Building Security In
Editors: John Steven, jsteven@cigital.com
Gunnar Peterson, gunnar@arctecgroup.net

Introduction to Identity Management Risk Metrics

GUNNAR PETERSON
Arctec Group

Good metrics aggregate both objective and quantitative measures, and should be consistent, cheap to gather, and expressed as numbers (see www.csoonline.com/read/070105/metrics.html). Typically, building metrics begins with collecting data and writing things down, reporting to management when you've got something with enough of a scientific base on which to stake your job, reputation, or more accurately, your organization's budget. As the numbers start coming in, you begin to see the tip of the iceberg, and eventually, you can build the total picture with the available metrics. Objective measurement will help you report on past progress, forecast future scenarios, and respond to real-time events.

In this installment of Building Security In, I present some *identity management* risk metrics that highlight the distribution, quality, affiliation, and governance of identity in a system as well as inform the decision-making process.

On metrics

Chief information officers (CIOs) typically grade their staffs' infrastructural efforts by collecting process metrics such as cost effectiveness, staff productivity, process efficiency, and cycle time (see www.apqc.org/portal/apqc/site/content?docid=112082).

Identity and security architects can use metric benchmarks to show the relative effectiveness of currently deployed security mechanisms and processes as well as to identify hot spots and areas for concern. Security decision makers often use metrics to drive prioritization, examine a deployment's feasibility, and support security architecture planning. Metrics also provide hard evidence of the existence, regular monitoring, and effectiveness of policy for either internal or external compliance audits.

On risk

Risk metrics quantify a system's assets, threats, countermeasures, and vulnerabilities. The combination of these elements ultimately yields a risk model in which you can manage, transfer, accept, or mitigate various risks. Risk differs from uncertainty in that it can be measured and managed whereas uncertainty can't. Risk management efforts hinge on this important distinction because it highlights the differences where a team could be more proactive. Many vulnerabilities are known, which means the security team can measure and manage them; however, the threats to a system contain a greater degree of uncertainty in that the threat environment contains numerous elements such as threat actors that the organization can't directly control. Vulnerability management processes consider such things as the cycle time required for vulnerability remediation, the amount and distribution of unremediated vulnerabilities, trend data, and so on. Measuring the likelihood of a threat successfully attacking a system is extremely difficult. Countermeasures can mitigate certain threats and vulnerabilities in defense of one or more asset, but more important, countermeasure metrics can help you focus on their presence.

Assets have value in a risk calculation: you can derive this value from their loss, misuse, disclosure, disruption, replacement value, or theft. For identity, all of these derivations could apply from the perspective of the identity provider, the relying party, or the user. Each type of risk metric yields a portion of the overall risk equation. Information security must combine the relative weightings of threats, vulnerabilities, assets, and countermeasures to successfully manage the risk the enterprise faces.

On identity management

At its simplest, identity provides the basis for access control decisions, and as such, the enterprise security architecture should reflect the quality of identity information on which it acts. Enterprise architects should report the rules, audit logs, filters, approvals, and delegations to which digital identities are subjected. These measures are chief informants to identity governance.

Provisioning systems (which create, edit, and delete accounts), virtual directories (which broker queries for identity data across disparate repositories), and metadirectories (which consolidate policy and management across identity systems) are rich sources of metrics because they typically contain critical metadata about identity definition, locale, and status. Identity can be a unique identifier in a

Table 1. Identity provider metrics. Parentheses indicate trend data over three months.

DIRECTORY	ACCOUNTS	GROUPS	AUTHENTICATION CLAIMS	AUTHORIZATION CLAIMS
IBM LDAP	2,156,968 (+2,401)	52 (–3)	1	4
Active directory	12,951 (+88)	2,047 (+17)	3	4
RACF	523 (+45)	54 (–1)	2	2
Federated identity store	81,203 (+502)	25 (+5)	3	3

virtual directory or metadirectory, a fact that drives several primitives:

- *identity repository size*—the number of accounts, groups, and other objects of interest clustered by the identity provider (individual, company, business unit, or certificate authority) or relying party (application, Web service, or network domain);
- *requesters/approvers*—the number of requesters and approvers for a given identity management process;
- *user account statistics*—the minimum, maximum, mean, and standard deviation of the number of accounts associated with a user across all users in a given domain, business unit, or location;
- *workflow branches*—the amount of serial and parallel branches in account management workflows (for example, the decision branches involved in adding a new user account to a payroll system);
- *authentication claims*—the number of claims for authentication purposes, such as user ID, password, and cryptographic keys, organized by account type and role;
- *authorization claims*—the number of claims for authorization accepted or rejected, subtotaled by role, user attribute, and group membership;
- *sensitive claims*—the number of claims whose release would injure the subject, identity provider, or relying party, subtotaled by account type and role;
- *potential and actual mapping points*—the number of logical locations that the workflow can provision based on its authority and how many places it's actually mapped to; and
- *provisioning geodesics*—the number of requests per approvals and the lengths a digital identity claim traverses to complete processes such as account creation, registration, and provisioning.

Relationships among the digital subject, the identity provider, and the relying party can be static or dynamic. Identity architects can use static and dynamic views of identity claims to show how identity services are used in the systems and their relationships to other services.

Locating the metrics

Many identity management metrics inform risk management; thus, identifying useful metrics demands collaboration among identity architects, directory architects, and identity management staff. To get pragmatic help, you should involve development and operations staff as well. Report the metrics that show development's use of identity services but also those that show the usefulness of such services to development. Adding instruments to a system with a metrics harness that provides objective measurements on process and systems inputs, outputs, and cycle times lets identity and security architects automate the metrics analysis steps and identity management operations staff can periodically monitor the quality of identity in the environment.

Examples

Let's look at some enterprise examples of identity metrics in three scenarios: measuring identity providers, provisioning process, and making claims about identities.

Identity provider metrics. Table 1 shows commonly used directory technologies' repository sizes, the number of claim types for authentication and authorization, and their trends over the previous quarter.

Identity and security architects can map identity provider metrics to the assets they're protecting, to compare, for example, the claims an identity provider makes on behalf of its subjects against asset value. The accounts and groups fields show the current data counts and a three-month trend. The parentheses indicate the trend, showing the risk profile for the assets over time. We can further delineate authentication and authorization claims to show how many claims are secret (such as passwords and cryptographic keys) and which are more public (such as social security numbers and account numbers). In this example, the system holding the largest amount of accounts, which could be presumed to also contain valuable assets, is protected by a single claim—for example, a password. The federated identity shows the account mappings for identities federated across the network, meaning the 8,203 accounts can be managed centrally, thus supporting multiple runtimes.

Provisioning process metrics. Table 2 shows the number of requesters and approvers required for certain identity provisioning processes to complete, the number of potential and actual mappings the process can access, and the geodesic length of all the processes' associated events. Cycle time calculates the average cycle time for the process to complete, but it can measure minimum

Table 2. Provisioning process metrics.

PROCESS	REQUESTERS	APPROVERS	POTENTIAL MAPPING POINTS	ACTUAL MAPPING POINTS	GEODESIC LENGTH	CYCLE TIME
Register user in system	4	4	16	7	11 nodes	4 hours, 12 minutes
Apply governance rule set	8	4	32	12	18 nodes	2 minutes
Terminate user	4	24	16	4	7 nodes	3 hours, 18 minutes

and maximum values as well. The apply governance rule set approver shows the additional rigor that this process provides at a relatively low cycle time.

Identity metrics. Table 3 shows the number of claims made for a set of roles and the geodesic length of the process to generate those claims. The claims required by the database administrator role show the strong form of identity representation that's commensurate with increased power.

Using risk metrics

Now that you've defined and gathered the base identity management risk metrics on a regular schedule and over time, you need a way to use them for decision support. Identity risk metrics can help in several ways:

- reporting (historical analysis, dashboards, forensics),
- predictive modeling (forecasting, scenarios, planning support), and
- real-time decision making (fraud detection, alerts, security events).

A metric's usage governs its collection, storage, and coverage. Incident response, for example, prefers large data sets for forensic analysis, whereas analysts working on a forecast prefer slices of specific data sets to help them identify likely outcomes and patterns. Similarly, fraud detection requires data slices, but in an online, interactive session with wire-speed execution time.

As a general-purpose decision-support tool, risk metrics provide a many-sided view of design options and risk management trade-off analyses: predictive models that forecast modes let risk management analysts differentiate tolerable failure from intolerable failure. This might be useful in scenarios in which a certain loss level is tolerated up to a given threshold, thus giving security architects and enterprise risk management staff a way to bound the loss of accounts.

Risk reporting

Identity and security architects can use reports to improve processes, measure organizational efficiency, and provide dashboards and scorecards for executives and business units. Reports help people understand their organization's identity landscape. Reporting metrics generally focus on

- *provisioning process efficiency*—the length of time from the initial request to the time the subject is provisioned or terminated in the identity repository;
- *provisioning process coverage*—the number of identity repositories and subjects the provisioning system governs versus repositories that rely on custom or ad hoc provisioning processes;
- *identity proofing strength*—the strength of secrets used for authentication;
- *identity assurance*—the amount of identity repositories that use proven, standards-based authentication and authorization protocols versus those that rely on custom-coded access control mechanisms; and
- *audit system usage*—the success and failure of authentication events, requests, authentication, and authorization.

These reports essentially let the analyst measure the system's state at a given time and date range as well as describe trend data to show the change trajectory over time.

Predictive modeling

A forecast's fidelity relies on both data quality and the modeler's domain knowledge and skill. The models themselves inform many different types of decisions.

Weighing the value of integrating to the identity source. In David Reed's law, which says the value of a network is derived from the many-to-many relationships of groups in a network (www.reed.com/Papers/GFN/reedslaw.html), security and identity architects look at the overall identity provider and relying party system as groups to determine the identity network's value. Federated identity lets organizations interoperate with their customers and partners without the burden of managing individual users; the cost–benefit analysis manages static individual users instead of the connectivity to a dynamic source of identity (as in federation). To make this sort of decision, you should compare the cost of iden-

Table 3. Identity metrics.

JOB FUNCTION	AUTHENTICATION CLAIMS	GEODESIC LENGTH	AUTHORIZATION CLAIMS	GEODESIC LENGTH
Database administrator role	11	18 nodes	7	15 nodes
Developer role	8	14 nodes	10	14 nodes
Business analyst	4	7 nodes	3	5 nodes

tity life-cycle management, including registration, provisioning, and termination, with the cost of managing a federation.

Estimating resource requirements for new systems. When new systems are under development, identity and security architects can use metrics to forecast the likely impact in cost, management, and technical resources needed to support the new system's identity architecture. This type of metric requires multiple types of data sources such as time and billing systems in addition to identity systems.

Identity architecture planning. The metrics for the assurance level an identity provider supports are very useful. If a consumer-facing system that relies on username and password combinations wants to strengthen its authentication process, for example, how much confidence should identity architects have in the consumer's desktop? What percentage of the consumer system's are infected with malware?

Deployment logistics. How are systems, users, groups, and management distributed in the system, and what standards and governance models are supported? These metrics help ensure that you take logistics into consideration when deploying new identity systems.

Asset protection. Given a set of assets, identity management metrics for the type of authentication used will show the level of assurance the asset can expect. An enterprise resource planning system might require two-factor authentication from the Web interface, for example, but allow password authentication from its rich client interface.

Identity incident response planning. Since February 2005, more than 84 million identity data records have been breached (www.privacyrights.org/ar/ChronDataBreaches.htm). For many organizations, an identity data breach is a matter of "when" not "if." Identity management metrics greatly inform how to plan the business impact, the cost to recover, and associated efforts such as notification.

Design for failure. Systems fail for different reasons and in different ways; identity management risk metrics help you analyze failure scenarios, such as the theft of a session versus the theft of a set of authentication claims.

Real-time decision making

Real-time decision-making systems rely on data that's small enough to be analyzed at wire speed. The areas of concern in such a system are fraud detection, diagnostics, security alerts, and security events. Security operations staff can configure the system to automatically respond to or propagate these events. From a metrics viewpoint, the data is useful for providing a current-state snapshot, as well as

- *validating access control requests*, in which authentication mechanisms use the claims about an identity and their source to assign a confidence level to an access control request;
- *detecting fraud*, whereby measuring behavior and events such as denied requests, the system can detect fraud attempts and cancel a user's session; and
- *determining access control availability*, in which metrics can measure standard statistics along with process the cycle time for an authentication request, from the initial request time until it's validated.

As with all risk metrics, understanding who will use the metrics and how is key to the metric definition and design process. Identity and security architects must understand the constraints in metric usage. The technical constraints, such as execution speed in real-time decision making, as well as human pattern recognition, such as with risk reporting, weigh heavily in the development of effective metrics.

The use of metrics is relatively new in security, but the progress toward quantitative assessment of identity systems and risk is positive. Security and risk metrics are gaining momentum as a way to understand and communicate risk across the enterprise, but many areas in software security architecture could benefit from a metrics-based approach. □

Acknowledgments

I thank Dan Blum, Kim Cameron, Dan Geer, Peter Lindstrom, and Elizabeth Nichols. Errors and omissions are my own.

***Gunnar Peterson** is a managing principal at Arctec Group. His research is focuses on distributed systems security. Contact him at gunnar@artecgroup.net.*