TECHNICAL PAPER

Cyber Risk Quantification Using the Balbix® Platform

Contact us at *info@balbix.com*

https://www.balbix.com

Balbix®

Introduction

The enterprise attack surface has exploded. There are practically unlimited ways in which enterprise networks can be breached. Figure 1 shows an abstract view of the attack surface.

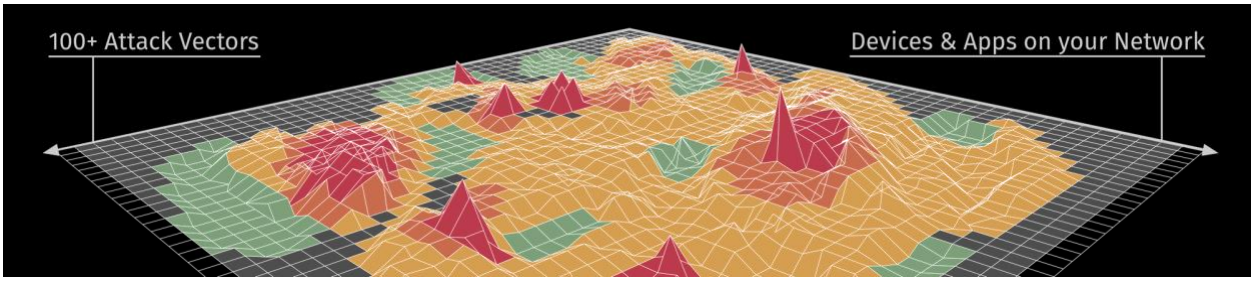

Figure 1: The enterprise attack surface

The right axis has your assets (maybe tens or even hundreds of thousands). These are your devices, applications, users and the data you are trying to protect. These elements can be on-prem, in the cloud, mobile or in your supply chain.

On the left axis, we have the various attack vectors – 100s of them, ranging from simple things like weak passwords, to more complex things like phishing, unpatched software, password reuse, encryption issues, misconfiguration, etc.

Each point on the graph represents one way in which the adversary can attack your network. The z-axis is risk, which estimates the expected loss from an attack. The red areas in the picture represent areas of greatest risk; areas where the infosec team of the enterprise needs to focus first to mitigate cyber risk. The yellow areas represent the next set of security issues that need to be addressed. In an ideal world, infosec teams would identify and fix red and yellow areas before any attacker could find and leverage these areas to breach the enterprise.

Unfortunately, as you probably know, the above picture is not easy to get. The reason? It is complicated to calculate cyber risk. There are numerous reasons for this.

Cyber Risk Quantification

Risk is defined as the probability of a loss event occurring in a given unit of time (Likelihood) multiplied by the expected magnitude of loss resulting from that loss event (Impact). *Cyber* risk is the expected loss resulting from a cyberattack or data breach. This definition of cyber risk is shown in Figure 2.

In the risk equation, the units of Impact are monetary *units,* e.g., dollars, euros, yen, pounds etc. The units of Likelihood is probability, typically expressed as a percentage value. Multiplying these two factors together gives us the units of risk as *monetary value of* the expected loss event occurring in a quarter or a year or some other defined time period.

Figure 2: The risk equation

As you can imagine, the ability to quantify cyber risk accurately is key to making the right decisions about your cybersecurity posture. There is no such thing as zero risk, and the leaders of each organization must decide for themselves if their organization's residual cyber risk is acceptable. It is critical that cyber risk is quantified in monetary terms so that your CFO, CEO and board can appreciate the amount of risk in business terms. A risk score that is denominated as high/medium/low or on a scale of 1-10 is not very useful for decision making. For example, your senior stakeholders will not be able to appreciate that a score of 7/10 corresponds to a high likelihood of a data breach with a $25M price tag. The red-yellow-green coloring of Figure 1 must be driven by monetary value thresholds that represent the risk appetite of the organization.

Similarly, a risk value that does not take all the above factors into account: vulnerabilities, exposure, threats, security controls and business criticality, won't be considered trustworthy by your infosec team members who will question the incomplete logic. Also, a risk calculation that does not cover the vast majority if not all the attack surface of Figure 1 will be misleading and lull you into a false sense of security.

Finally, it is important that the risk calculation both explains itself, i.e., what factors are driving the risk, and be actionable. The risk calculation must also be accompanied by a prioritized set of mitigating actions, with price tags and expected values of post-action risks.

Computing the risk equation for each point of the attack surface of Figure 1 is not easy. Let's see why, and how it impacts your ability to quickly remediate those risks.

3 Brutal Truths

Let's say you have invested in dozens of cybersecurity tools in the last few years. You are trying to manage your vulnerabilities aggressively, have deployed an endpoint detection and response (EDR) solution and next-gen IDS/IPS and have invested in a SIEM and have implemented detection and containment playbooks. You do regular pen-testing, perhaps augmenting this with a Breach Attack Simulation (BAS) tool and an outside-in risk scoring tool. Figure 3 shows an example of a cybersecurity tools stack with all the bells and whistles.

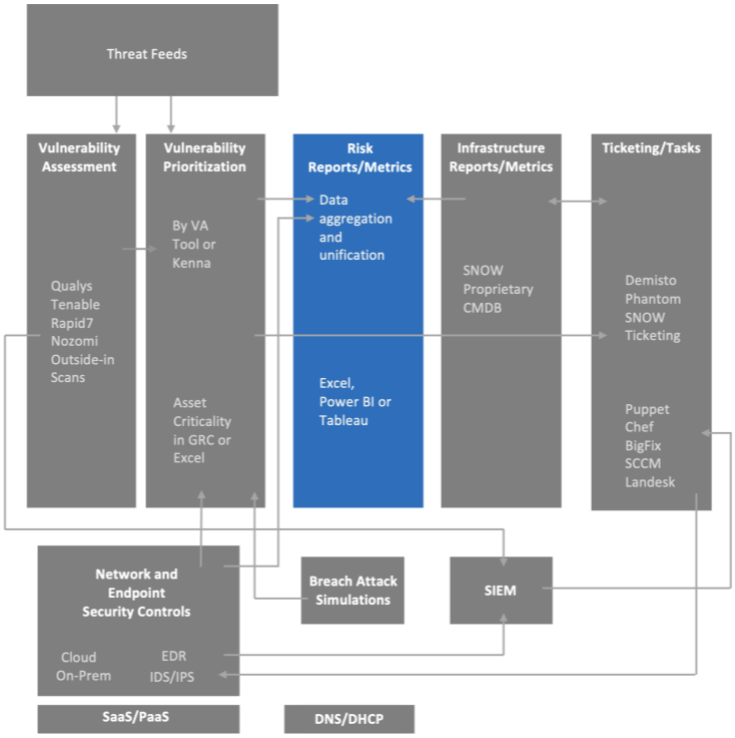

Figure 3: Cybersecurity tools stack

The blue block in the center of Figure 3 is akin to the "brain" of your cybersecurity program. This is where you are trying to get an overall view of your cyber risk and use this information to make tactical and strategic cybersecurity decisions.

Here are three brutal truths about the risk calculation and how that affects cybersecurity posture.

1. **Data ≠ Visibility.** Your dozens of tools may be generating lots of data. However, mountains of data do not result in better cybersecurity visibility. It is difficult and time-consuming to sift through terabytes or petabytes of data to find what you are looking for. In addition, organizations continue to struggle with gaps in asset inventory—they can't accurately enumerate the right axis of Figure 1. As a result, security leaders constantly worry about unseen cyber risks and vulnerabilities. They also don't know the effectiveness of deployed security tools in mitigating risk.

2. **Can't Calculate Risk.** Tools have different formats and semantics for the same attributes of assets or users. They will often surface contradictory information. Cybersecurity context tends to reside mostly in specialized cybersecurity tools. IT context is spread across multiple IT tools such as AD, CMDB and ticketing systems, while business context tends to be spread across a 3rd set of databases and spreadsheets. Unifying this data into a common schema is very difficult. Different

stakeholders also speak using different terminology, and it is nearly impossible to reconcile to a commonly understood risk metric.

3. Partial Remediation. New vulnerabilities and security issues emerge at a very rapid rate. It is very hard to keep up. Because risk cannot be calculated, security issues and risk items cannot be quickly identified, prioritized and remediated. In most organizations, the mean time to mitigate security issues is weeks or months, and during this time the organization is open to compromise by attackers. You are also constantly wondering if your existing security controls compensate for these emerging risk items, or not. Or if you need to invest in additional tools?

Quantifying Cyber Risk with Balbix

Balbix ingests data from your various security tools via connectors or through direct observation from specialized sensors (Figure 4). The data is normalized, de-duplicated and analyzed to produce a comprehensive view of cyber-risk for each asset, group and for the entire enterprise, along with relevant context and recommended action items.

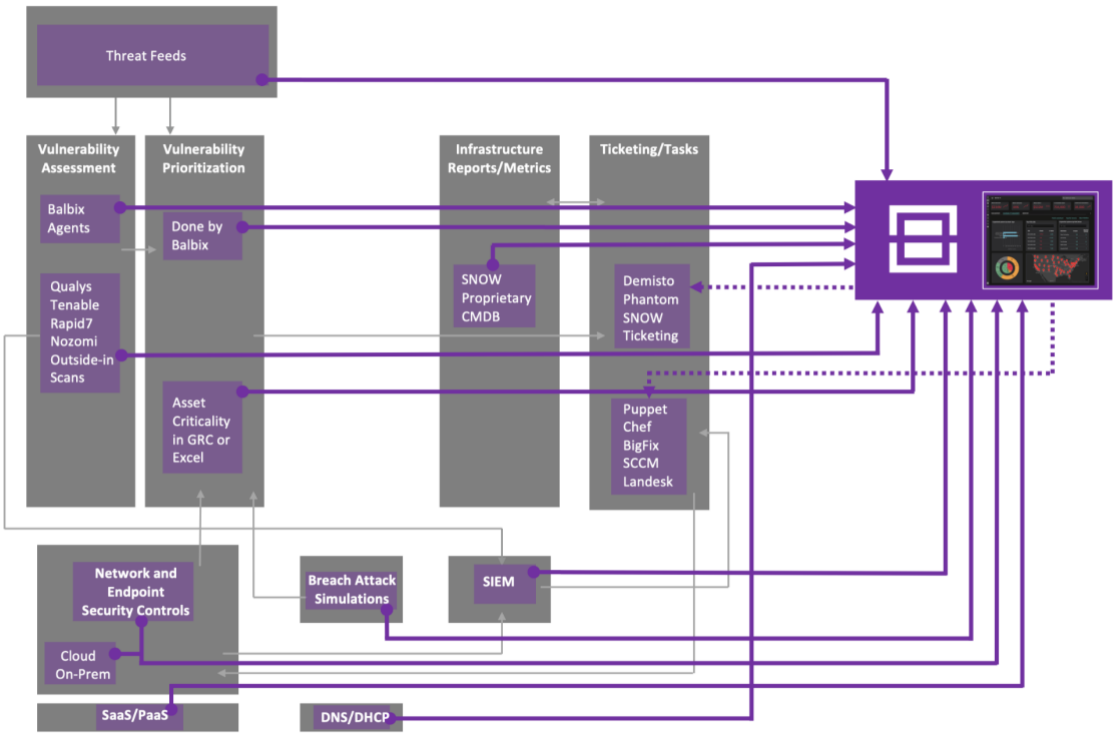

Figure 4: Balbix ingests and analyzes cybersecurity data

Figure 5 shows a "brain scan" of Balbix. On the left of the picture, we have an example of typical inputs to Balbix- various IT, cybersecurity and business data sources. On the right, we have the outputs – risk metrics, mitigation plan and actions, alerts/notifications, benchmarks, scorecards and trends. This output

is available to users via online dashboards, customizable based on role. The Balbix output is also available via APIs to other tools and systems.

Each interior node represents an ensemble of specialized ML models that has been purpose-built to solve a specific problem. For example, the host enumeration (*Host Enum)* node performs deduplication of assets across all data stream signals that are fed into the Balbix brain and provides this information to all nodes in the system. As we will see below, *Host Enum* is a powerful self-learning ML system that would be impossible for humans to replicate.

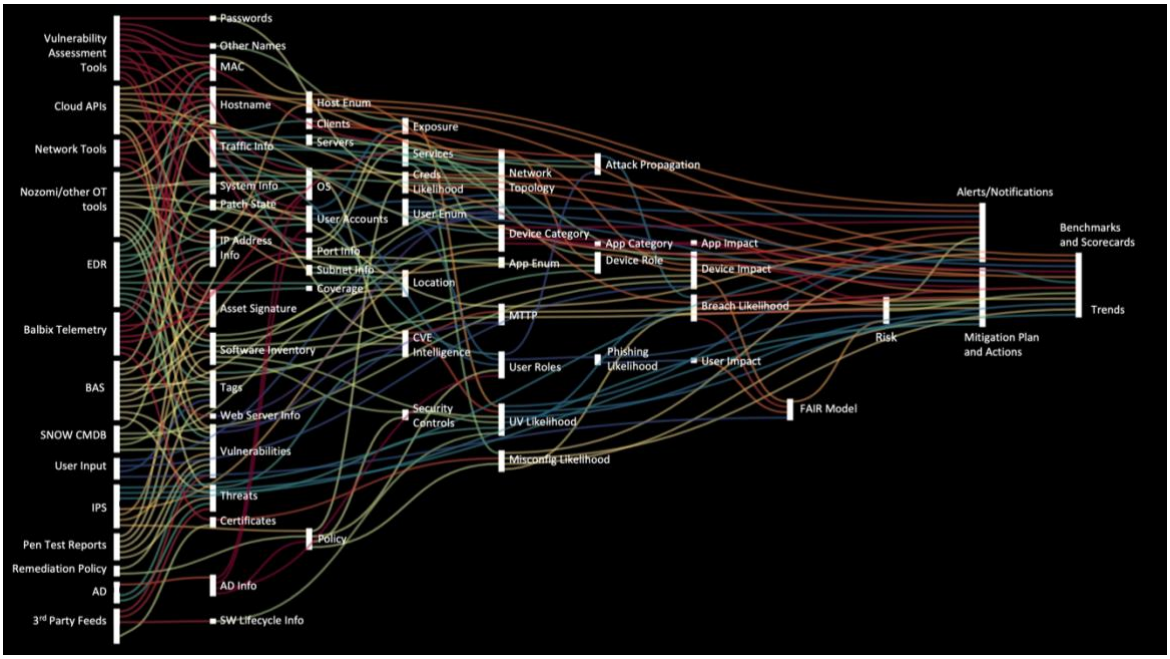

Figure 5: Converting data to insights in Balbix

Data Sources

Balbix can ingest data from common IT, cybersecurity or business tools using our library of connectors and flexible connector framework. Example data sources include vulnerability assessment tools, CMDB, EDR, firewalls, SIEM, MDM systems, AppSec systems, OT/IoT management systems, Active Directory, DNS/DHCP and cloud infrastructure APIs. Balbix can also quickly ingest data from proprietary tools.

There are two types of connectors: *streaming connectors* and *snapshot connectors*. Streaming connectors connect via API to the data source and pull in data on a specified schedule. Snapshot connectors ingest data dumps from your tools in .csv or other formats. You can mix and match streaming and snapshot data sources.

An initial configuration of data sources from a typical Balbix customer is shown in Figure 6. From this screenshot, you can see the variance in the coverage of the different tools and the significant data sizes that need to be processed. Data sources can readily grow to encompass 15-30 tools in full production with data volumes of several 100s of terabytes/day for ~250K assets.

Data Source	# Assets	Data Processed Per Hour
Nozomi	27K	227MB
Tenable	6000	250MB
Rapid7	125K	125GB
Kenna Security	17K	63MB
ServiceNow CMDB	221K	12GB
CrowdStrike Falcon	69K	5GB
Tanium	121K	10GB
IPAM System	25K subnets	100MB
Policy Information	121K	25MB

Figure 6: Example data sources

Not all data sources for Balbix need to come from tools deployed within your enterprise environment. It is straightforward to take a penetration test report or outside-in scan report in *.csv* or *.xlsx* format and ingest it into Balbix.

Balbix also provides flexible and secure choices for data movement. Data can be made available to Balbix via SaaS APIs, S3/Blob storage, Doc reps as well as on-prem data sources via a special collector.

The ingestion layer can be programmed with

- authentication info
- the shape and size of the data pulls (which rows and columns to fetch)
- the desired frequency of data pulls.

This is shown in Figure 6.

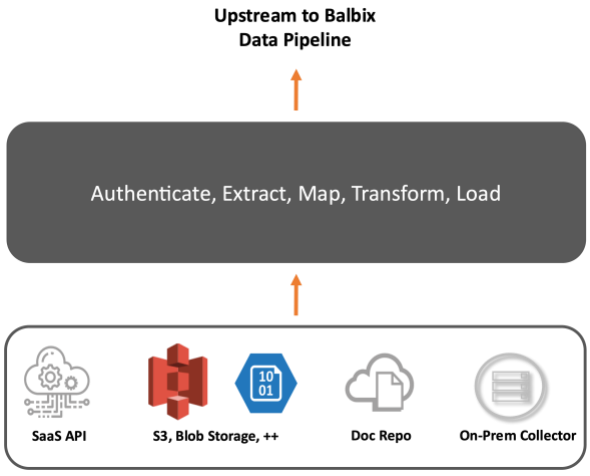

Figure 7: Data Plumbing

With Balbix, you can also optionally ingest data into the Balbix brain by using Balbix's specialized network and host sensors. These sensors are useful when you don't have security and IT tools of the type shown in

Figure 6 already deployed. Balbix sensors can perform *asset discovery and inventory*, as well as *continuous vulnerability assessment* in a non-intrusive manner.

Asset Inventory

You can't quantify what you don't know about. Asset Inventory is the foundational component of the risk calculation. This corresponds to the enumeration and categorization of assets (the right axis of Figure 1).

Host Enumeration Logic

The first challenge in asset inventory is matching signals from different data sources about a given asset to the same "asset" entity in the Balbix brain.

Consider the following challenges:

- IP addresses may be dynamic or stable attributes of a system depending on the asset type
- Multi-homed systems may have multiple IP addresses
- Roaming systems will have different IP addresses at different times
- DNS, MDNS and DHCP names are generally synchronized but there may be delays, and duplicates.
- The presence of virtual machines (especially transient VMs) complicates asset identity mapping.

These are all challenges that are known to anyone who has ever tried to correlate events from different IT and cybersecurity tools.

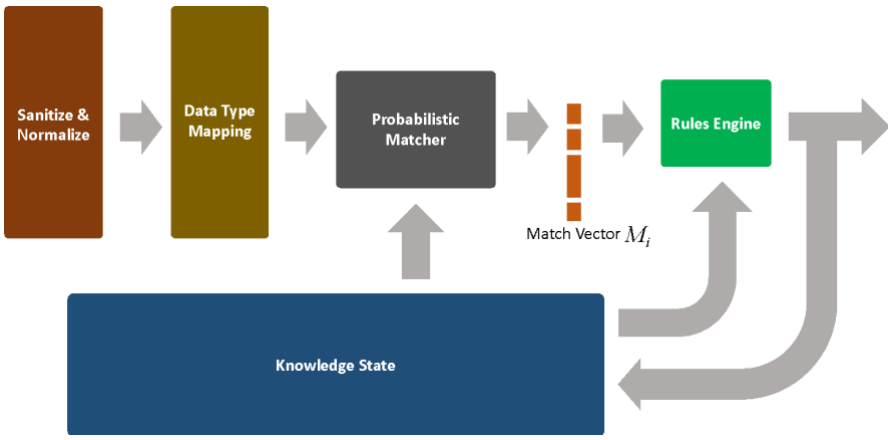

Figure 8: Host Enumeration Logic

Balbix's Host Enumeration Logic capability (otherwise known as HEL) implements the deduplication and mapping of signals to unique asset IDs. As signals from different data sources reach the Balbix Brain. First the signal is cleaned up by recognizing, sanitizing, normalizing and rewriting IPs, hostnames, mac addresses and timestamps. Then a finite state machine runs through a pipeline of statistical matches for each new signal against the Brain's learnt knowledge of asset IDs, subnets and attributes at different times to produce a match vector of probabilities for the signal against assets. A rules engine uses this

probability vector and additional learnt knowledge about the enterprise network to update the knowledge state of assets, subnets and attributes.

In summary, there are two things that happen in HEL:

1. Sanitization and normalization
2. Intelligent probabilistic matches that consider global knowledge state not just exact attribute matches

Balbix's HEL has significant differences from other approaches to deduplication. First, Balbix's observation-to-asset mapping process **is stateful**. It uses learnt knowledge. For example, the knowledge that a specific subnet is a server subnet in a data center is utilized by the probabilistic matcher. Secondly, Balbix's HEL is **self-learning**. For example, the system initially may not know that a system is multi-homed, but after sufficient observations indicate that this is so, the system learns the knowledge that these 4 assets are really the same asset. Finally, unlike other systems that are either purely probabilistic or completely deterministic, HEL is both **probabilistic and deterministic**.

In practice, HEL works very well to map signals from a broad variety of data sources to unique assets. HEL eliminates the drawbacks of **traditional methods which are fragile, require complete and consistent information schema across the enterprise, and an impossible amount of ongoing product and vendor-specific tuning to stay accurate.**

Asset Categorization

We know that the best human experts can put together an accurate picture of the type and category of a device on the enterprise network by manually looking at a broad variety of data sources. For example, from Layer 3 packet analysis, an expert might be able to extract MAC address OU information that indicates that a device is a Cisco device. At Layer 4 they might see transport headers and protocol behavior that are consistent with the device being a switch with a management portal available at ports 80 and 443. From Layer 7 analysis of protocol behavior and a study of artifacts rendered in the web browser, we might be able to say that port 80 does not automatically redirect to port 443, and that the device is a WLAN controller made by Cisco.

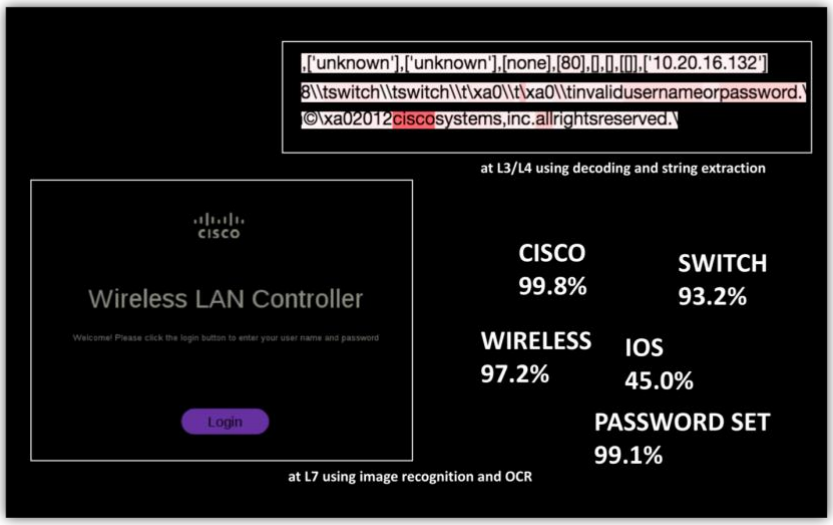

Figure 9: Asset categorization

At Balbix, we try to mimic this type of intelligent analysis by throwing in this type of L3, L4 and L7 data from different vantage points on the enterprise network into a tensor-to-tensor deep neural network which categorizes all devices, users, and applications on the extended network. The system is near real-time, with entities analyzed the second they show up on the network. Once the confidence level of the algorithms is above a threshold, they system surfaces an analyzed device in the appropriate part of the model, indexed with all the relevant attributes. A visual for this is shown in Figure 4.

Another example of hyper-dimensional machine learning is the algorithm Balbix uses to figure out the users associated with certain types of unmanaged devices. Balbix is generally able to identify users associated with laptops and desktops either by linking them with Kerberos authentication or domain logon sessions. This technique does not work very well to identify users of unmanaged devices.

Now as you can imagine, when you walk into the office with your laptop and smartphone together, and then these devices start transmitting and receiving on the network in a somewhat time-correlated fashion. One day, you might walk in with a colleague, confusing the correlation, but after a few days a clear picture of linkage between your devices should emerge. The Balbix system uses time-domain correlation of DHCP lease renewals, begin-of-day traffic time, on-off times, and timing of Gmail push notifications to different devices to figure out if a smartphone and a laptop/desktop have the same owner.

For typical enterprise networks, as long as the system gets 24 hours' worth of reasonable telemetry data, the accuracy of Balbix's categorization module is over 99.9%. When Balbix is unable to categorize with high confidence, the system marks the asset(s) as "partially categorized" while highlighting the low telemetry areas of the enterprise network.

Breach Likelihood Calculation

In Balbix, the overall Breach Likelihood of an asset is calculated as a weighted sum of Breach Likelihood from individual attack vectors. This is shown in Figure 10.

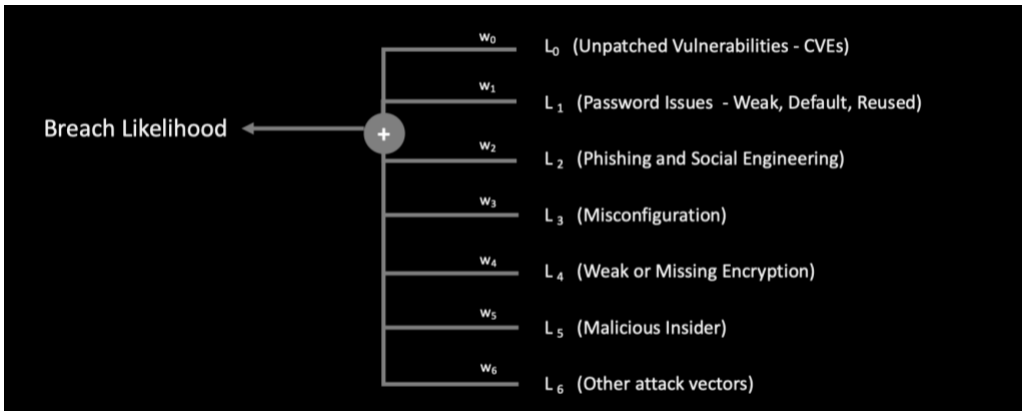

Figure 10: Breach likelihood calculation

Each Likelihood factor L_i is computed using a function of the form below with 4 inputs:

$$L_i = f\,(\text{Vulnerabilities, Exposure, Threats, Mitigating Security Controls})$$

Each of these inputs is computed by a separate ML model which is specific to the attack vector, and whose inputs are observations specific to the attack vector. These models can also be bypassed by ingesting calculated factor values from 3rd party systems.

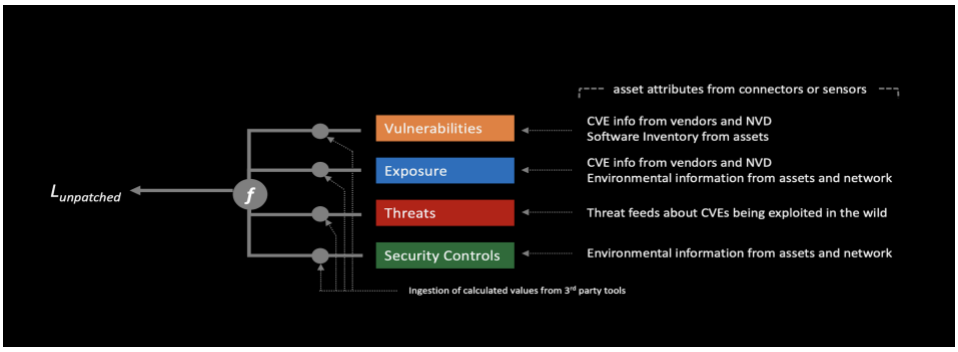

Figure 11: Breach likelihood due to unpatched software vulnerabilities

Balbix uses Probabilistic Graphical Models (PGMs) for estimating vulnerability levels for various attack vectors.

- PGMs encode domain knowledge
- Independent distributions are learnt through continuous self-learning
- Dependent variables updated using conditional probability distributions

An example PGM (for phishing) is shown in the picture below.

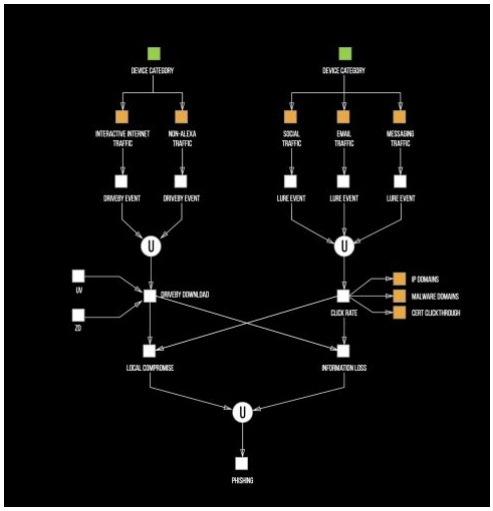

Figure 12: Probabilistic Graphical Model for predicting phishing likelihood

Breach Impact Calculation

In Balbix, breach impact is estimated as described in Figure 13. First, Balbix first uses the ingested data to compute relative asset criticality of various assets automatically. This algorithm is described in more detail in a later subsection below. As part of this automatic step, Balbix understands how devices and applications are related to other assets, and the infrastructure components that each asset are depends on for normal operation. Users with privileged access and high-value users from a business standpoint are also identified as part of Step 0.

After the automatic Step 0, the CISO (or equivalent or delegate) provides a best estimate for the impact of a major data breach in the enterprise in monetary units ("*I believe a major breach here will cost us at least $25-$30M*"). Balbix uses this estimate to recalculate the breach impact values of each asset and asset group using the relative asset impact distribution that was calculated automatically in Step 0.

Now every risk owner is asked to review the estimated breach impact value in monetary units for the asset groups that they are responsible for. They can adjust these values, and Balbix automatically updates the relative impact distribution of the enterprise as well as the overall impact and risk values. This step happens in a distributed manner with (hopefully) all risk owners chiming in.

Finally, the CISO can make a final top-down adjustment if necessary. Now we arrive at a breach impact distribution for all assets, including infrastructure nodes. These breach impact values are understood and appreciated by the risk owners (because they were part of the estimation process), and correctly influence the breach impact values of the infrastructure assets that the various business applications and processes rely on.

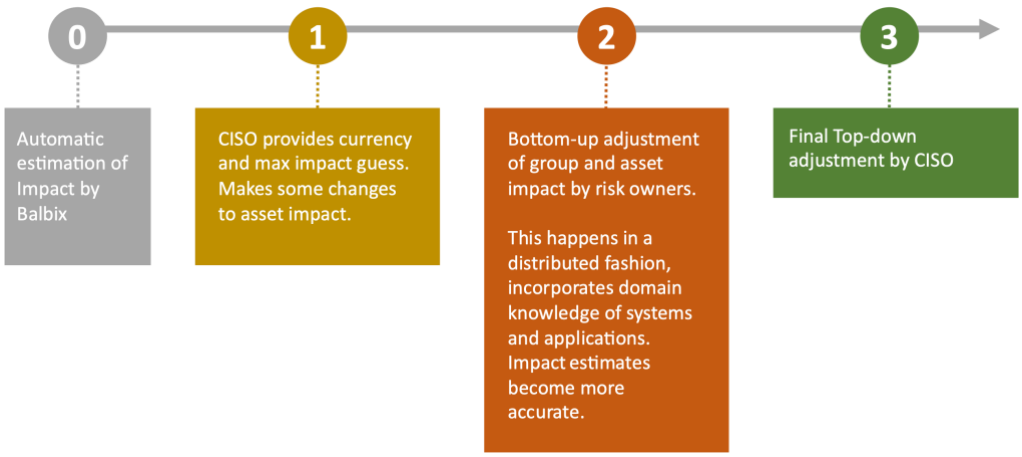

Figure 13: Asset impact calculation

Automatic Impact Modeling

In Balbix, automatic impact modeling takes into account both inherent properties (e.g. asset category, business unit) and contextual properties of the asset (roles, applications, user privilege and interaction with other assets). The impact model at Balbix is powered by Bayesian ranking system which comprises three key steps:

1. Prior Impact Model

First, based on expert knowledge, a coarse-level prior estimate of impact for each category of asset is derived. For example, servers and core routers are more important than display devices or smartphones. Additional impact points are accrued based on special privileges of the asset users (e.g., network, OS, application administrators, or C-level executive user machines are considered more important). Host roles and applications are also considered - a GIT server is deemed more important than an intranet web server.

2. Contextual Impact Model

Next, Balbix incorporates information about the duration and frequency of use, and the volume of traffic flows between pairs of assets across ports, protocols and applications. Interactions are modeled using Dirichlet Rank, a Bayesian version of the well-known Google Page Rank algorithm.

3. Effective Impact Model

Finally, the prior model and the contextual model are combined using a probabilistic programming technique called variational inference implemented on a deep learning stack.

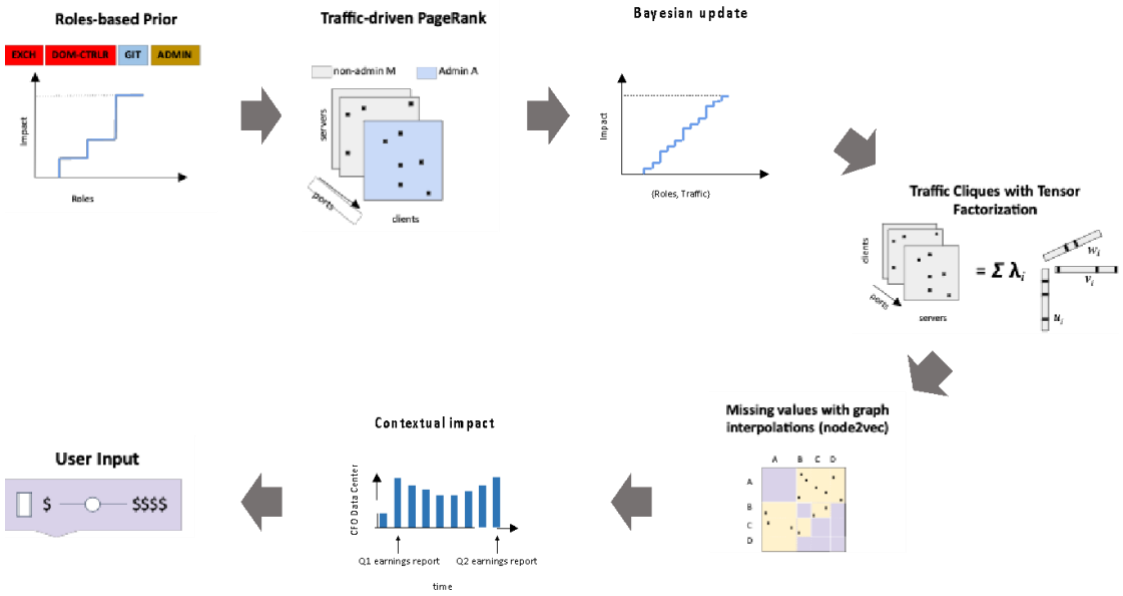

Figure 14: Automatic impact modelling in Balbix

The resulting impact model is completely unsupervised, requiring no input from the user about the criticality of a given asset. However, it is designed to easily accept implicit and explicit user signals regarding business criticality using group memberships, user-assigned tags. The final output is an impact score for each asset, along with a confidence measure, which is a real-time estimate of its business criticality in context. Figure 14 shows a high-level flow of Impact Modeling in Balbix.

Impact Adjustment

Impact adjustment for an asset group or even a single asset can be accomplished easily in Balbix as shown in Figure 15.

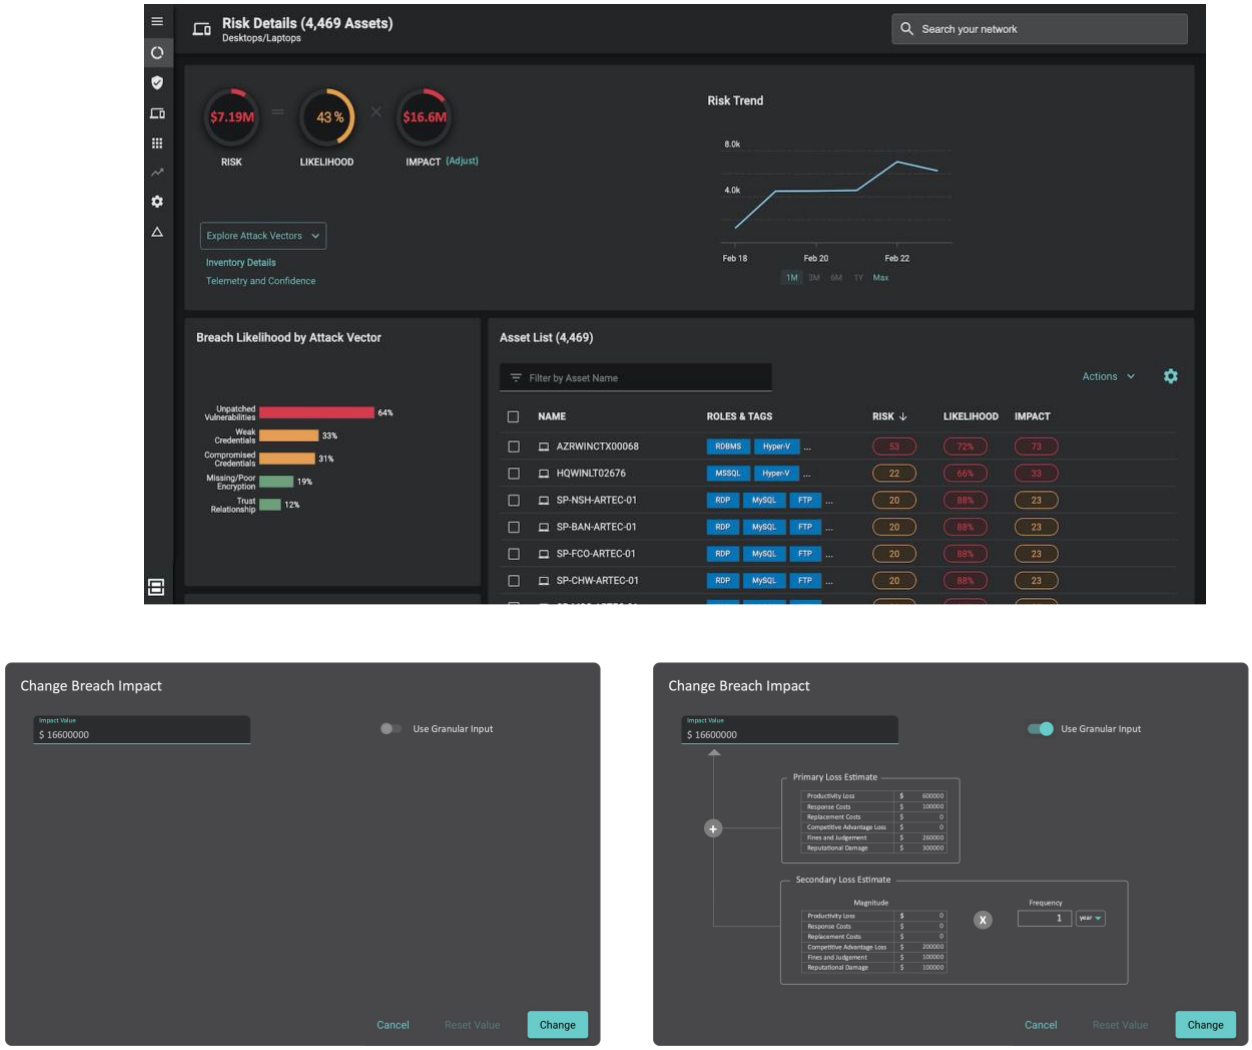

Figure 15: Asset impact adjustment

Breach impact values and tags can also be adjusted in Balbix by ingesting impact values in a file (.csv or similar format), or via API from a CMDB or GRC system.

Putting this All Together

Going back to Figure 2, Balbix autonomously and continuously ingests data from a wide variety of sources and performs correlations across hundreds of dimensions. The ensemble of models shown in Figure 5 is chained together and works in coordination to build an overall picture of your organization's cybersecurity posture and breach risk, device by device, group by group and site by site. The ensemble consists of both unsupervised as well as supervised models.

Getting user input into the models is built into the system in two flavors: via connectors from your high confidence IT and security tools or via the UI.

In essence, your Balbix instance develops the following categories of intelligence:

1. Understanding of every relevant detail (configuration, usage, etc.) of your extended enterprise inventory – all devices, users, and applications, on-premises and off.
2. Context around business criticality of each asset and user.
3. Up-to-date knowledge of global and industry-specific threats – aka what is fashionable with the adversary on a daily and weekly basis.
4. Understanding of the various security products and processes you have deployed as part of your overall breach risk mitigation strategy.
5. Calculation of your effective risk that incorporates all the information gathered above and predicts how and where you are most likely to be breached.
6. Prescriptive insights into how you might best configure and enhance your security controls and processes to mitigate risk and improve your cyber-resilience

Use-cases

What can you do with Balbix?

The Balbix platform is used by organizations ranging in size from startups to the Fortune 10 for dozens of use-cases. The top 4 use-cases are:

1. Unified cybersecurity asset inventory
2. Enterprise vulnerability prioritization
3. Automated risk calculation
4. Security analytics and Board-level reporting

Once these use-cases are up and running with Balbix, we have seen our customers use Balbix for some very interesting additional use-cases listed below:

1. Periodic Asset Criticality Analysis
2. Operational Dashboards for Vulnerability Prioritization and Dispatch to Risk Owners
3. Gamification of Cybersecurity Posture
4. Threat hunting
5. Risk-aware incident response
6. Planning of cybersecurity architecture enhancements
7. Cybersecurity stack optimization

There are many key cybersecurity questions that Balbix can help you answer for various stakeholders in your organization. A sample of these is shown in Figure 16: Key questions that Balbix helps you answer.

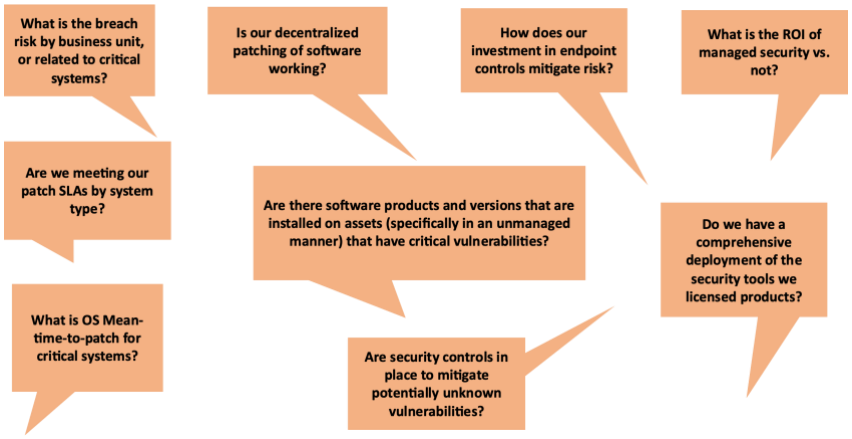

Figure 16: Key questions that Balbix helps you answer

Business Outcomes

Ultimately, the objective of cybersecurity is to stay a few steps ahead of the adversary. Being able to quantify cyber-risk enables organizations to streamline the use-cases described above. Hard risk metrics denominated in monetary terms with clear action items enables everyone involved make better decisions faster. The net of this is that the organization is able to identify and mitigate vulnerabilities and security issues faster, even as they emerge.

The effect of faster risk mitigation on cyber risk can be seen in Figure 16.

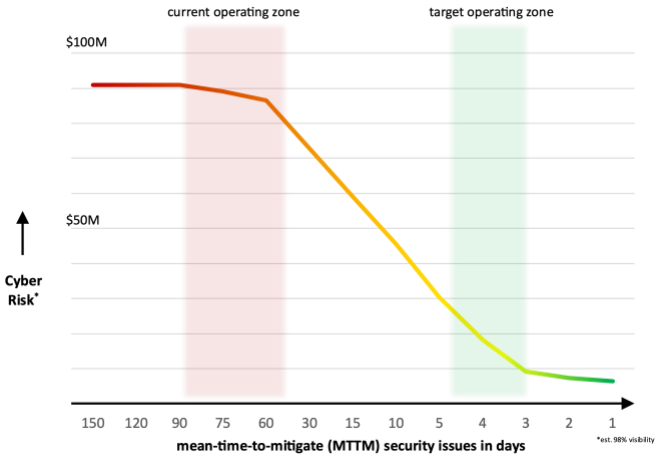

Figure 17: Cyber risk reduction with faster mitigation

Besides reduction in cyber risk, there are secondary benefits of cyber risk quantification with Balbix. Your infosec team has 100x more accurate visibility than before. Many cybersecurity use-cases, e.g., creating a cybersecurity update for executives or prioritizing CVEs take 10x less human effort. Balbix customers have saved millions of dollars each year.

Getting started

Getting started with Balbix is easy and quick (Figure 17). After deciding to proceed with Balbix, our customers typically take a week to identify the data sources and organize data snapshots for the initial deployment. Activation and initial data ingestion take less than an hour. It takes the Balbix brain a few days to understand and baseline your organization. In typically 7 days or less after we start, we can provide the first readout to operational and executive stakeholders, typically surfacing numerous previously unknown insights and risk issues showcasing capabilities for the use-cases discussed earlier.

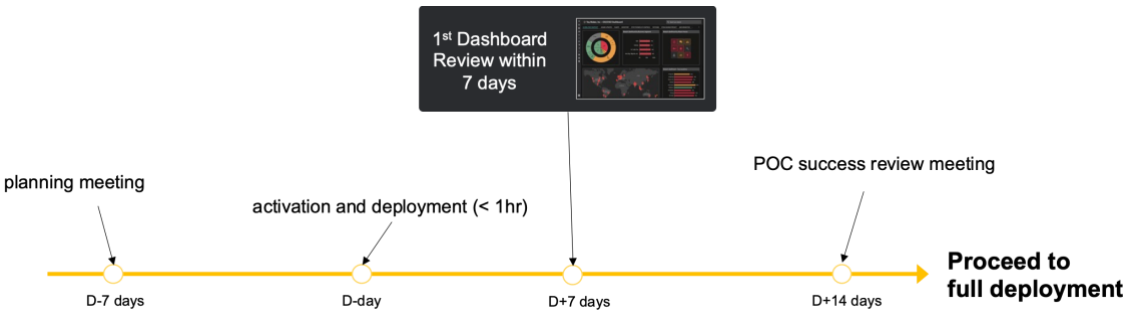

Figure 18: Time to value with Balbix

For more information contact: *info@balbix.com*.