology consulting for CSC's global business solutions unit, provides a wakeup call in this issue for the security of endpoints. Although endpoints primarily refer to servers, desktops, laptops, smart phones, and embedded systems, today a newer set of endpoint concerns has emerged, related to virtualization, mobility, and social networking.

In social networking, for instance, where server endpoints may host customer data, as opposed to just marketing collateral, on public clouds managed by sites such as Facebook or Twitter, companies need to be concerned about data security and regulatory compliance issues related to the vertical industries involved, Kajeepeta notes.

Gartner Group has identified a market for Endpoint Protection Platforms and issued a Magic Quadrant report on the segment. Despite improvements and a variety of security products, malware is on the rise in general.

So now the corporate security analysis challenges corporate IT professionals and system integrators to conduct a broader plan. The author suggests two recommended frameworks.

Writer Umesh Jain, founder and president of Merging Elements, an advisory firm focused on customer management and IT strategy, clarifies the picture of big data as having three critical elements important to recognize. The first is the volume of data, being driven by the growth of social networking sites such as Facebook, and the growth in mobile phones and associated data. Today, big data is terabytes and even petabytes of information.

The second is the variety of data sources generating data that results in big data today, including Web sources such as clickstreams, RFID data from supply-chain applications, unstructured text data from contact centers, geospatial data, and multimedia data. The third crucial driver is the velocity, or speed, at which this data is being generated, making it difficult to analyze the data and take action.

Jain outlines good reasons this data needs to be analyzed, and he identifies a "talent gap" between the managerial skill sets needed to analyze big data and the available supply. He cites a McKinsey Global Institute study saying the U.S. faces a shortage of 140,000 to 190,000 managers with deep analytical skills, and 1.5 million managers skilled in analysis to be able to make decisions on the findings. This is a challenge and an opportunity.

We hope these insights will help IT professionals in our audience to do a better job. Please let me know what you think of the two pieces and what other topics would interest you.

Regards,

John P. Desmond
Editor, Software Magazine
jdesmond@softwaremag.com

DIGITAL SOFTWARE MAGAZINE
The Software Decision Journal

www.softwaremag.com

Winter 2011
Software Magazine Vol.30, Issue 4

EDITOR
John P. Desmond
jdesmond@softwaremag.com

CONTRIBUTING WRITERS
John Humphrey, Umesh Jain, Jim Johnson, Sreedhar Kajeepeta, Kim Pries, Jon Quigley, and Steven Rogers

MANAGING EDITOR
Lana Gates
lgates@softwaremag.com

ASSISTANT MANAGING EDITOR
Sue Bencuya
sbencuya@comcast.net

ART DIRECTOR
Martha Abdella
mabdella@softwaremag.com

ILLUSTRATOR/ARTIST
Chuck Mackey
chuck@chuckmackey.com

ONLINE PRODUCTION EDITOR
Martha Spizziri
mspizziri@softwaremag.com

SOFTWARE 500 PROJECT CONTACT
Kim Evans
kevans@softwaremag.com

EDITORIAL ADVISORY BOARD
Dr. V. P. Kochikar
Associate VP, Principal Consultant
Infosys Technologies

Neil Florio
Director of Marketing
QlikView

Software Magazine is a property of King Content Co.

KING CONTENT CO.

P.O. Box 135
E. Walpole, MA 02032

Phone: (508) 668-9928
Fax: (508) 668-0442

www.king-content.com

PRESIDENT
John P. Desmond

For print subscriptions, go to:
www.softwaremag.com/register

The Software Magazine digital version is delivered in QuVu (patent pending) format, which is licensed from and used by permission of Qiosk.com.

From: Andrew Blencowe [mailto:CEO@Argent.com]

To: 'All_VMware_Users'

Subject: **Get Total Control Of Your VMware In Under Four Hours**

Hi,

Spoke to your boss

He's pleased with your progress with VMware

He is concerned there's no automation or monitoring

We discussed installing Argent for VMware

I mentioned us winning Best of Show at TechEd beating Microsoft

Pls download from **www.Argent.com/VM**

Regards,

ARB

ARGENT

By John Humphrey and Steven Rogers

On the Edge of a Breakthrough

By and large, companies have been forced to react to recent economic pressures by slashing budgets, reducing staff, and postponing needed IT infrastructure upgrades. Frequently, these decisions were based purely on the need to cut costs, with little consideration of how such actions would impact business and revenue growth goals.

During times of economic uncertainty, needed IT infrastructure upgrades and transformations are regularly postponed. In spite of this reprioritization, however, businesses frequently fail to determine how to proactively direct resources toward maximizing the value of their existing IT systems by choosing simpler, less expensive options designed to improve productivity and deliver a faster time to market and speed to value.

This can be accomplished by redirecting IT investments from the "core" infrastructure and applications, such as enterprise resource planning (ERP) solutions, to enterprise applications and solutions with technologies that deliver value at the "edge."

"Edge" IT applications include portals, e-commerce platforms, mobile applications, and distributed business intelligence (BI) platforms that enable real-time (or near real-time) interaction with employees, customers, suppliers, and other stakeholders. By design, edge IT applications improve the customer and employee experience, reduce costs, and maximize efficiencies at investment levels much lower than what would be required for core enterprise system customization.

Focusing on the edge directs investments toward solutions that create customer and shareholder value in a shorter timeframe; this strategy also prepares businesses to address future economic downturns. To determine how to maximize IT investments to ensure that consistent, incremental value is delivered to shareholders, stakeholders, and employees even in the most challenging economic circumstances, executives need to consider how *the industry value chain, economic cycles,* and the choices in *IT application architectures* impact their organizations.

The Industry Value Chain

Throughout history, economies have changed along with theories about value chains—the combination of participants (both suppliers and customers) that come together to create, purchase, and sell materials and finished goods to the customer. Once upon a time, businesses like Ford Motor Company contained an entire value chain within the organization—Ford owned the foundry that manufactured the steel, the factories that built the cars, and the dealerships that sold them.

When increased competition from international steel companies and foreign car manufacturers drove companies like Ford and General Motors to specialize in order to drive efficiencies and improve the quality of their vehicles, value chains were forced to evolve. The value chain within almost every industry was affected in a similar fashion, as component manufacturing became outsourced to specialized organizations.

In both good and bad economies, companies look for ways to reduce the length of time it takes for a product to reach the customer by collapsing the value chain. In the past, enterprises would spend billions of dollars on core ERP and electronic data interchange (EDI) systems, the soundest investment possible given the alternatives.

This is no longer the case—orga-

John Humphrey is a co-founder of Pariveda Solutions and currently serves as its chairman of the board. Steven Rogers is a vice president at Pariveda Solutions and has extensive experience with Web technologies and service-oriented architectures.

nizations now have tremendously rich enabling technologies that can be used to improve the experience and create additional value for stakeholders. Changing the IT investment approach to edge solutions supports rapid time to market and reduces the cost of development, while still delivering on business requirements.

The Economic Cycle

A common error that organizations make during an economic growth phase is to direct the majority of available funding for capital IT projects toward core IT infrastructure and ERP. However, because core IT systems require tremendous capital investment and time to implement, directing a majority of the budget toward these solutions takes the focus away from more nimble, lower-cost alternatives. Organizations that thrived during the downturn captured as much value as possible out of the systems they already owned.

Investing in capital projects that yield incremental value for stakeholders frequently provides organizations with a competitive advantage. Organizations benefit from a greater number of solutions that can be delivered during all parts of the economic cycle; they also develop a more cost-effective enterprise architecture that can continue to adapt during the downturns. In the process, dynamic IT cultures are created and resources directed to seek out opportunities to quickly deliver solutions to market and rapidly react to changing business demands.

The Core vs. the Edge

If the core is about transactions, the edge is where the biggest opportunity exists to improve the customer and employee experience, increase automation and accuracy, and collapse the latencies in the value chain. Edge applications move transactions closer to the "moment of value"—the expedited exchange of goods, services, and/or information among value-chain participants (employees, suppliers, and customers).

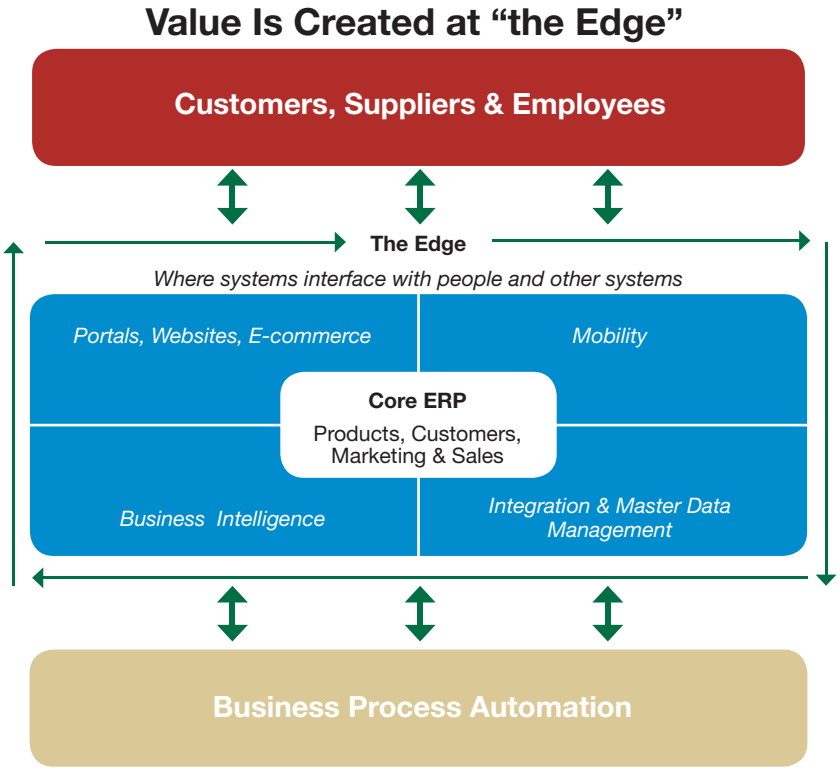

Figure 1 Source: Pariveda Solutions

Investing at the edge enables efficiencies that promote better decisions (BI), deploy solutions to employees, customers, and suppliers (e-commerce, portals), drive those efficiencies deeper into the value chain (mobility), and deploy data and integration through various Web services across the value chain. (See Fig. 1.) Edge applications are also more self-contained and allow organizations to rapidly deploy and/or adjust as business conditions change.

Increased investments in core applications result in intensive and expensive labor demands; projects can last for years and cost tens of millions of dollars. Conversely, edge applications can be delivered incrementally and create new value in a very short time period.

When economic tides turn, there is little time to react. The best-prepared organizations will have the ability to quickly respond and deliver new solutions to meet the needs of their customers, partners, and employees. As we enter the beginning stages of an economic recovery, companies should explore how to direct capital investments to the edge.

The delivered solutions will drive latency from the value chain and allow the kinds of efficiencies that will promote reductions in labor costs in the future. Maximizing the edge also addresses the moment of value and mitigates the risks associated with attempting to deliver all business needs with core IT applications. Speed, value, and flexibility—each are addressed through the edge. SW

The 2011 Software 500 CD

www.softwaremag.com/SW500CD

Software Magazine's **Research Team Delivers to you the 2011 Software 500 CD**

Inside you'll find:

- More than 5,000 executive names from the 2011 Software 500 companies (Excel format)
- More than 100 companies new to the list this year
- The 2011 Software 500 (Excel format)
- Data cuts from the 2011 Software 500 (Excel format)
- And the 2011 Software 500 editorial coverage (PDF format)
- Additional editorial content on mergers and acquisitions from 2009 to 2010, the ranking year

APPLICATION DEVELOPMENT

Estimating the Duration of Project Tasks

IS ESTIMATING THE DURATION OF A PROJECT ALWAYS A MATTER OF GUESSWORK? The answer to that should be no. Yet it frequently seems that it is guesswork, especially if we didn't learn much from previous projects. Although it often seems like every time is the first time, project estimation does not have to be a novel experience for the parties involved. Historical records can help to reduce the risk to the estimating process—if you have recorded and subsequently review this information.

Ultimately, estimation is and will always be an educated guess, although prior experience does go some way toward mitigating risk. (Even so, it's hard to understand how executives who have no idea of what it takes to achieve project results are able to tell you that you're budgeting for too much time.) Invariably, as a team marches through a project, they find that the accuracy (or lack thereof) of their estimates becomes clearer and clearer, in the same way that a GPS's arrival estimate becomes more accurate as we approach the target location. This article will guide you through how to scope a project and use some techniques to maintain control over the schedule estimates.

Steps in Estimating

Some information is essential to creating a meaningful estimate. For example:

- A statement of scope or a scope document that defines what the project is and is not
- A task list in the form of a work breakdown structure (WBS)
- Task details defined (i.e., not simply a list of task names)
- Duration estimations provided by the team
- Task dependencies (schedule and risks) clarified
- Schedule risks and task variations identified via Critical Path or other methodology
- Planned schedule risk mitigation

Insufficient time spent on schedule development is a key risk to project success; you can eliminate the target-date tango by building a schedule defense that manages the risks

By Kim Pries and Jon Quigley

Example of PERT Using an Excel Spreadsheet

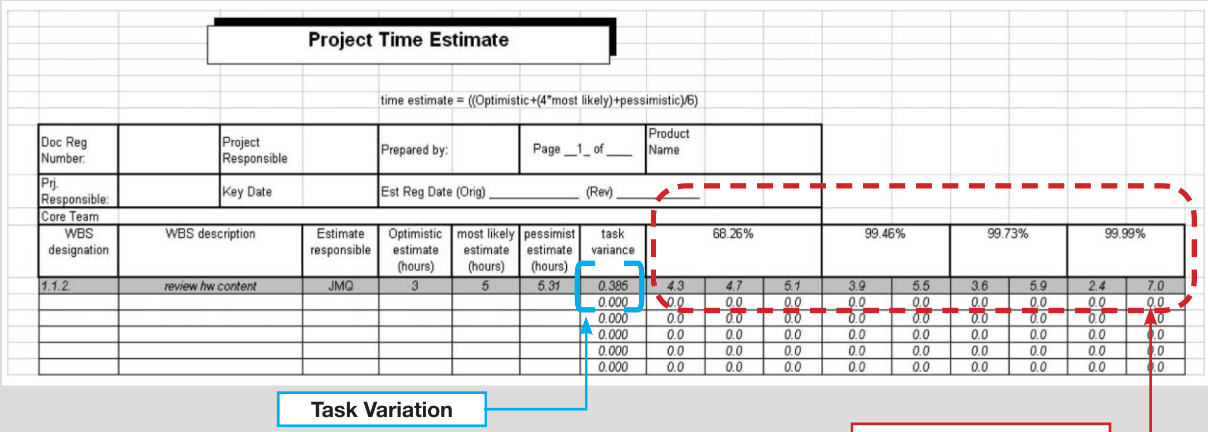

Project Time Estimate

time estimate = ((Optimistic+(4*most likely)+pessimistic)/6)

Doc Reg Number:		Project Responsible	Prepared by:		Page __1_ of ____	Product Name
Prj. Responsible:		Key Date	Est Reg Date (Orig) ____________ (Rev) ____________			
Core Team						

WBS designation	WBS description	Estimate responsible	Optimistic estimate (hours)	most likely estimate (hours)	pessimist estimate (hours)	task variance	68.26%			99.46%		99.73%		99.99%	
1.1.2.	*review hw content*	JMQ	3	5	*5.31*	*0.385*	*4.3*	*4.7*	*5.1*	*3.9*	*5.5*	*3.6*	*5.9*	*2.4*	*7.0*
						0.000	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*
						0.000	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*
						0.000	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*
						0.000	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*
						0.000	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*	*0.0*

Task Variation

The division by 6 suggests that PERT uses the assumption of a normal distribution (Six Sigma covers 99.98 percent of the possible variation), which may not be warranted by the data.

Impact of variation on range of possibilities

Figure 1

Source: Kim Pries and Jon Quigley

All of the above items should be in at least preliminary form before attempting an estimate of a project schedule.

Defining Project Scope

The project scope actually defines the constraints on the project. Without this definition (agreed upon by customers both internal and external) of project boundaries and expectations, the team has no way to re-estimate time durations or assess risks if the scope changes later. (Note that "internal customers" are essentially players in projects that involve deliveries from one organizational function to another. It is necessary for such players to understand and agree upon what constitutes "good," useable input from their and other functions.) The very heart of scope is the WBS, which can be formatted in any of several ways.

Essentially, the scope definition allows us to quantify project success by providing the boundaries within which the team will work.

WBSs are often hierarchically broken down as cost centers (subsystems that cost money and for which we have accounting), particularly if the organization is following the dicta of MIL-STD-881C, the U.S. Department of Defense (DOD) standard that defines WBS format and content. In the context of this standard, cost center or task names can originate from organizational processes, known and proven best practices, expert and experienced opinions, or major deliverables.

In addition to breaking down hierarchically, this activity breaks down deliverables into the steps it takes to produce those deliverables (sub-deliverables). This smaller task/deliverable is easier to estimate than the larger tasks. The scope is clearer and it fits into the recommended 40- to 80-hour work package, allowing more refined tracking of progress when it comes time to monitor and control the project.

Just as the overall project has a scope, each task objective will have a task scope, defined as that which constitutes successful achievement of that task. By what measures will you determine completion of that task? The DOD expects these individual measures to be defined in a WBS dictionary, which typically provides a textual definition of each WBS line item. This also helps to delineate conditions for success in each task.

Duration Estimation a la PERT

Program Evaluation and Review Technique (PERT) is an estimating method that originated in the 1950s during the USS Nautilus nuclear submarine project. The PERT technique attempts to provide a quick approach to assessing the schedule with a pseudo-normal distribution of expectations. PERT has a terminology and concepts all its own:

- Optimistic = O
- Most likely = ML
- Pessimistic = P
- Task variance
- Normal distribution
- Task duration as a continuum of possibilities (probability)
- PERT equation = [(O + 4 x ML + P) / 6].

The result of the PERT equation is a weighted average that attempts to represent the joining of the three varieties of estimate—i.e., optimistic, most likely, and pessimistic. An optimistic schedule occurs earlier than the most likely schedule, which occurs earlier than the pessimistic schedule or milestone.

It is possible to downplay the most likely estimate by reducing the multiplier for any given component—for example, software development duration is often "optimistic" among all team members. Additionally, the approach assumes that all components of the scheduling model are sensible estimates themselves. (See Fig. 1.)

Task Variance

The task variance, the root of which is the standard deviation (normal distribution), is the delta between pessimistic and optimistic durations in its roughest form (statisticians may

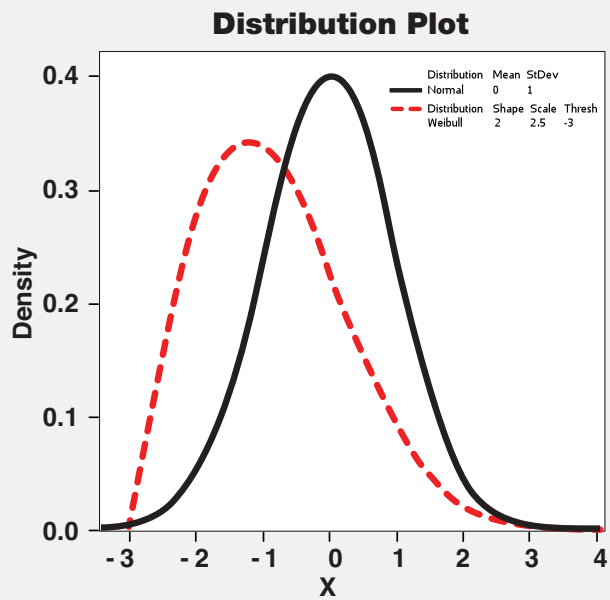

Figure 2 Source: Kim Pries and Jon Quigley

observe that this value is actually the range, a coarse measure of dispersion). The task variance, being derived from the previous duration estimates of "optimistic" and "pessimistic," provides an envelope of possible durations for the task to be completed. The larger the variance (or envelope), the higher the degree of uncertainty assumed by those doing the estimating.

It is incumbent upon a project manager to pay attention to this variation and the specific tasks affected—especially if these tasks turn out to be on the critical path. The critical path is the longest, contiguous, slackless path in the schedule and therefore provides the estimate of the shortest possible time to completion. The project cannot be delivered before the date this path indicates. Therefore, duration extension of any task on this path will result in a later delivery date.

Task Dependencies

Task dependencies are another part of this stew. Task dependencies are found in sequences of tasks in which one task *cannot* start until another task is complete. For example, a facetious set of dependencies would follow the sequence egg > chick > hen > fryer. If any task other than the project kick-off (first task) or the project closing meeting (last task) has no dependencies, then that task can presumably be executed immediately, if the resources to do so are available.

Large variation in the task estimate portends significant impacts on the schedule analogous to "tolerance stack-up" in the mechanical world (another form of variation). Each one of these variances is a schedule risk. If the project has been properly baselined in the project management software, the variance can be measured.

In cases in which the team experiences large variation on dependent tasks on the critical path, the project will see a "ripple effect" on schedule risk and an increase in overall project variance. If this ripple occurs on the critical path, the risk to the estimated schedule is immediate and recovery is potentially unattainable.

Network Diagrams and GANTT Charts

We use network diagrams to understand dependencies and schedule impacts. The network diagram is a representation of dependencies—in effect, a directed graph that considers all tasks to be nodes. Commonly, the nodes in the graph will show a variety of information, such as resources, start and finish dates, and budgetary information. The collection of task finish dates allow for a forecast of the probable completion date of the project.

Gantt charts are the best-known graphical representation of projects, but this approach has some significant limitations. These charts are, for example, not good at showing dependency impacts, and they don't offer the mathematical resources needed to graph theoretical calculations.

Estimation and Probability

Any time an upstream manager requests single, "hard" dates—which imply 100 percent likelihood—he or she is requesting an absurdity. A rational response would be to offer a span of dates. This could be done using PERT, with the estimated mean plus standard deviations indicating assumption of "normal" distribution. A Rayleigh distribution would also work here.

A Rayleigh distribution is a Weibull distribution with a shape factor whose value is the integer 2. (See Fig. 2.) The Rayleigh mean and variance is not the same as that for the normal distribution, but rather:

$$\mu = \gamma + \eta\Gamma\left(1+\frac{1}{2}\right) = \gamma + 0.88623\eta, \textit{where}$$

$$\gamma = \textit{position or threshold}$$

$$\eta = \textit{scale factor}$$

$$\sigma^2 = \eta^2\left[1-\left[\Gamma(1.5)\right]^2\right] = \eta^2(1-0.7854) = \eta^2 0.2146$$

where the position can be negative. The scale is the point at which we have 63.2 percent completion of the specific task.

Difficulties with probability are multitudinous, and include, for example:

- Lack of project history
- Failure to baseline previous projects
- Failure to scrutinize previous projects
- Tendency to be either too optimistic or too pessimistic
- Assumption of distribution can be totally wrong (especially for tasks with no history)
- Incorrect dependencies (joint probabilities)
- Existence of a critical path
- Organization's demand for "a" specific date (See Fig. 3.)

Each of these issues can be overcome, and some of them may not even be that important, if team estimates are solid. But in some cases, errors are unavoidable. Some of the errors we have seen are:

- Insufficient up-front time generating estimates (point source)
- Underestimates of test time
- Estimates provided by personnel with no experience or responsibility for the task
- Underestimates of the impact of lateness on the chain of dependencies
- Overestimates of the benefit of "crashing" the schedule—which usually means overtime and the possible addition of resources, especially people
- Underestimates of the human cost of overtime.

Expanding the Interval

Expanding the confidence interval is another action that—up to a point—increases the probability of a meaningful estimate. As we expand the interval in which our estimate falls, we increase the confidence in the result—in other words, 80 percent confidence has a narrower interval than 90 percent confidence. The issue can become meaningless—for example: Beginning of universe > end of time = 100 percent confidence.

Project estimates have a cone or triangle of uncertainty. (See Fig. 4.) The corollary to that cone of uncertainty is the impact on the target date estimate. An additional factor is measurement uncertainty, which is equal to the estimate variance. The following equation represents one model for estimate uncertainty:

$$\sigma^2_{estimate} = \sigma^2_{assumptions} + \sigma^2_{scope} + \sigma^2_{environment} + \sigma^2_{procedure} + \sigma^2_{task} + \sigma^2_{analysis} + \xi$$

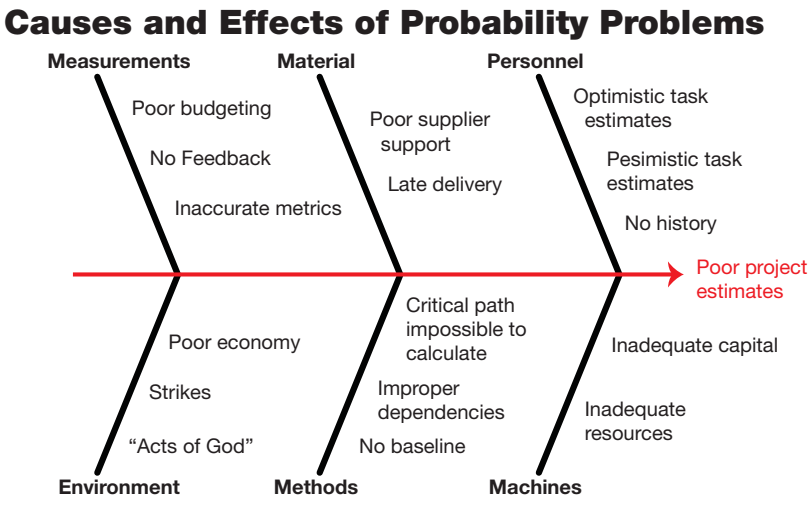

Figure 3 Source: Kim Pries and Jon Quigley

Managing Slack, Risk, and Deliverables

Managing the slack, or ostensibly idle time, may be perhaps the single most important factor in project success. When there is no slack, the team may move into the "death march" phase, which is inevitably followed by project doom. Instead, manage the slack through control mechanisms (feedback) and monitoring. This will help the team to identify key task (critical path) metrics and then track them to predict task conclusion, at the same time sounding the alarm when slack time vanishes.

Without some kind of risk mitigation, estimates are likely to fail. We suggest that managers perform a project failure mode and effects analysis, activate contingency plans to keep on track, and above all, pay attention to the details.

The project manager might consider managing deliverables rather than directly managing the tasks. After all, what customers receive is a deliverable they can see or hold or use. Delivering a product is only part of the story; documentation is also a deliverable, and so is support work. We introduce errors when we don't account for these items; hence it is best to develop the schedule cross-functionally.

Re-estimation

If all else fails, we can re-estimate the course of the project. Re-estimation should be routine for any change in scope (schedule, budget, feature set/quality) and "noise" (floods, power outages, hurricanes; strikes; losing key players).

Note that few projects scale linearly; that is, increasing scope increases complexity. We believe re-estimation should be linear only if the scope increase is very small. Without history from other projects, nonlinear adjustments are difficult.

What About Reviews?

Reviews are a primary feedback mechanism. For reviews to really work, they must occur frequently (no more than 30 days apart). The client should be updated after each review, and wherever possible, surprises should be eliminated (that way, ideally, the project will go no more than 30 days out of whack).

We can also activate a set of prophylactic schedule responses:

- Alter task sequence/dependencies where possible
- Control the method of achieving specific tasks (i.e., wherever possible, eliminate risks that cause high variation)
- Account for task variation in the project delivery schedule
- Use a capacity resource planning approach (i.e., critical chain)—an iterative approach to resource planning that always provides a good, but may never converge to an ideal, solution.

Do *not* build schedule crashes into the estimates. Crashing the schedule attempts to decrease the sched-

ule to deliver the project by "throwing" human resources at the project tasks. It is analogous to the idea that nine women can produce a baby from scratch in one month. In order for a crash to work, the team will need to maintain some slack and have a convenient management reserve (more capacity in terms of dollars and human resource availability).

Once a project is completely on the critical path, the target date is unachievable. Crashing may sound great to an upstream manager, but crashing often means the project already has exhausted its contingencies and the schedule itself is now out of control.

Conclusion

Can the target *ever* be met? Yes—look at the Empire State Building. It was completed 1.5 months ahead of schedule at less than 90 percent of the budget, with no project software or electronic spreadsheets!

We say that the best schedule is one that is accurate enough that it doesn't have to be changed. Barring that, contingency plans are the order of the day, and they must be exhaustive—in fact, it is not unreasonable to have multiple or layered contingency plans. Additionally, constant monitoring of project deliverables and schedule dependencies lets the project manager know the state of his or her project and when the contingency plans should be evoked (or the stakeholders notified).

We suggest that project managers express targets in probabilistic terms and that neither they nor their managers expect the "old school try" will save a fiasco. In short, comprehensive preparation is the surest way to moderate the effects of accidents, stupid decisions, and irrational expectations.

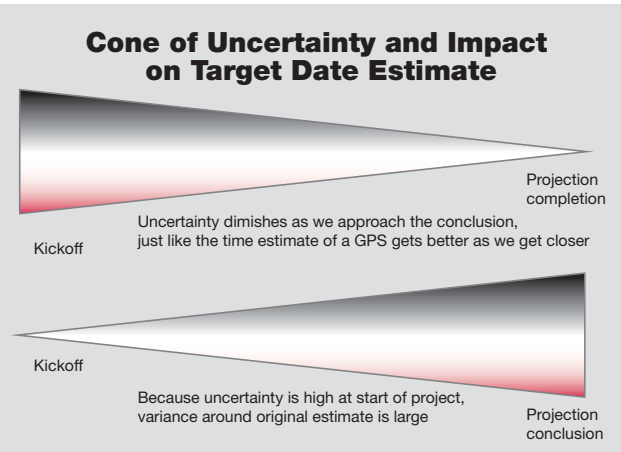

Figure 4 Source: Kim Pries and Jon Quigley

Buyers Guide: Project Cost Management Software & Services

Company	Focus Areas	Products	Services
A-1 Enterprise	ACLC, BC,COL, ONL, PRP, PRT, TET, WBS	x	x
Accord Software & Systems	COL, PRP, PRT, RSC, WBS	x	
CA	BC, COL, PRP, PRT, RSC, TET	x	x
Celoxis	COL, ONL, PRP, PRT, RSC, TET	x	
Charismatek Software Metrics	EST, PRP, RSC	x	x
Cost Xpert	EST, PRP	x	x
EcoSys	BC, PRP, RSC	x	x
Eigasoft	BC, WBS	x	
Galorath	EST, PRP, WBS	x	x
IBM Rational	ACLC, COL, PRP, PRT, RSC	x	x
Metrics Software	COL, EST, PRP, PRT, RSC, TET, WBS	x	
Microsoft	BC, COL, PRP, PRT, RSC, TET	x	x
Oracle	ACLC, BC, COL, PRP, PRT, RSC, TET, WBS	x	x
PRICE Systems	EST	x	x
PrimaSoft	BC, PRP, PRT, TET	x	
QSM	ACLC, EST, PRP, PRT	x	x
RASS Tools	EST, PRP, RSC	x	
Scoutwest	PRP, RSC, TET	x	
Software Productivity Research	EST, WBS		x
Tassc Ltd.	EST, PRP, PRT, RSC, TET	x	x
Tecolate Research	ACLC, EST, PRP, PRT, TET	x	x
Tenrox	PRP, RSC, TET	x	x
VIP Quality Software	COL, PRP, RSC, TET	x	

Key to Focus Areas

ACLC	Automatic Calculations	ONL	Online Software	TET	Time/Expense Tracking
BC	Budget Control	PRP	Project Planning	WBS	Work Breakdown Structure
COL	Collaboration	PRT	Project Tracking		
EST	Estimation	RSC	Resource Management		**Bold = Sponsor**

Kim Pries, APICS CPIM, and Jon Quigley, PMP CTFL, are principals with Value Transformation, LLC, a product development training and cost improvement or ganization. They have written Project Management of Complex and Embedded Systems *as well as* Scrum Project Management *and* Testing Complex and Embedded Systems, *all available at Amazon. They are working on two more books. Contact them at kim.pries@valuetransform.com and jon.quigley@valuetransform.com.* SW

BUSINESS INTELLIGENCE

Discovering the Big Deal About Big Data

THE DEFINITION OF "BIG DATA" VARIES DEPENDING ON WHOM YOU ASK. At a very simplistic level, "big data" refers to datasets whose size is beyond the ability of typical database software tools to capture, store, manage, and analyze.1 This definition, though, is rather subjective, focusing primarily on the volume of data. Volume, of course, is important, but consider that having a lot of data about just one thing may not be that big of a deal, from a value-creation viewpoint. Moreover, what's "big" today will not be "big" tomorrow—making it tough for people to grasp what we really mean when we use the term.

The definition of "big data" varies depending on whom you ask. At a very simplistic level, "big data" refers to datasets whose size is beyond the ability of typical database software tools to capture, store, manage, and analyze.[1] This definition, though, is rather subjective, focusing primarily on the volume of data. Volume, of course, is important, but consider that having a lot of data about just one thing may not be that big of a deal, from a value-creation viewpoint. Moreover, what's "big" today will not be "big" tomorrow—making it tough for people to grasp what we really mean when we use the term.

I personally prefer an alternative definition of big data that combines three attributes: volume, variety and velocity—commonly referred to as the "three Vs."[2] (See Fig. 1.)

Today, with a significant decrease in cost of storage and processing—along with a slew of new tools for capturing, storing, organizing, and analyzing data—big data has become big opportunity

By Umesh Jain

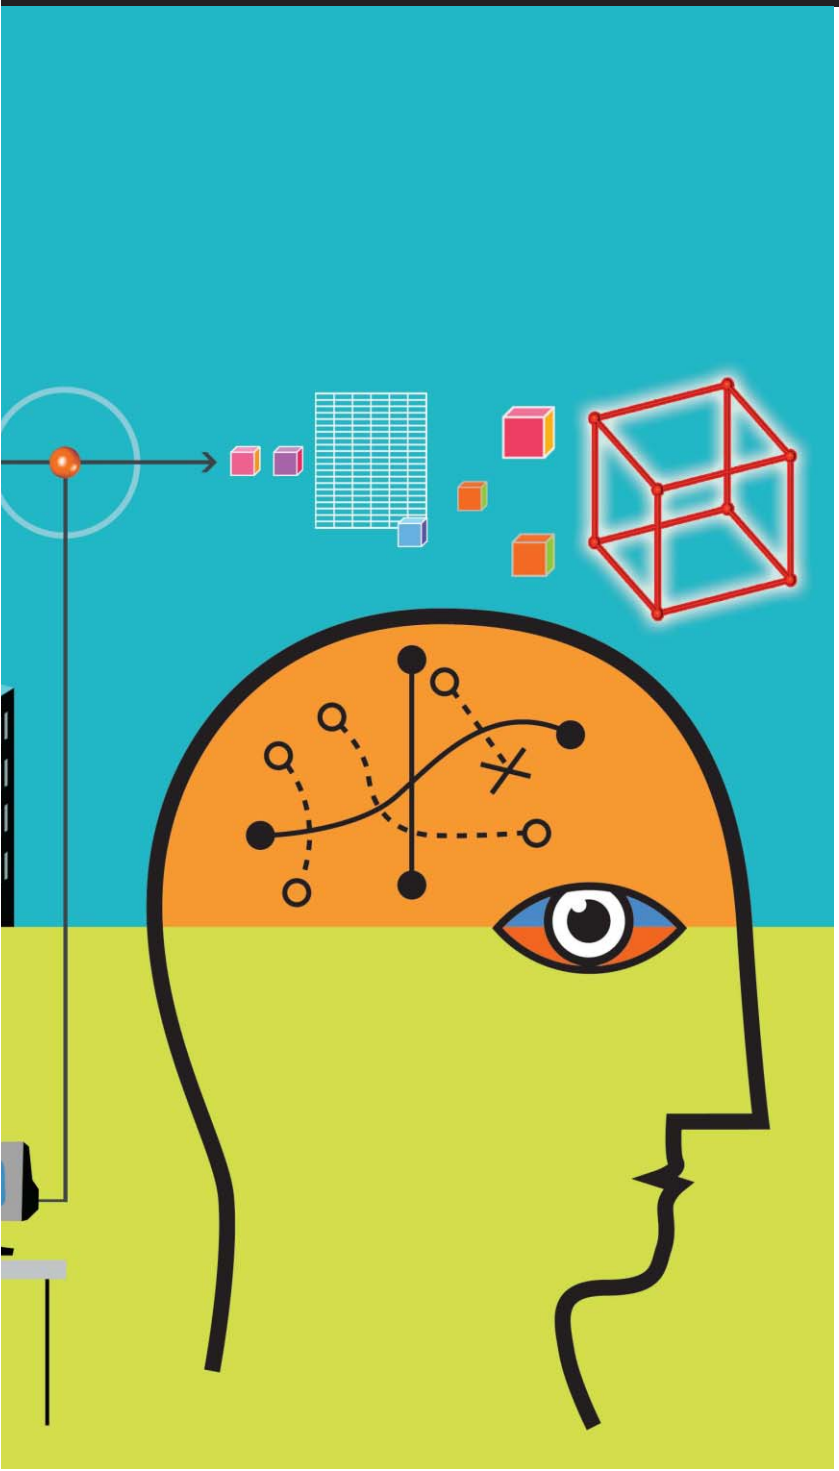

Volume

In 2011, the amount of information created and replicated will surpass 1.8 zettabytes (1.8 trillion gigabytes)—growing by a factor of 9 in just five years.[3]

Thirty billion pieces of content (photos, Weblinks, notes, and so on) is shared on Facebook every month.

As of September 2011, the Library of Congress had collected about 254 terabytes of data (1 terabyte = 1,024 gigabytes). The archives grow at a rate of about 5 terabytes a month.

In the 20 years from 1990 to 2010, worldwide mobile phone subscriptions grew from 12.4 million to more than 4.6 billion, penetrating the developing economies and reaching the bottom of the economic pyramid.

It is clear that data volume is the primary attribute of big data. With the growth of volume in social networking sites such as Facebook and the growth in mobile phone and associated data usage, managing the tremendous growth in data volumes has to be a critical part of any big-data strategy. Just a few years back, big data would have been defined as gigabytes of information. Today, we have moved that to terabytes or even petabytes of information to define big data. Currently, most users define big data as anything above a few terabytes.

Variety

The second attribute of big data is the variety of data sources involved. Historically, companies had just a few operational and transactional systems from which data was aggregated. Today, we have many more sources and formats, from clickstreams and social media to RFID from supply-chain applications. We get geospatial and multimedia data, unstructured text data from contact centers, and context-aware data from mobile sources. The net effect of this wide variety is that it has made data aggregation and the analysis of unstructured data more complex, while also fueling the growth in data volume.

Velocity

The speed at which all this data is being generated—its velocity—is yet another attribute that helps in defining what we mean by "big data." Another way to think of velocity is in terms of the frequency of data generation and delivery. The collection of real-time and near real-time data is nothing new: Companies such as Amazon

The Three Vs of Big Data

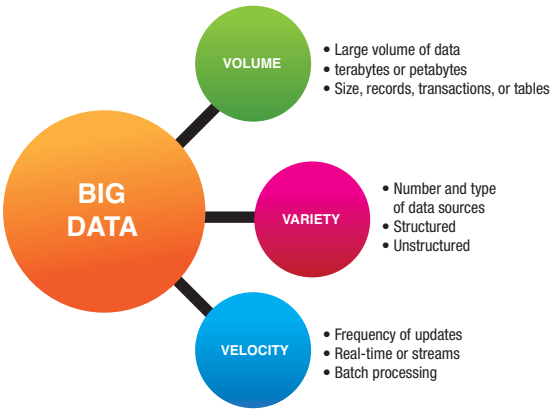

Figure 1 Source: Umesh Jain

and others have been collecting clickstream data for years to make purchase recommendations to Web visitors. However, the velocity at which it is generated makes this data difficult to analyze—and possibly respond to—in real time.

With the growth of volume in social networking sites such as Facebook and the growth in mobile phone and associated data usage, managing the tremendous growth in data volumes has to be a critical part of any big-data strategy.

The Importance of Big Data

Organizations that use data-driven decision making exhibit a 5 percent to 6 percent higher output and productivity than would be expected, given their other investments and usage of information technology.[4]

According to a recent estimate by McKinsey Global Institute, if U.S. healthcare could use big data creatively and effectively to drive efficiency and quality, more than $300 billion could be saved every year, and national healthcare expenditures could be reduced by about 8 percent.[5]

With the amount of research and analysis seen lately backing big-data analysis, it is not surprising that this area of business intelligence is getting so much attention.

Analyzing big data can improve organizational performance. Using data to understand organizational performance trends, drivers, and variability in performance makes it possible to quickly identify root causes of performance issues and thus drive higher performance. An example of this is in contact centers, where accessing and analyzing a bigger dataset at the individual call and agent level enables organizations to reduce call volumes, improve agent performance, and drive better productivity, while improving customer experience.

Big data can customize the brand, product, and service experience. As Amazon does with its recommendation engine, organizations can use big-data analysis to segment and micro-segment their customers to improve real-time personalization. This personalization allows organizations to improve existing products and services, identify new opportunities, and invent entirely new business models, based on the ongoing analysis of data around changing customer needs and wants.

Large statistical samples provide better results. In general, the larger the sample size, the more accurate the results of statistical analysis. Hence, the use of big data significantly increases the confidence in statistical analysis used to make data-driven decisions.

Technology costs and capabilities have reached an acceptable level. Most of the analysis being done on big data is not new; it's just more accessible now to a broader set of companies due to significant improvements in the cost, capabilities, and simplicity of technologies that support this effort. The number of vendor choices and technology options for big-data analysis has grown astronomically

Buyers Guide: Big-Data Management and Analysis Tools & Services

Company	Focus Areas	Products	Services
1010data	CLB, PA	x	x
Accenture	IMD, PA		x
Actian	PA, RTD	x	x
Apache	DP, HDP, MR	x	
Cloudera	HDP, RTD	x	x
Greenplum, a division of EMC	AGA, HDP, MPP, SUD	x	
IBM	AGA, HDP, SUD	x	x
Impetus Technologies	HDP, MPP, NSQL		x
Infobright	COB, DC, HDP	x	x
Informatica	DP, HDP	x	x
Kognitio	CLB, IMD	x	x
Microsoft	HDP, MPP, MR, SUD	x	x
Oracle	HDP, MPP, NSQL	x	x
ParAccel	AGA, COB, HDP, MPP	x	
QlikTech	DC, IMD	x	x
SAND Technology	DC, IDA, PA, SUD	x	x
SAS	DVI, IDA, IMD, PA, SUD	x	
Sybase, an SAP company	AGA, COB, DC, IMD, PA	x	x
Tableau Software	DVI, IMD	x	
Teradata	MPP, MR	x	x
Vertica, an HP company	COB, DC, HDP, IDA, MPP	x	

Key to Focus Areas

AGA Agile Analytics
CLB Cloud-Based
COB Column-Based
DC Data Compression
DP Distributed Processing
DVI Data Visualization
HDP Hadoop
IDA In-Database Analytics
IMD In-Memory Database
MPP Massively Parallel Processing
MR MapReduce
PA Predictive Analytics
RTD Real-Time Dashboards
SUD Structured and Unstructured Data

Bold = Sponsor

over the last few years. Today, the market is full of vendors offering tools for advanced data visualization, predictive analytics, real-time dashboards, no-SQL DBMS, Hadoop, MapReduce, in-memory database, and visual discovery, among others. (See sidebar.) The rapid proliferation of these tools has made it easier for organizations to find the best tool to address their specific needs.

Big-data analysis offers transparency that was not available before. Making big data more easily accessible to relevant stakeholders in a timely manner can create tremendous value. For example, deep insights from raw contact-center data around what customers are calling about today can allow marketing, sales, operations, product development, and finance to fine-tune and improve the productivity and performance of the product, service, and overall business.

Challenges to Address

As with any major opportunity, big data involves challenges that need to be addressed. Organizations looking to harness the power of big data need to understand the issues and lessons learned by the early adopters in this space. Here are some key points to consider:

•*Talent gap.* There is a significant shortage of the analytical and managerial talent necessary to make the most of big data. According to the McKinsey study mentioned above, the United States alone faces a shortage of 140,000 to 190,000 people with deep analytical skills as well as 1.5 million managers and analysts to analyze big data and make decisions based on their findings. (See Fig. 2.) This is one of the biggest challenges facing organizations that want to move toward data-driven decision making. Organizations that recognize this challenge should consider investing in educating the current talent and look at the possibility of partnering with an outsourced analysis service provider to bridge this gap.

•*Technologies and techniques.* There are a number of technology decisions that are required to support big-data analysis. When selecting the technologies, first consider going with what you already know, as long as it fits the requirements; most vendors, including Microsoft, IBM, and Oracle, have solutions for big data. Beyond that, look for massively parallel processing (MPP) platforms, column-store databases, in-database processing techniques, in-memory solutions, predictive analytics, and advanced visualization technologies when looking for the best combination of technologies to meet your big data analysis requirements.

•*Data security and privacy policies.* Data security and privacy are critical considerations when designing your solution. With big data, you will be looking at data being extracted from multiple sources internal and external to your organization. Moreover, with data privacy concerns, particularly as they relate to consumer data, make sure that your policies and procedures are in place before embarking on big-data initiatives.

•*Access to raw data.* To enable greater insights and value from your initiative, you will very likely need to integrate data from external sources. Some of this data will need to be purchased from external entities, and gaining access to this third-party data is often not a straightforward process. Consider the options and requirements. Furthermore, this may be the first time you integrate some of your *internal* data sources into the overall data structure. These will come with their own complexities of integration and performance. Both internal and external sources will need to be considered in order to develop optimal data aggregation and access design.

•*Speed versus scale.* Balancing the speed of data access and process with the size and scalability of your storage and processing platforms is another vital area of consideration. Focus on starting with the questions that you are trying to answer through analysis, and extract only the data that's required to answer those questions. A lot of companies have more data than they need. Despite the fact that storage and processing costs are coming down, wasting resources and unnecessarily increasing the complexity of your solution is pointless.

•*Data versus value.* The process of capturing data has become so easy that organizations have gotten very good at collecting and storing it over the last 10 to 15 years. The problem that most of them are struggling with is figuring out how to extract value from all of this data. It's more important to spend some time figuring out what questions need to be answered to improve the business *before* turning to the data required for the answer.

Remember the three attributes of big data—volume, va-

Supply and Demand of Deep Analytical Talent by 2018

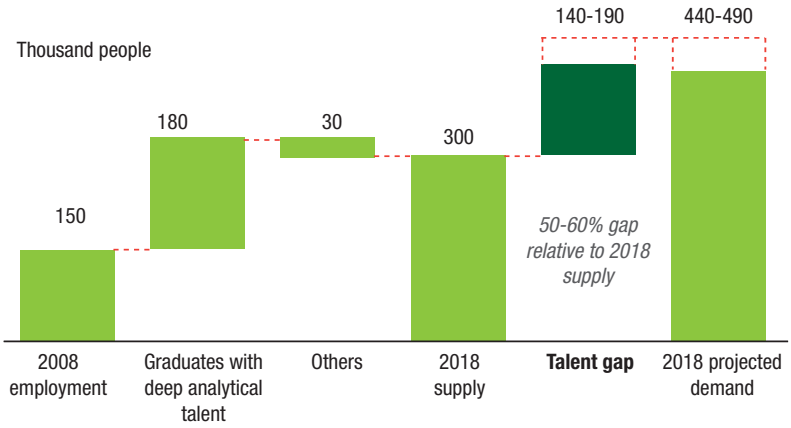

Figure 2 Source: McKinsey Global Institute, June 2011

Glossary: Common Terms Used in Big-Data Discussions	
Advanced Analytics	Refers to options such as predictive analytics, data mining, and statistical analysis—a combination of analytic tools and techniques to use on your big-data sets.
Complex Event Processing (CEP)	Used to process multiple events happening across all the layers of an organization, identifying the most meaningful events, analyzing their impact, and taking subsequent action.
Hadoop Distributed File System (HDFS)	Originally developed by Yahoo! as a clone of Google's MapReduce infrastructure (but subsequently open-sourced), Hadoop takes care of running your code across a cluster of machines. It chunks up the input data, sends it to each machine, runs your code on each chunk, checks that the code ran and passes on the results, sorts data between the map and reduce stages, sends each chunk of the sorted data to the right machine, and writes debugging information on each job's progress, among other things.
In-Database Analytics	Allows data processing to be conducted within the database by building analytic logic into the database itself. Doing so eliminates the time and effort required to transform data and move it back and forth between a database and a separate analytics application.
In-Memory Database	Can serve many purposes, but in business intelligence (BI), it usually supports real-time dashboards for operational BI and stores metrics, key performance indicators (KPIs), and sometimes OLAP cubes. The biggest benefit of an in-memory solution is the response rate and performance benefit that can be gained by eliminating disk input/output and other speed bumps.
MapReduce	A relatively new analytic option, MapReduce is an algorithm design pattern that originated in the functional programming world. It makes a distributed file system such as HDFS addressable through analytic logic.
Massively Parallel Processing (MPP)	Also known as a parallel database system, an MPP system runs on more than one machine, where each machine has its own disk storage. The database is physically located in several disk storage systems that are interconnected to each other.
NoSQL DBMS	(Sometimes expanded to "not only SQL") A broad class of database management systems that differs from the classic model of the relational database management system (RDBMS) in some significant ways. These data stores may not require fixed table schemas, usually avoid join operations, and typically scale horizontally.
Real-Time	A number of operational BI solutions utilize data in a real-time or near real-time basis. As more of these solutions move toward true analytics as compared to reporting and dashboards, a number of the big data analysis solutions will need to support real-time updates.
Unstructured Data	Data from social media, call centers, and other natural-language data sources that needs access to text mining, audio mining, and text analysis engines to be converted to structured data for analysis. The resulting data can be applied to a host of applications, including customer sentiment analysis, call reason analysis, and competitive intelligence applications. With big-data analysis, the need to harness the untapped value in unstructured data is very important.
Visualization	Data visualization is one of the fastest growing areas of BI. Advanced data visualization (ADV) is the perfect complement to big-data analysis, as it can handle visualizations to represent thousands or millions of data points—unlike standard pie, bar, or line charts.

Source: Umesh Jain

riety, and velocity. When deciding your big-data analysis solution, be sure you understand all of these attributes rather than focusing on just one of them.

As with many technology initiatives, it pays to start small and learn as you go, rather than take a big-bang approach. Remember that, given the improvements in technology and the constant decrease in storage and processing costs, what you consider big data today may be small data tomorrow. Don't overinvest in technology for technology's sake. Ensure that you clearly understand the value your organization needs from big-data analysis, because getting a high return on investment from your initiatives is critical to getting stakeholder and executive buy-in.

Umesh Jain is the founder and president of Merging Elements, an advisory firm focused on customer management and IT strategy services. He has more than 15 years of experience in contact centers, CRM, analytics, and technology in a variety of roles. He can be reached at umesh.jain@mergingelements.com.

[1] "Big data: The next frontier for innovation, competition, and productivity," McKinsey Global Institute, May 2011.

[2] In a 2001 research report, "3D Data Management: Controlling Data Volume, Velocity and Variety," and in related conference presentations, META Group (now Gartner) analyst Doug Laney defined data growth challenges (and opportunities) as being 3-D—increasing in volume (amount of data), velocity (speed of data in/out), and variety (range of data types, sources). Gartner continues to use this model for describing big data.

[3] IDC, "The 2011 Digital Universe Study," June 2011.

[4] Erik Brynjolfsson, Lorin M. Hitt, and Heekyung Hellen Kim, *"Strength in Numbers: How does data-driven decision-making affect firm performance?"* Social Science Electronic Publishing, April 2011.

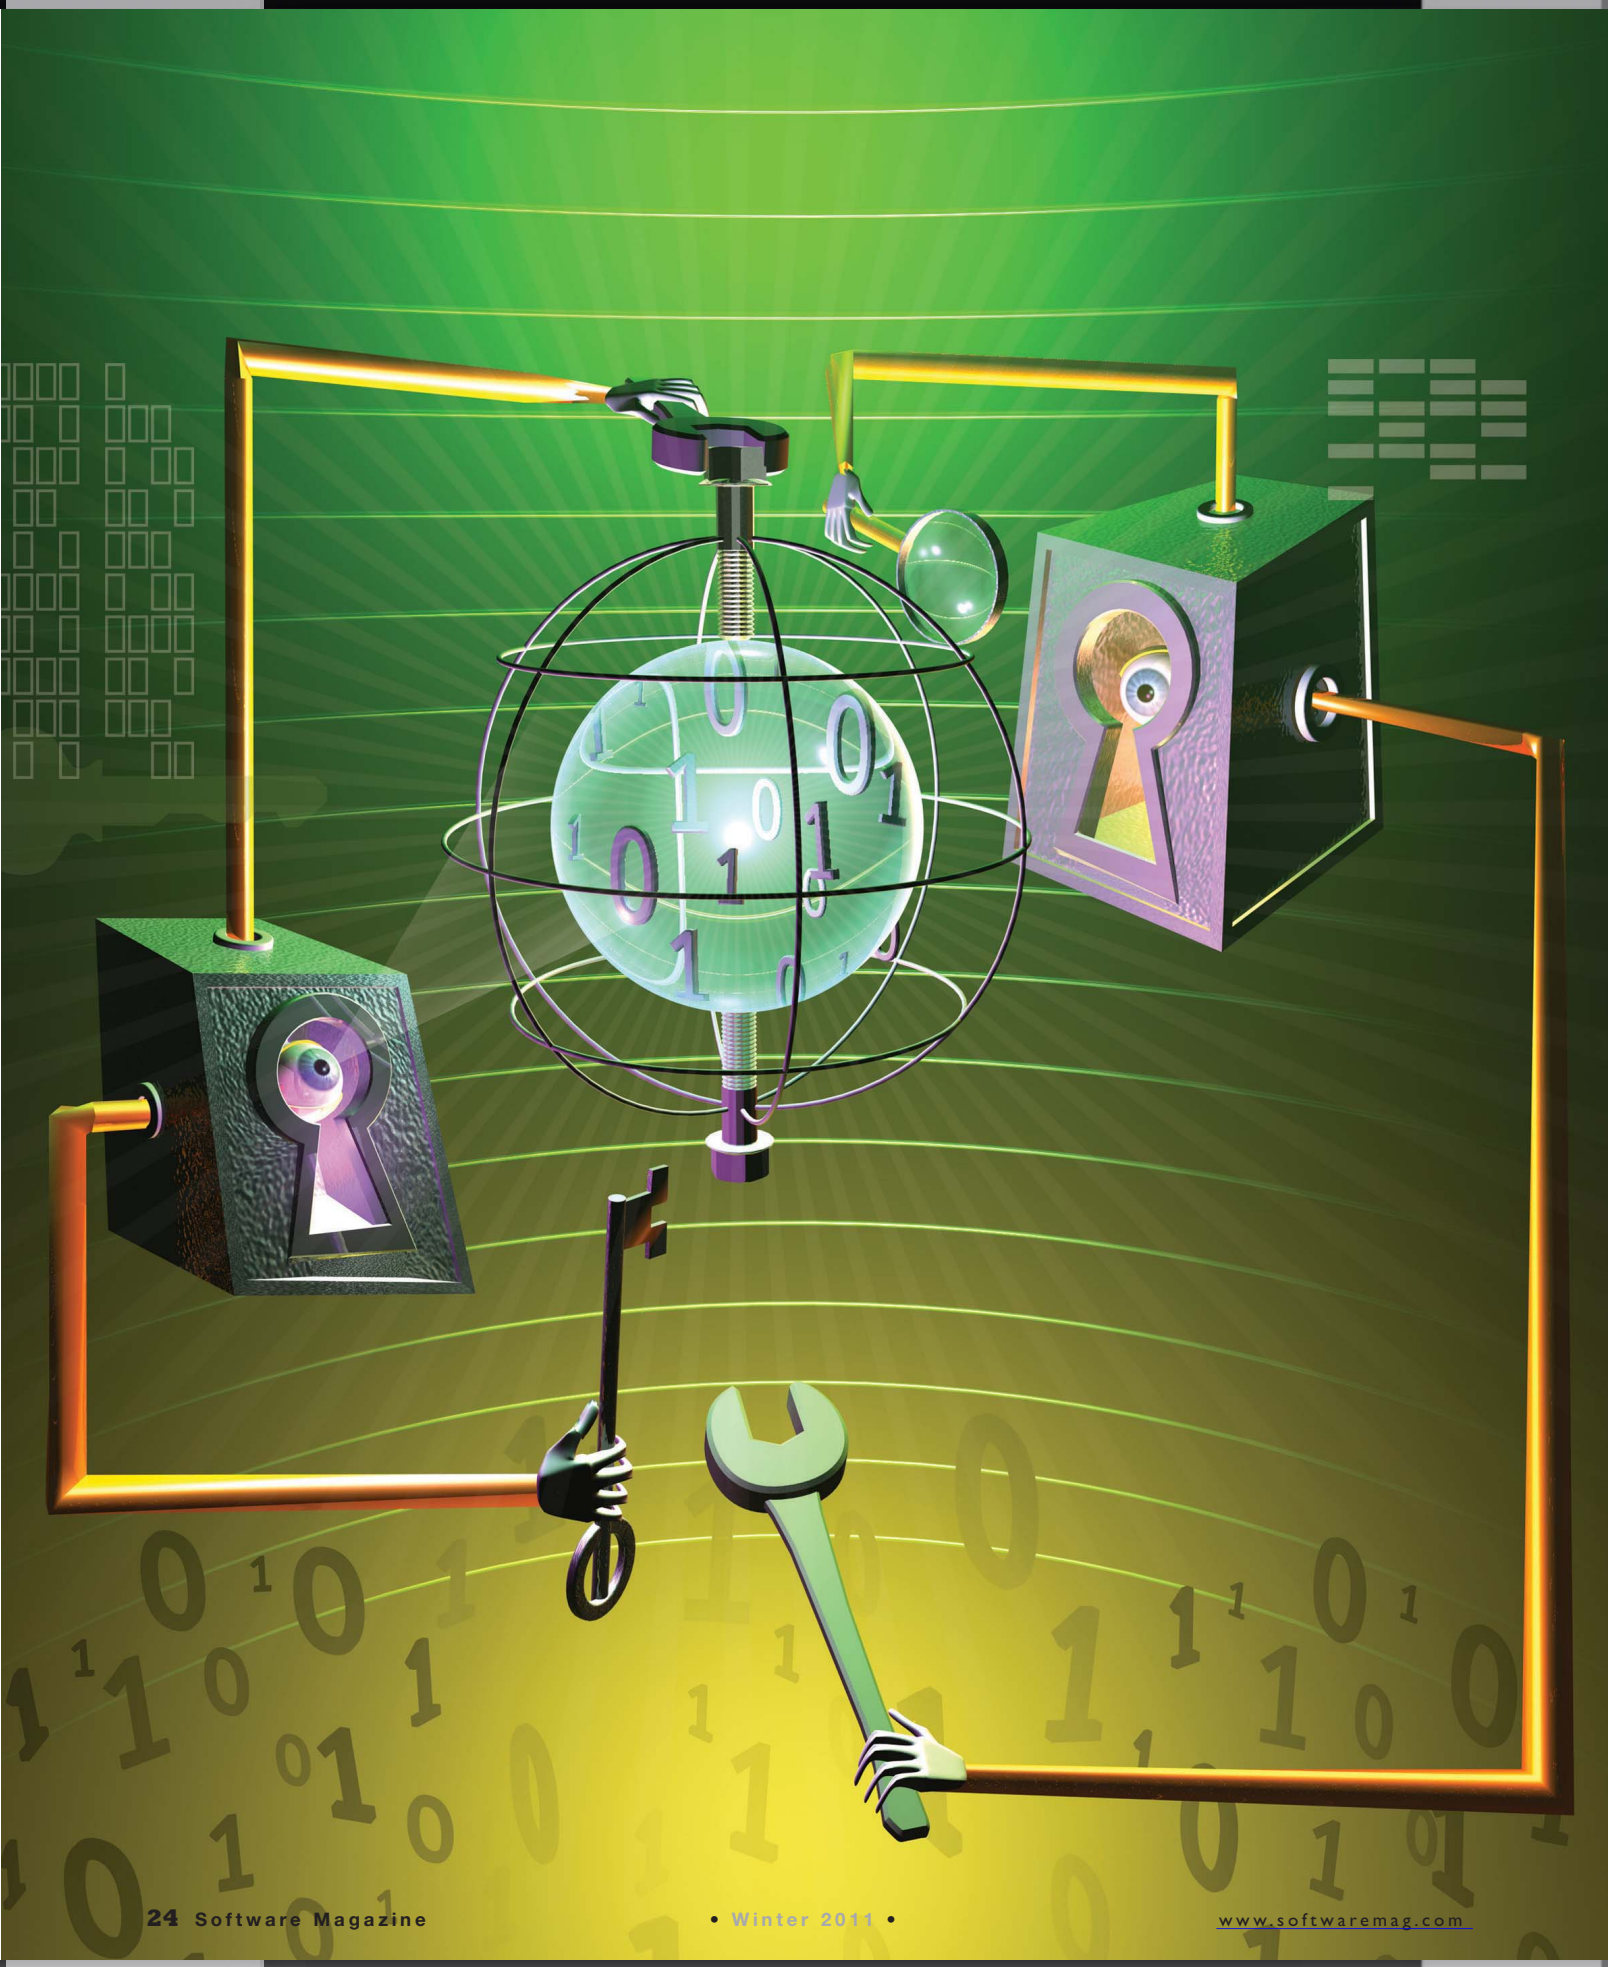

SOFTWARE MAGAZINE
Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page
Qmags THE WORLD'S NEWSSTAND®

SECURITY

Addressing the New Urgency of ENDPOINT SECURITY

Lately, thanks to improved technologies and compliance in perimeter security, many of us corporate denizens of the Net, those who work behind enterprise firewalls, have not been severely hurt or otherwise inconvenienced by the exploits of cybercriminals. Arguably, the last three to four years have been much quieter when it comes to email cyber-attacks.

Unfortunately, that's no indication that the digital world has suddenly become a safer place. Instead, the sad reality that confronts us is one in which the threats are increasingly more targeted and the criminals are becoming seasoned activists and professionals. As the criminals up the ante, their prime targets—the establishment and the enterprise—find new vulnerabilities exposed.

The challenge of safeguarding against this latest breed of advanced attacks is further compounded as IT (in public and commercial sectors alike) tries to embrace emerging trends that are inherently weaker against attacks. Such trends, which in fact are business imperatives, range from the infrastructural (as in wanting to be more virtual and more mobile) to the social (as in wanting to be more available on social networks).

Cybercriminals are fully exploiting this situation and are unleashing a new wave of attacks by targeting a different layer of enterprise security: the endpoints, which, as it turns out, were where cyber-attacks first started back in the 1980s, mostly through break-ins and sneaking in "Trojan horses" as part of software releases on mini-computers and mainframes.

Before we examine the nature of these new attacks and the way technology and related cyber-security solutions are rising to the occasion, let us first gain an understanding of the vulnerabilities we just talked about.

The Common Vulnerabilities and Exposures (CVE) system tracks publicly known information security vulnerabilities and exposures. This reference database is maintained by MITRE Corp. with the endorsement and support of the National Cyber Security Division of the United States Department of Homeland Security. Other federal agencies that recommend and/or require the use of cyber-security products that use CVE identifiers are the National Institute of Standards and Technology (NIST) and the Defense Information Systems Agency (DISA).

In 2010, a total of 4,651 vulnerabilities were identified in the CVE system. (See Fig. 1.) Endpoint operating systems (OSs)/applications such as Windows, Internet Explorer (IE), Java, and Adobe were all hit by these vulnerabilities.

Where Endpoint Security Fits In

The CVE represents just one quantitative analysis that we can use to un-

Unless enterprise security architecture addresses endpoint security comprehensively (i.e., taking into consideration things like virtualization, mobility, and social networking), cyber-security will remain elusive

By Sreedhar Kajeepeta

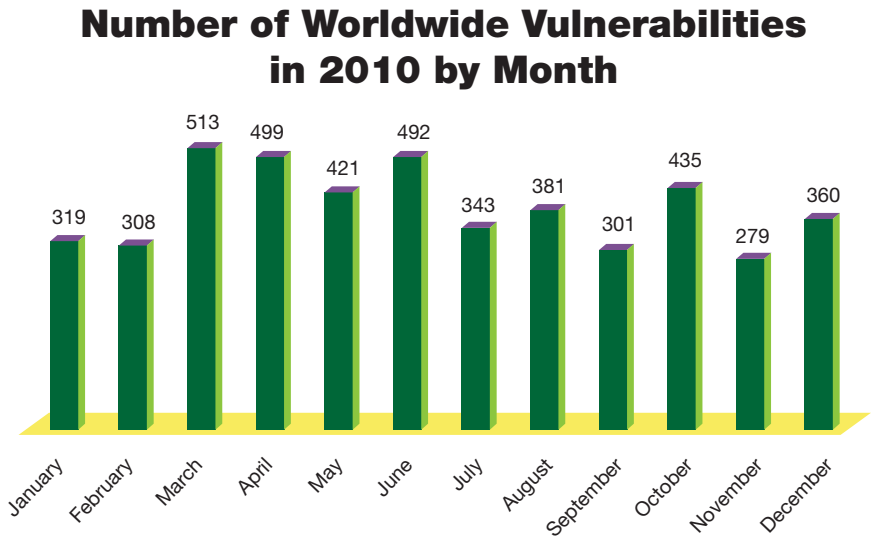

Figure 1 Source: Trend Micro on CVE database

derstand and appreciate the persistent nature of this stealth activity. To get back to our broader discussion, let us first restate what endpoints are and the position that endpoint security occupies in the broader map of enterprise security architecture.

In the strictest sense, endpoints are defined as network devices with an IP address and a port—in effect, any device that is formally attached to the corporate network. For the purposes of this discussion, however, endpoints are defined as servers, desktops, laptops, smart phones, embedded systems, and the like.

Technically, endpoint security is just one of a handful of security layers, and seemingly an independent one at that (in the sense of "separation of concerns"). However, it is also a layer in which trust can be quickly eroded, and one that is difficult to keep tabs on.

Take IBM's security framework, for example. The layers in that framework include (from top to bottom): People and Identity; Data and Information; Application and Process; Network, Server, and Endpoint; and Physical Infrastructure.

The People and Identity layer can have effective checks and balances built into it with identity and access control (and an accompanying audit log of who did what, when). Things can get very murky, however, when that layer comes in contact with the Network, Server, and Endpoint layer underneath. (See Fig. 2.)

When people use and manage endpoints, there can be a lot of room for contamination, whether willful or unintended—contamination that can never be tracked or accounted for. Surreptitious access to other enterprise assets, including systems and data through the many devices (laptops, desktops, smart phones, gaming consoles), servers (especially public-facing Web servers), and related (high-speed and high-capacity) multimedia ports available on the new and emerging endpoints, leave enough room for criminals to conduct their activities expeditiously, and without ever leaving fingerprints.

The April 2011 attack on Sony's gaming networks, which reportedly involved the theft of millions of records of consumer data, is a great example of how clear and present (and yes, very enormous) the danger still is at this layer.

These attacks were followed by attacks on Sony Online Entertainment and on Sony's Greek website. Indeed, all of the biggest security breaches of 2010 were attacks on endpoints. The details of these attacks are:

• *Aurora/Hydraq*, in January 2010, was targeted at high-tech companies (such as Google, Yahoo!, and Rackspace) and defense contractors (such as Northrop Grumman). It exploited an IE loophole to deliver malware capable of modifying applications.

• *Stuxnet*, in July 2010, was targeted at industrial software and programmable logic controllers (PLCs) of Siemens control systems. It was Windows-based, and its many variants attacked uranium enrichment infrastructure in Iran. In a related development, Iran reported that it had uncovered a new espionage virus, called Stars, that is aimed at damaging its government institutions.

• *WikiLeaks*, in October and November 2010, was targeted at U.S. defense and state departments. Exploiting unprotected (downloading without encryption) peripherals (as described above), WikiLeaks managed to steal 400,000 classified documents and more than 2,000 sensitive cables. The leaks continued into April 2011 with the publication of 779 documents related to the Guantanamo Bay prison camp.

• *LizaMoon*, in October 2010, was targeted at consumers and websites. Using SQL injection, it spread "scareware" and encouraged users to install rogue antivirus software. The attacks, according to McAfee, continued into April 2011, and affected more than a million sites.

A 2008 study conducted by the European Network and Information Security Agency (ENISA) found that the most common infection methods used in the preceding years were browser exploits (65 percent), email attachments (13 percent), OS exploits (11 percent), and downloads (9 percent).

That confirms that the attackers are indeed becoming adept at bypassing the perimeter to aim at the endpoints. Their efforts are getting a boost from malware or crimeware toolkits such as FakeAV and Zeus (which help in building botnets, which can control computers remotely; and zero-day attacks—for known (to hackers, but so new that they are unknown to most developers) vulnerabilities that are still being fixed by the software vendor.

A Rogue's Gallery of Endpoint Threats

Let us now examine the full range of endpoint vulnerabilities, and some of the attack types that can target them.

Attacks Related to Virtualization

Inefficient and sub-optimal as they were, physical servers nevertheless offered a level of built-in security that was inherent in their segregation and dedicated functionality. They had their own unique access, security controls, and administration. Unless similar segmentations are implemented using virtual local-area networks (VLANs) with appropriate role-based policies to restrict unauthorized access to a VLAN, virtual data centers and cloud infrastructures will be vulnerable to attacks.

Full-disk encryption on such mobile assets as laptops has been standard practice for certain types of users, and it covered OSs, program files, temp data, and user data. But virtual machines (VMs) can be much more mobile than such physical assets; they can be moved around at the click of a mouse to enable dynamic provisioning. Full-disk encryption must be applied to sensitive virtual images as well. In addition, moving to a policy-driven, data-centric encryption will ensure protection against copying through multimedia ports.

VMs are only as safe and risk-free as the host. Limiting the host's attack surface area (with fewer OSs and an optimal number of general-purpose endpoint applications) will make them that much safer.

Virtualization software itself could become a vulnerable area as attacks on endpoints get deeper and more persistent. Regular patch management of the software is a basic defense measure against that threat.

Attacks Related to Mobility

The ever-increasing computing power, convenience of form factor (made only more attractive by the new wave of tablets led by the iPad), and perpetual connectivity that most mobile devices offer these days have contributed to a significant growth of mobile endpoints that the enterprise must now worry about.

So it won't be long before even-newer OSs, such as iOS and Android, are the focus of targeted attacks. They do come with built-in local/remote wipeout features, as well as 256-bit Advanced Encryption Standard (AES) encryption, but they can be very vulnerable to targeted attacks unless they are required to operate within the perimeter. This can be done through virtual private networks (VPNs) and enterprise mobility server connections, or through corporate virtual desktop infrastructures (VDIs) in the case of tablet computers. The C-level clamor for tablets often sidesteps measures to have a formal mobile device management (MDM) system in place before allowing access to mobile enterprise assets.

Local data sitting on mobile devices is another significant area of vulnerability. With the abundance of native mobile apps, more and more mobile devices are storing data locally, thereby subjecting them to the same security threats as standard desktops and laptops.

Although some mobile OSs offer application "sandbox" capabilities that isolate an application and its data from other apps within the mobile device, not all apps are designed to take advantage of OS features. This situation is further exacerbated by the fragmentation of mobile OSs in the market today. All of it leads to inconsistent security.

An emerging and rapidly growing area in mobile apps is mobile commerce and payments. Banks and payment processors are already implementing solutions to make mobile devices into "virtual wallets," leveraging OS support and mobile hardware innovations such as near-field communication technology in modern smart phones, such as Google's Nexus S. These new features will also make it necessary for mobile devices to become compliant with, for example, the Payment Card Industry Data Security Standard (PCI DSS), to protect credit cardholder's personal data.

Wi-Fi sniffing has become a new concern. Software such as "Firesheep" has exposed the vulnerability of our Wi-Fi networks. In a typical coffee shop, which normally has a common, shared Wired Equivalent Privacy (WEP) password, an attacker running Firesheep can easily sniff out and decrypt the cookies of folks accessing Facebook and other social networks.

Threats Related to Social Networking

In social networking, the threats we need to worry about are server endpoints that now host customer data of some kind (as opposed to just marketing collateral and catalogs) and are public-facing—and, in many cases, on public clouds, which are owned/managed by

Layers Make Up the IBM Security Framework

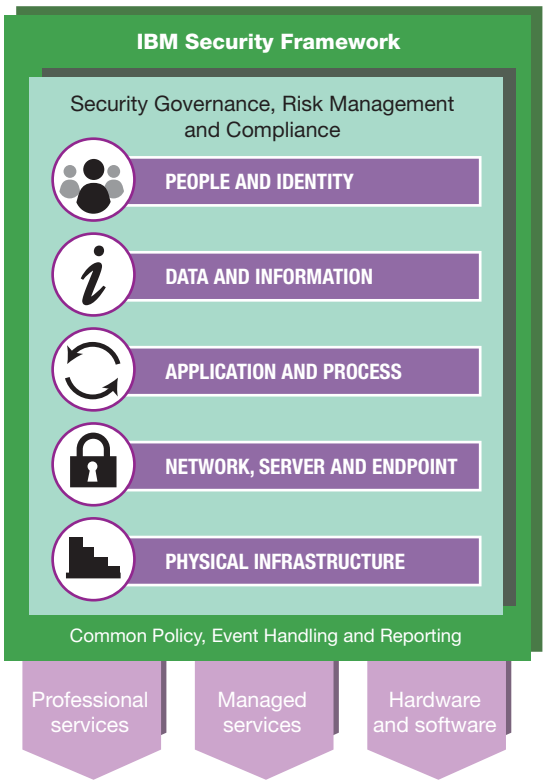

Figure 2 Source: IBM Redpapers, January 2011

the dominant social sites, such as Facebook and Twitter. So the concerns here would be about security of data and regulatory compliance issues related to the vertical industries involved.

E-commerce sites have had many years of experience in safeguarding themselves against theft of credit card information by complying with the rather exhaustive PCI DSS. This standard covers a broad range of topics related to security of data, network, and computers (endpoints) in a given industry. Similar regulatory standards include, but are not limited to:

• The Health Insurance Portability and Accountability Act (HIPAA), for healthcare

Buyers Guide: Security Applications & Services

Company	Focus Areas	Product(s)	Service(s)
ActivIdentity Corp.	AUTH, IDM, NS	x	x
Afilias	WS		x
Application Security, Inc.	COM, DBS, R, VA	x	
Arxan Technologies	EN, IPD, MALP	x	x
Aspect Security	APS, COM		x
Axway	EN, ES, IDM	x	x
Barracuda Networks	ES, NS, WS	x	
Beta Systems Software AG	APS, AUTH, COM, IDM	x	
Beyond Security	NS, SECT, VA, WS	x	
BeyondTrust Software	AUTH, COM, IDM	x	x
BMC Software	COM, IDM	x	
Catbird Networks	COM, VS	x	
Centrify	AUTH, COM, EN, IDM, VS	x	x
Check Point Software Techs	AV, ES, IPD, MALP, MBS, NS, VS, WS	x	x
Cigital	APS, R, SECT, VA		x
Cisco Systems	ES, MBS, NS, WS	x	x
Crosscheck Networks	COM, IDM, SECT, SOAS	x	
EdgeWave	ES, WS	x	
Entrust	AUTH, COM, DBS, EN, ES, IDM, WS	x	x
F-Secure Corp.	AV, MALP, MBS, WS	x	
F5 Networks	APS, COM, DBS, NS, R, WS	x	
Fiberlink Communications Corp.	AV, EN, IPD, MBS, VA		x
Fidelis Security Systems	COM, DBS, IDM	x	
Gemalto NV	AUTH, DBS, EN	x	x
GFI Software	ES, VA, WS	x	
Hewlett-Packard	APS, SECT, VA, WS	x	x
HP Fortify	APS, COM, R, VA	x	x
IBM	APS, AUTH, COM, DBS, EN, ES, IDM, IPD, MSGS, VA, WS	x	x
Imperva	APS, COM, DBS, R, VA, WS	x	x
Integrigy Corp.	COM, IPD, VA, WS	x	x
Juniper Networks	IDM, NS	x	x
Klocwork	SA, VA	x	
Layer 7 Technologies	COM, SOAS	x	
Lieberman Software	IDM, MALP	x	
Lumension Security	AV, COM, DBS, R, VA	x	x
McAfee, Inc.	AV, COM, DBS, ES, IPD, MALP, MBS, NS, R, VS, WS	x	x
Microsoft	AV, MALP	x	x
NCC Group	APS, DBS, NS, SECT, VA, WS		x
Oracle	APS, COM, DBS, IDM, R, VA, WS	x	x
Palamida	APS, COM, OSS	x	x
Panda Security	AV, WS	x	x
Parasoft	APS, COM, SA, SECT	x	x
PreEmptive Solutions	APS, EN	x	
Proofpoint, Inc.	AV, COM, EN, ES, MSGS, NS	x	
Protegrity Corp.	APS, DBS, EN, IPD	x	x
Quest Software	AUTH, COM, DBS, IDM, MSGS	x	x
Radware Ltd.	APS, COM, IPD, MALP, NS, WS	x	
Reflex Systems	COM, VS	x	
RSA, the security division of EMC	AUTH, COM, DBS, EN, ES, IDM, WS	x	x
SafeNet, Inc.	AUTH, COM, DBS, EN, NS, R, SOAS, WS	x	x
Security Innovation	APS, EN, R, SA, SECT, VA	x	x
SonicWALL	ES, IPD, MALP, MBS, NS, VS	x	x
Sophos	AV, COM, DBS, EN, ES, IPD, MALP, WS	x	
Sourcefire, Inc.	COM, IPD, MALP, NS, R, VS	x	x
SPAMfighter	ES	x	
StillSecure	COM, IPD, MALP, NS, VA, VS	x	x
Symantec	AUTH, AV, COM, ES, IDM, MSGS, R, VA, WS	x	x
Technology Nexus AB	AUTH, EN, ES, IDM	x	x
Thawte	WS	x	
Trend Micro	AV, DBS, ES, IPD, MALP, MSGS, WS	x	x
Trustwave	AUTH, COM, DBS, EN, ES, IPD, SECT, VA, WS	x	x
Veracode	APS, COM, OSS, R, SA, SECT	x	x
VMware, Inc.	COM, VS	x	x
WatchGuard Technologies	AV, IPD, MALP, NS, WS	x	x
Webroot	AV, ES, IPD, MALP, WS	x	
Websense, Inc.	DBS, ES, IPD, MALP, MSGS, WS	x	x

Key to Focus Areas

APS	Application Security	MBS	Mobile Security
AUTH	Authentication	MSGS	Messaging Security
AV	Antivirus	NS	Network Security
COM	Compliance	OSS	Open Source Security
DBS	Data/Database Security	R	Risk
ES	Email Security	SA	Static Analysis
EN	Encryption	SECT	Security Testing
IDM	Identity Management	SOAS	SOAS Security
IPD	Intrusion Prevention/Detection	VA	Vulnerability Assessment
MALP	Malware Protection	VS	Virtualization Security
		WS	Web Security

Bold = Sponsor

• The Federal Information Security Management Act (FISMA), for the public sector

• The Gramm-Leach-Bliley Act (GBL), for financial services.

Corporate websites that connect to the social-networking sites (with the famous "Follow us on" invitations and/or "Like" buttons) should reconfirm their authenticity by complying with Extended Validation SSL (EV SSL). This will demonstrate a commitment to the security of customers, and it will combat phishing as well.

Technology/Vendor Options and Solutions

In an effort to offer comprehensive protection against emerging and traditional endpoint threats, endpoint security products have evolved into suites offering the following set of features:

• Anti-malware, which broadly includes protection against malware, viruses, and spyware

• Endpoint firewalls, as a second level of defense behind the perimeter

• Host Intrusion Prevention Systems (HIPS), to prevent malicious attacks on servers and PCs

• Centralized management for patches, configuration, and reporting.

Gartner's term for these suites is "Endpoint Protection Platforms (EPP)," and in its December 2010 Magic Quadrant for EPP, the firm laments the fact that in 2010, malware effectiveness was on the rise in general, gaining an upper hand over recent enhancements in EPP products. The firm shares a concern that most of today's EPP vendors are working more on reactive signature-based detection techniques and less on attacking related root causes proactively. The Magic Quadrant of EPP vendors shows that Symantec, McAfee, Trend Micro, and Sophos lead the pack.

The report that accompanied this Magic Quadrant presents a thorough evaluation of EPP products from across the industry using an exhaustive set of business, technical, and financial criteria, and it highlights the strong and cautionary aspects of the products included in the analysis.

To build comprehensive endpoint security solutions, corporate IT groups and system integrators (SIs) first have to conduct a broader business and IT security analysis (including business drivers and threats). This will help them plan, build, and sustain a strategy for endpoint security. The approach should be an integral part of a broader enterprise cyber-security initiative and unified threat management (UTM) solution.

IBM strongly recommends, in its Redpaper on the IBM Security Framework,[1] that companies adopt relevant aspects of internationally accepted frameworks and best practices for IT governance. Two such recommended frameworks are Control Objectives for Information and related Technology (COBIT) and International Organization for Standardization 27002:2005 (ISO 27002:2005).

In addition, the vertical-industry security compliance standards that we listed earlier should be studied and strategically adopted.

The resulting planning, building, and sustaining/strengthening activities for endpoint security would involve:

• Business goals and related threat modeling and analysis

• Endpoint-security blueprint and related total cost of ownership and return-on-investment calculations

• Endpoint security use case analysis

• Endpoint security architecture (with an eye to how it fits with the company's overall enterprise security architecture) and implementation

• Ongoing vulnerability detection and patch management

• Audit and compliance enforcement

• Reporting.

PCI DSS applies to any application that either stores, processes, or transmits the primary account number.

Prognosis

Gartner's concerns around EPP playing catch-up to the exploits of the cyber underworld underscores the need of the hour and is a call to action for greater shared responsibility and accountability across the ecosystem. This catch-up is supposed to have taken the EPP products market to $3.96 billion in 2010, resulting in a 30-plus percent growth year over year. This growth speaks well of the efforts made, but the effectiveness of endpoint security must go up as well in 2012.

EPP vendors should certainly be analyzing root causes and developing remedies, but they should not be doing so in isolation. Collaboration is needed to plan and execute an industry-wide, unified counter-attack on cybercrime. Yes, scanning, patching, and personal firewalling through EPP is fine, but the overall security of the stack on the endpoint cannot be an afterthought.

Infrastructure vendors have to adopt "secure by design" principles (an evolving discipline that is actively being promoted by IBM and Citrix, among others) when they develop/enhance platforms that constitute the surface area of attacks on endpoints. And corporate IT and the SI community must do the same with the applications layer, showing some leadership in defining and implementing comprehensive UTM systems and policies. SW

Sreedhar Kajeepeta is global VP and CTO of technology consulting for GBS at CSC. CSC's consulting groups specialize in cloud computing, SOA, enterprise transformation, data warehousing and business intelligence/analytics, and application consulting (open source, JEE, and .NET). Kajeepeta is based in Farmington Hills, Mich., and can be reached at skajeepeta@csc.com.

[1] www.redbooks.ibm.com/abstracts/redp4528.html

2011 Software 500 Eprints & Reprints

Eprints & Reprints

Order Your Reprint Now!

- Reproduce your 2011 Software 500 ranking from the magazine
- Your company ranking is highlighted
- Includes your company description and contact information

THE 29TH ANNUAL SOFTWARE 500 RANKING

Bold URLs indicate sponsors

RANK	Private ●	Company / U.S./World Headquarters	Website	Total Software/Services Revenue: Worldwide 2010 ($ Millions)	Total Software/Services Revenue: Worldwide % Growth '09 to '10	Total Corporate Revenue ($ Millions)	% Growth Corporate Revenue '09 to '10	% Revenue on R&D	Total Employees	Software Business Sector
222	●	Flexera Software, Inc. Schaumburg, IL	www.flexerasoftware.com	$142.0	25.3%	$142.0	25.3%	13.5%	438	Software License Management
223		Fundtech Ltd. Jersey City, NJ	www.fundtech.com	$141.9	20.5%	$141.9	20.5%	15.4%	1,024	Enterprise Resource Planning (ERP)
224		Merge Healthcare Inc. Chicago, IL	www.merge.com	$140.3	109.9%	$140.3	109.9%	14.3%	[illegible]	
225		Infinite Computer Solns. (India) Ltd. Bangalore Karnataka, India	www.infinite.com	$140.1	31.1%	$140.5	[illegible]	[illegible]	[illegible]	
226		Basware, Inc. Stamford, CT	www.basware.com	$136.9	5.9%	$136.9	[illegible]	[illegible]	[illegible]	
227	●	ExactTarget Indianapolis, IN	www.exacttarget.com	$134.0	41.1%	$134.0	[illegible]	[illegible]	[illegible]	
228		Motricity, Inc. Bellevue, WA	www.motricity.com	$133.4	17.3%	$133.4	[illegible]	[illegible]	[illegible]	
229		Ebix, Inc. Atlanta, GA	www.ebix.com	$132.2	35.3%	$132.2	[illegible]	[illegible]	[illegible]	
230		Actuate Corp. San Mateo, CA	www.actuate.com	$131.5	10.2%	$131.5	[illegible]	[illegible]	[illegible]	
231		Sourcefire Inc. Columbia, MD	www.sourcefire.com	$130.6	26.2%	$130.6	[illegible]	[illegible]	[illegible]	

About Flexera

Flexera Software is the leading provider of strategic solutions for Application Usage Management. Flexera Software helps application producers and their customers strategically manage application usage to achieve continuous compliance, optimized usage and maximized value. The producer and consumer communities have moved beyond just tactical tools for application administration to now seek strategic solutions for Application Usage Management. Application usage has turned strategic; in the enterprise, application usage has moved from aiding the business to running the business; the competitive edge in hardware is increasingly based upon [illegible]

FLEXERA SOFTWARE

Flexera Software, Inc.
1000 E Woodfield Road
Suite 400
Schaumburg, IL 60173
PHONE (847) 466-4000
FAX (847) 619-0788
WEBSITE www.flexerasoftware.com
E-MAIL areist@flexerasoftware.com

Contact:

Carol Samost
National Sales Manager
csamost@softwaremag.com

Continued from p. 32

thing you can do is tell the team how to solve their problems. The best thing to do is help them identify that they *have* problems and issues. You can help develop solutions for these issues, but it really has to be their resolution, because they have to own it and make it work. If they think it is your solution, they can blame you if it does not work. As the ES, you want the team to own the solution. The only thing I can manage is my time; otherwise, I am micromanaging.

Johnson: What is the difference between being a coach and a manager?
Coleman: Managing implies that you are telling somebody what to do: "You are going to stand at this place in this production line. You are then going to pick up the wrench, and when this part comes by you, turn the screw to the right until you see 4.4 pounds of pressure. You are then going to stop and wait for the next part to come down the belt." If you have managed everything about that job, you have paid not one cent for the intelligence of the individual. In the world we are in, you are hiring people for their knowledge and intelligibility.

Johnson: Tell us about one of your early projects.
Coleman: One of my first projects was the VisiOn project for VisiCorp. It was the first window system ever done for the PC in the early '80s. Terry Opdendyk was the president, and I was the director of product development. I was presenting the project plan to Terry, and he started to ask me the "what if" questions. I began to think we didn't have backup plans. These were the days of the waterfall development process, and by the time you learned you had the wrong approach, you were way down the development cycle. So, not only didn't I understand if things would succeed, but I didn't know how I was going to determine if things were on the right track. I determined that we should build a core product and make sure it worked, and then scale out functionality in an iterative process.

Johnson: What is your secret for getting things done?
Coleman: Scott [McNealy] had just reorganized Sun from products into functions, and I ended up with all six software divisions. The Solaris project was a hallmark for me in what we actually had to do. Before then, software had reported to hardware, and we had six versions of the SunOS, one for every hardware type. Every engineer in the company thought they could make changes to the operating system without formal testing. Sun, like Google today, equated process with bureaucracy and wanted no part of it. My feeling is bureaucracy is the reaction when the organization does not know how to get things done. So my first task as owner of the new software division was to create the software development framework [SDF], which was really a group of processes, but we kept that a secret.

Johnson: Give us an example of how you got things done.
Coleman: SDF had to do two things: guarantee a new release every six months, and each release had to improve both functionality and quality. SDF was very simple; it separated all the releases into a "train." Then we divided all the software groups into fewer than 10 people, including a marketing person and testers. The group had to own something that could be identified as to its competitive performance in the marketplace. So the team would own the product, and it was their responsibility to be the best in the market. A train would leave every six months, and no product could hold up the train. You could not get on the train unless you passed the quality and performance tests. No team wanted to be left off the train, and no one wanted to be responsible for their team not being on the train. It was all about individual and team accountability.
Johnson: So individual ownership was a big part of the success of the project?
Coleman: So we had these small groups, but they might be part of a bigger subsystem. The subsystem would also have to qualify to be on the train that left every six months. In order to manage this process, we did a software build every night and published the results online. Every engineer in the company could see if they were going to be part of the train or if they needed to do something to make sure they were on the train. If something was denied entry, it was flagged, and the group that owned it would have to make sure it was corrected or completed.

Johnson: What are the responsibilities of an ES?
Coleman: The leader of an organization only has three first-order responsibilities: the three Vs, hiring, and organizing. The three Vs are the most important—vision, value, and valuable. The leader must establish a compelling vision of why you are in business or why you are doing the project. The leader must prove why the project has differentiation. The leader must show how this differentiation has value to the customer. Then the leader must prove that the project is valuable to the organization and supports the vision.

Johnson: How do you handle the communications problem?
Coleman: If you develop a good process that is repeatable and measurable, it comes with a common language. It also comes with a set of metrics. Because you are so focused on having to deliver objective results, it forces a common vocabulary. Part of that vocabulary is building transparency and accountability into the process. A common vocabulary provides objectivity so there can be no ambiguity in the results. SW

softwaremag.commentary

HISTRIONICS

By Jim Johnson

Q&A: How to Be a Good Executive Sponsor

The single most important person in any project is the executive sponsor (ES). His or her actions and decisions have a huge impact on the success and delivery of a project. However, Standish Group research shows that the majority of executive sponsors have little or no formal training on how to be an ES.

In fact, most organizations expect their ESs to get their training "on the job." It is, therefore, the project manager assigned to the project the ES is supporting who supplies most of this education.

There are a number of problems with this practice: reliance on the project manager, the project manager's subordinate role, and differing views, approaches, and opinions about how to be a good ES.

The Standish Group set out to discover, map, and assess the skills needed to be a good ES. Part of that quest included an interview with Bill Coleman, founder of BEA Systems and venture angel.

Johnson: Tell us about one of the projects where you were an ES.

Coleman: I was the ES for the Solaris project at Sun Microsystems. Solaris was a brand-new operating system and would change the whole company—it would be the center of all Sun products in the future. We were under a mandate to build a brand-new operating system from scratch. Every line of code was to be written from the ground up. We used nothing from SunOS. The first thing I did was to select Steve Bourne as my program manager. The first thing Steve did was to ask me what the success criteria were. My number one success criterion was predictability.

Johnson: What was the single most important thing you learned from the Solaris project?

Coleman: The Solaris project was the future of the company [Sun Microsystems]. Our plan called for the project to last two and a half years, ending in June 1992. I wasn't concerned about shipping on that exact date. What I was concerned about was that we put a program in place that was predictable, so that we would know when we would ship, what we would ship, and the state of what we shipped. We needed to know this because the entire company was basing its plans around when and what we shipped. So if we got to six months before ship date and found we needed another year, that would sink the company. However, if a year before ship date we found we would have to make some tradeoffs with scheduling for function or cost, then it would be in the company's control. Steve Bourne, program manager, and I would have a one-on-one meeting every week to go over the progress of the project. The number one thing that an executive sponsor can do is set the expectations and clearly state how the team is being measured.

Johnson: What is your approach to being an ES?

Coleman: An ES is a coach, not a manager. As coach, you want to be a sounding board to listen to the team's issues and problems. The first thing the team needs to do is to write down their goals and objectives. So they take this huge, fuzzy problem and deconstruct it into manageable chunks. However, they always have to come back to the vision. The vision has to be from the top down, because you always have to come back to the reason you are doing the project in the first place. This starts from first principles—the most important goals. These principles then turn to metrics and measurables that turn into processes that can be predictive and quantified.

Johnson: What is the worst and best thing you can do as an ES?

Coleman: As project mentor, the worst

Continued on p. 31

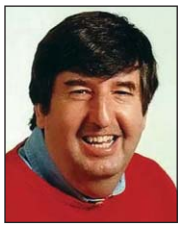

Jim Johnson is the founder and chairman of the Standish Group. He has been involved in the computer industry for more than 40 years and has a long list of published papers, articles and speeches. Go to www.standishgroup.com.

www.softwaremag.com

Advertising Sales

National Sales Director
Carol Samost

csamost@softwaremag.com
PHONE: 617-962-4725
FAX: 508-668-0442

MAILING ADDRESS:
Software Magazine
P.O. Box 135
E. Walpole, MA 02032

Contact Sales Director for information on:

- Advertising packages
- SoftwareGram email newsletter sponsorships

Eprints • Reprints

WINTER 2011 ADVERTISER INDEX

DIGITAL SOFTWARE MAGAZINE The Software Decision Journal

This index is a service. The publisher does not assume any liability for errors or omissions.