Limiting Insider Effects Through Micro-Zoning

Fred Cohen – CEO – Management Analytics Senior member, IEEE, Livermore, CA 94550 USA

Abstract—**This paper focuses on the use of micro-zoning to limit the effects of insider acts. Background is provided on components of micro-zoning and typical internal controls and their limits. A micro-zoning approach and its effect on insider effect reduction are explored, including effects on "why" and "ANA" issues. A set of insider threat properties are introduced, four example insider threats are identified, and the effects from those threats are discussed. Implementation strategies within three different sorts of environments are described, and conclusions presented.**

Keywords-component; insiders; micro-zones; virtualization; zoning architecture.

I. Introduction and Background

Malevolent or inadvertent acts by insiders may, in some cases, lead to potentially serious, grave, or catastrophic consequences. In cases where zoning architecture is in place and content controls are available, risk aggregation from insider acts may still reach levels where additional controls are desired.

Virtualization is becoming accessible to more enterprises and is increasingly easy to manage, deploy, and use. To the extent that visualization may realize protective value, that value typically lies in three key areas; (1) spatial separation, (2) temporal separation, and (3) controlled communication.

File sharing and representation methodologies (e.g., the fuse filesystem),[1] encrypted filesystems within files (e.g., the Linux loopback interface in encrypted mode)[2] and directories, and similar mechanisms have long offered the potential for finer grain than file structure area controls over content without high overhead or complexity normally associated with fine grained access controls.

Encrypted tunnels, virtual local area networks (VLANs) and other similar mechanisms also offer methods for better, finer grained, and easier to setup and tear-down separation of content in transmission.[3]

This paper focuses on the combination of these various separation and control approaches to form micro-zones and the potential effect and decision processes surrounding the use of micro-zone approaches in mitigating insider effects.

A. *Background on the insider threat*

When considering the "insider threat", a few basics must be defined. Insiders include anyone authorized beyond the authority of the general public. As such, there are typically many insiders, all of whom may be considered threats. We generally grant insiders authority based on trust, and in most cases historically, that trust is justified. Insiders acting commensurate with the trust placed in them by the organization granting access are "loyal". If and as insiders change loyalties while still retaining authority, we may call them disloyal.

While disloyal insiders may intentionally violate the trust placed in them, a far broader range of acts by insiders that cause harm are not intended to cause such harm. Examples include, without limit, errors and omissions, lack of policy leading to imperfect decisions, being too helpful or not helpful enough in hard-to-discern situations, and acting to get the job done despite impediments. Even the most obvious acts like browsing to Web sites, using email, or participating in a shared desktop activity in order to do your job can lead to the introduction of Trojan horses, allow remote control, or access content in undesired ways, that, with that insider assistance, gain access to and/or control of internal content and mechanisms.

B. *Background on typical internal control limitations*

Content control exists in many forms to limit the introduction of undesired and useless content and control content with business utility. However, these controls are often unable to effectively and timely differentiate desired from undesired content and are less effective when access is supported by the normal activities of insiders.

Zoning strategies applied at the enterprise level typically include zones (e.g., Internet, DMZ, Trusted, Restricted, Control, and Audit zones are in common use) for gross-level control and subzones used for risk disaggregation and functional separation within zones.[4] However, these controls have proven largely ineffective at preventing within subzone introduction and spread of malicious content and are particularly ineffective in the trusted zone where insiders have to mix access to external resources and collaboration with outsiders with their use of internal systems and mechanisms in order to do their jobs. Furthermore, at the infrastructure level, where zones and subzones are typically defined, aggregated risks are still often quite substantial.

While other controls, like auditing, background checks, better supervision, and so forth are available, history does not support their effectiveness against insiders, especially in prevention. Accurate attribution is fundamental to insider detection and proper disposition, and yet this is problematic in many insider cases, even when no subversion is attempted.

C. *Background on micro-zoning support technologies*

We hypothesize that a micro-zoning strategy using increasingly lightweight mechanisms may have substantially more benefit than cost in limiting the effects of insider acts, whether from loyal or disloyal insiders, and regardless of their intent. But this depends largely on the availability and ease of deployment and use of these technologies.

1) *Virtual machines and their properties*

Virtualization is becoming accessible to more enterprises and is increasingly easy to manage, deploy, and use. Well-working products are freely downloadable and relatively low cost per system, there is price substantial competition in the market from several large and stable competitors, protection from programs within virtual machines (VMs) altering the external environment through other than authorized paths is effectively limited to covert channel effects, and there are extremely lightweight versions of similar capabilities in place and increasingly coming available for wider uses. As a component of a micro-zoning strategy, this is key, because it enables rapid deployment of VMs on servers and, endpoints, and lightweight versions are usable from mobile devices. The protective value offered by VMs stems largely from their support for controlled separation and communication.

Separation comes in two ways. Spatial separation stems from the fact that the VM can support independent internal storage, either in memory, on disk, or over infrastructure to servers, so that VMs don't communicate with each other except through controlled interfaces. This has proven more effective than operating system protection and almost as effective as hardware protection, except for access initiated and performed by the enveloping operating environment. Temporal separation is also offered in that VMs need only operate while in use, and can be shut down and restarted at will. This typically includes the option of retaining or not retaining changes, and without such retention, the only effects other than covert channel effects of the VM on other VMs or other aspects of the environment stem from the controlled communications. Thus a Trojan horse or virus in the VM cannot infect other VMs, and in cases where the VM does not retain state past reboots, the effects will be limited in time.

Controlled communication in most VMs today come in one of two forms. Communications through network connections and other peripherals are controlled just as in any other environment. But in addition, there is usually a way to share files with the host operating environment. This file sharing provides the means to support controlled movement of stored content between VMs and other areas in a very clean and direct way, and using the methods already known and in use by the users in their normal operating environments. For example, in a shared desktop environment with external workers, a file needed for the session can be readily loaded into the VM environment via drag-and-drop, and a file to be moved from the VM into the external environment can be moved back in the same way. The key is that this is controlled only from the external operating environment and cannot be done by any mechanism within the VM itself. Thus an explicit act must be undertaken in order to communicate stored content to be retained or made available elsewhere. A Trojan horse or virus in the VM cannot move itself outside of the shared file area into the enclosing environment, it can only make itself available for such a move. Of course this does not eliminate the mechanism, it just makes it require explicit action by the insider.

A VM provides a complete operating environment that can be initiated in a on the order of 30 seconds on an end-user system (e.g., a laptop) and faster on servers. As a result, for tasks that take from minutes to hours, the overhead associated with operation is relatively low.

2) File sharing and more granular controls

File sharing and representation methodologies (e.g., the fuse filesystem), provide the potential for use in a wide variety of applications, including the temporary provisioning and tear-down of small (micro-) file systems on short notice. An example of this is a remote mount of a filesystem within a user's directory (e.g., a subdirectory associated with a project mounted as the root of a remotely accessed filesystem) or the creation of a filesystem representation of a set of records from within a database based on a query (e.g., using the fuse filesystem to treat the results of the query as a filesystem). Such increased granularity of access for short durations allow for the creation of micro-filesystems associated with VMs for the duration of use so as to isolate the VM from accessing content outside of the micro-filesystem while still facilitating ready access to the content area(s) needed for the particular use.

Another approach to this is to use encrypted filesystems within files (e.g., the Linux loopback interface in encrypted mode). In this approach, an independent and dynamically growable filesystem can be used for specific uses, with access limited based on the quality of the encryption, to particular sessions, VMs, or other uses. Encrypted directories have also been in use for many years, but the means to readily access them from within a VM without exposing them elsewhere has historically not been available. Another alternative is areas created on file servers on a case-by-case basis, such as the creation of an identity to a particular purpose on a cloud file resource, and its use for the purpose and subsequent destruction, but this technology is not yet well suited for rapid setup and teardown or remote mounting in VMs with minimal effort. Other similar mechanisms have long offered the potential for finer grain than file structure area controls over content without high overhead or complexity normally associated with fine grained access controls.

Regardless of the specific benefits, risks, and costs of each of the available methods, there are clearly many such methods that could be used for micro-zone storage.

3) Encrypted tunnels, VLANs, and similar mechanisms

Encrypted tunnels are an example of a very widely used technology that is already applied to rapid setup and teardown at the session level, to the point where it is so standard, it is not readily disabled in many applications. Such mechanisms are readily available for applications like remote file system mounts (e.g., sftp mounted as a fuse filesystem), remote desktop applications (e.g., essentially all current such applications encrypt by default with pseudo-randomly generated session keys), and Web site access (i.e., SSL). For access beyond secured local areas, these are in common use, and require little or no configuration by the end user.

Another approach to separation at a fine granularity is VLAN mechanisms now widely available and used for zone and subzone separation. VLAN mechanisms are rarely used for micro-zone applications or anything of that fine a granularity today because the technology for provisioning and teardown is not yet rapid and lightweight enough to make it practical, and because the number of available VLANs within most routing and switching infrastructure is several orders of magnitude too small to use for this purpose. Such a technology implemented in such infrastructure might someday be an alternative to relying on encryption, or may be used in some high surety systems today for mer certainty.

The capability is now readily available for separation of micro-zones communicating outside of machines and between multiple VMs and associated storage areas.

II. A Micro-zoning Approach

A currently available micro-zoning approach involves the combination of these three technologies (VMs, encrypted tunnels, and micro-filesystems) in a structured way so as to reduce undesired effects of insider behavior on enterprises without substantially increasing the operational costs to those enterprises for the protective value afforded or reducing the performance of the insiders in their legitimate work.

A. The value of micro-zoning to insider effect reduction

Regardless of any other value brought by micro-zoning, the key value of interest in this paper is the reduction of aggregated risks from insiders. This aggregation of risks stems fundamentally from the fact that insiders need access to many mechanisms and much content in their overall work efforts, and granting such access under fine grained control is expensive and complex to the point where it is infeasible today and likely too expensive for the foreseeable future. This relates largely to the notion that “any is not all” [5] and the lack of a current capacity to limit access based on “why” as opposed to “who”, “what”, “where”, “how”, and “when”.

1) Any is not all (ANA)

As an example of the any vs. all issue, consider a systems administrator whose job it is to deal with any problem with any particular system at any particular time. That implies that they be able to do any of a large set of tasks that require a wide array of capabilities and access. However, at any given time and under any given situation, the actual access required is typically far less.

Another example is the chief marketing officer in a large enterprise who potentially must be able to access any information about any account, but who in practice, never needs access to all information about all accounts.

Each of these any any number of other examples lead to the problem of limiting access to those who must be able to access anything. Various efforts have been made and techniques developed for limiting volume over time, controlling aggregated access, and so forth, but none effectively solve the underlying problem that the underlying protective mechanisms are not designed to provide for the inherent conflict of access to any without access to all.

2) The problem of “why”

There is a closely related issue of the reason behind the use being linked to the access associated with the use. The fundamental problem is that computers are not designed to provide protection based on the reasons behind things, but rather the things that are more directly measurable and controllable by computers.

For example, a systems administrator may need to be able to read all files in a filesystem in order to do a backup and to write all files in a filesystem in order to do a restore, but if we try to control mechanisms to the point where each of the potentially unlimited activities of a systems administrator are explicitly controlled, we may soon find that the time and effort spent in managing the process far exceeds the time and effort associated with systems administration. Use of automation in administration has largely reduced the number of administrators required for large environments, but as a side effect, it has also dramatically increased the aggregated risk associated with each of them. Thus we have used automation to increase aggregated risk and now seek to reduce that aggregation by making them less efficient, which may ultimately lead to higher cost and less efficiency, the very reason for the automation in the first place.

A sales representative is typically limited to a limited number of accounts they deal with, and in some cases, the process controls associated with sales work flows are highly restrictive with respect to what a sales person can do as a function of the place in the sales process they are currently at with any given prospect. And yet this efficiency does not prevent the sales manager from being able to reallocate sales people to prospects, because if it did, the sales manager couldn't get the sales force to work on prospects when a sales person was ill or after they retired. Different reasons may even be provided to the system in the process of authorizing such activities, and such reasons are commonly misused in order to get the job done, when the specific reason is not on the list, or when the process is too much of an impediment to performing the job.

So we see that the “why” behind actions, even when highly controlled, produces a tradeoff between efficiency and effectiveness, and controlling it beyond some threshold becomes too complex and burdensome for effective and efficient operations. Furthermore, these examples are from highly constrained and mature examples of such processes, and many enterprises don't have that level of maturity over all of their functions.

3) Different sorts of insider behavior and effects

Different insiders have different access, motivation, skill, and jobs to do. For the malicious insider with unlimited access, high levels of motivation to act maliciously, great skill, and a high level of trust associated with their job, the challenge of limiting effects is likely to be harder to meet. But the vast majority of insiders are workers with limited access, a desire only to do their job, limited skill in relevant areas other than their job, and the normal level of trust of the average worker. The following characterization shows what effects might reasonably be expected based on these areas:

Access: More access implies that the aggregated effect is greater, approximately linearly with the quantity of content accessible.

Motivation: Benevolent insiders may be more or less diligent while malicious insiders are by a range of specifics that may make them directed or opportunistic in their acts.

Skill: The skill level of workers in terms of the ability to use and abuse protective and operational mechanisms range from knowingly acting only in ways they were trained to supported by teams of outside experts.

Trust: The actual trust placed in insiders varies from the amount of trust at the granularity of protection to the total loss of enterprise function for days to weeks.

If we consider the effects experienced from different insiders based on these properties, and the realities of a well-zoned enterprise with sound controls, the spectrum runs from collapse of several subzones for the least trusted minimal skill benevolent and reasonably diligent insider with limited access to enterprise collapse from the most trusted insider with unlimited ANA access to critical components in the interdependency chain, highly motivated, and with outsider expert assistance.

We will assume that the goal of micro-zoning is to reduce as much of this aggregated risk as it can without going to extremes that substantially disrupt enterprise operations and consider its effects across the insider spectrum as we understand the architecture and what it can and cannot do.

B. How micro-zoning disaggregates risks

Micro-zoning changes the fundamental risk aggregation at the level of granularity applied in time and space. It does this largely by separation in time and space. To the extent that during various relatively higher risk operations, micro-zones are in use, they can reduce the associated risk to the extent that separation can do so.

Various attempts have been made to characterize the effects of vulnerabilities, attack mechanisms, and so forth with regard to information technology, but none are really designed for the issues at hand. Rather than seeking to reinvent the wheel here, we will characterize risk in terms of insiders use and/or movement of inbound or outbound structured or unstructured data and/or executable content subject to effects on integrity, availability, confidentiality, use control, and accountability.

1) Completely unconstrained environments

In a completely unconstrained environment, insiders may use anything in the environment for any purpose, and move anything in any volume from anywhere to anywhere else within the enterprise and or to and from accessible areas of the Internet. No controls over data are in place, so anything can be put anywhere, and thus any mechanisms for control of integrity are largely ineffective, availability can be readily disrupted by any insider at any time, confidentiality is violable by any insider at any time and in any amount, there is no control over use, and no accounting is reliable enough to attribute actions to actors. These are only the direct effects of actions by insiders and the systems they use. Indirect effects may spread to the transitive closure of information flows, but without any controls this hardly makes a difference.

2) Common zoned enterprise environments today

In the most common case we encounter today, an insider using an enterprise endpoint is logged into their user account and granted all access relevant to their job function. If and to the extent they use other systems, such use may be undertaken by custom or customized interfaces to those systems (e.g., Web interfaces to transaction systems, Mail user interfaces to internal email systems, etc.) and for use of tools (e.g., document processors, spreadsheets, and code development tools) they run programs on their endpoint. If access to enterprise storage areas is required, they typically either operate through a document management systems, or use remotely mounted file systems logged in as a user or member of a group. Their overall effect is limited by zoning and subzoning to the extent this is well controlled, and their immediate effect is limited by their logged in account(s) and associated use privileges and access. Their workstation operates at least from the time they arrive in the office to the time they go home, with screen-lock or logout during extended periods of non-use.

Some resultant effects on the areas identified include:

Use: Any programs on the insider's computer are usable for any use they may be put to. This includes the ability to load more programs and use them. Use of enterprise controlled systems are limited to the uses authorized for the insider's use identity over time.

Movement: Any accessible content can be copied to any accessible and writeable area, all constrained by zones and subzones and potentially blocked by other content controls.

Inbound: Any inbound content can be stored and moved to any writeable area.

Outbound: Any readable content can be sent out of the zone and subzone to the Internet.

Structured or unstructured data: Any data can be moved within the zone and subzone and/or to connected external locations unless restricted by content controls.

Executable content: Execution has scope of direct effect limited by zone and subzone.

Integrity effects: Arbitrary corruption can extend to all accessible areas of the zone and subzone not under additional controls.

Availability effects: Executable content can directly deny services or destroy writeable content.

Confidentiality effects: Any accessible confidential data can be leaked in any volume unless limited by content controls.

Use control effects: Only zone and subzone accessible mechanisms and content can be used and use by any executable mechanism can potentially invoke any such use.

Accountability effects: Acts undertaken by the insider's computer are often attributed to the logged in user identity, including acts by software within that insider's computer.

All of these capabilities are typically present when the insider's computer is active, which is most of the time during most work days, and in some cases, nearly 24x7.

3) Micro-zoned environments

In micro-zoned environments, activities like shared remote desktop sessions, Web-based research activities, email, and similar tasks that interact with and potentially move information to and from external areas, are undertaken from within micro-zones. In essence, a VM is initiated with a standard set of software tools and access to micro-zoned file system shared areas or with upload and download capability. They operate within and augment zoned architecture extant.

The net user experience effect of this micro-zoning is an increase in startup and shutdown time for these activities and some overhead associated with virtualization, encryption, and the definition of micro-zone file areas.

The net effect on risk aggregation is that the direct use risks associated with these applications is reduced to the accessible areas over the timeframe of the activity, and indirect use risks are limited to side-effects of insider acts that move content between micro-zones and their enveloping zoned environments. There is also a potential for some increase in risks associated with the presence of the micro-zoning mechanisms themselves, but these mechanisms have not displayed substantial risks in this area over their history to date, and the mechanisms used for isolation of VMs, encrypted tunnels, and micro-filesystem areas (1) do not reduce the effectiveness of the other existing controls in the environments they operate in and (2) tend to be small and well controlled mechanisms that, in some cases go through extensive proofs related to security and separation to a level far in excess of what is done for their enveloping environments.

An example of micro-zoning in day-to-day operations for remote desktop use is applied as part of standard process for a consulting operation. In this operation, shared desktop sessions of one or more hours are commonplace. These sessions typically involve a presenter running a desktop with a number of participants who may participate telephonically or, in some cases, may directly work with one or more applications on the desktop. Files are loaded to and from the environment, programs run within the environment, and so forth.

In a non-micro-zoned environment, this potentially exposes the entire accessible area of the insider's operating environment to the direct and indirect effects of the activities during the session, and any residual effects on that environment left by the session.

In a micro-zoned environment, the session is started within a VM initiated for the purpose and configured to support such activities. The VM has access to internal zone and subzone network segments only in cases where cryptographic keys are made available to the VM. This is used for cases where micro-filesystem areas are accessed through encrypted tunnels. In these cases, the remote filesystem mount is undertaken by mounting only the area of the filesystem associated with the meeting. For example, in a meeting used to review documents related to a project, the remote mount will typically be restricted to a directory where information related to the particular project is kept, thus preventing access to all other projects potentially accessible to the insider. The VM has access to a shared file system area that is also directly accessible from the insider's desktop. At the end of the session, the VM is shut down, not retaining any internal state. The only retained content is in the shared filesystem area or, in the case of VPN connections to micro-filesystem areas, stored content in those areas.

For the purposes of understanding, let's assume that, operationally, insiders are trained to use the mechanisms and that the often used mechanisms are available as an icon on the desktop. For example, there is a commercial application called TeamViewer that performs remote desktop operations, and an icon on the insider's desktop labeled TeamViewer is invoked by the user to start the process. This icon causes a VM to start up using a configured version of an operating system (say Linux) that, at startup, returns to a retained state logged in as the insider, with the TeamViewer application just started and ready to operate, and a shared file system with the user's desktop mounted. Let's assume that access to a repository is needed for this session, and that the user then uses an icon to mount the remote micro-filesystem area, specifying the directory to be mounted. The system mounts this directory using a secure shell tunnel, and the system is ready to use for the meeting. If the insider wishes to move content in and out of the meeting micro-zone, it is done by using the shared desktop directory, and for storage within the meeting to be kept after the meeting, the micro-filesystem is used. At the end of the meeting, the insider closes the VM, which doesn't retain any information internally, and the session ends.

Some resultant effects on the areas identified include:

Use: Only programs on the micro-zoned VM are usable for any use they may be put to during the period of VM operation. This includes the ability to load more programs and use them, but only for the duration of the VM session, since the VM does not retain internal state information across sessions. Use of enterprise controlled systems are limited to the uses authorized for the insider's use identity during the period of operation of the VM.

Movement: Any micro-zone accessible content or areas can be copied to or from, all still constrained by zones and subzones, but far more tightly constrained by the micro-zone, and potentially blocked by other content controls.

Inbound: Any inbound content is constrained to the micro-zone unless and until it is moved from the micro-zone by the insider's explicit acts. It does not persist past the session.

Outbound: Any readable content within the micro-zone can be sent out to connected Internet locations only during the session.

Structured or unstructured data: Any data can be moved within the micro-zone and/or to connected external locations unless restricted by content controls, but only during the session.

Executable content: Execution has scope of direct effects limited by the micro-zone. That is, it can execute within the VM and access accessible areas of the micro-zone, but only micro-zone storage areas not within the VM will retain side-effects.

Integrity effects: Arbitrary corruption can extend to all accessible areas of the micro-zone not under additional controls and remain there for the duration of the session, except for micro-zone storage areas outside of the VM, which will retain effects.

Availability effects: Executable content can directly deny services within the VM and to the extent the VM effects its enveloping operating environment, within that operating environment, as well as through use of network bandwidth or direct effects on network-accessible areas, all during the session only. It can also potentially destroy retained writeable content within the micro-zone.

Confidentiality effects: Micro-zone accessible confidential data can be leaked in any volume unless limited by content controls, over the period of the VM session.

Use control effects: Only micro-zone accessible mechanisms and content can be used and use by any executable mechanism in the VM can potentially invoke any such use.

Accountability effects: Acts undertaken by the insider's VM are often attributed to the logged in user identity, including acts by software within that insider's VM. Attribution to the MV and session are also typically reflected in those logs.

All of these capabilities are typically present only during the insider's invocation of the micro-zone, which is typically limited to the functional period of use. With the ability to pause and continue operations, use and all of the micro-zone VM capabilities can be suspended during periods of non-use.

This reduction in potential effects is substantial. Instead of long-term residual effects on all content in the insider's endpoint, limited time effects on a small subset of content are the typical worst-case result.

The cost of such micro-zoning in a typical use case is quite low. For a 1-hour remote desktop session, startup of the microzone takes place during the same time used to start the remote desktop application and get the meeting started. It is unclear that this takes any additional time, since meeting startup times tend to be dominated by the wait-time associated with members joining the meeting. File movement and other similar activities are slightly slower in some cases, but this is not noticeable relative to delay times associated with communications between distant desktops. If and to the extent these meetings are commonplace, the operation and configuration of the micro-zone is standard and may be retained as part of the VM startup environment.

4) Limitations on ANA and Why

ANA effects are limited by micro-zones in that the “any” access granted is only a subset of the “all” access potentially available to the insider. The micro-zone effectively limits the scope of “all” within the micro-zone to the set identified for the use rather than the set for the user. While no direct “why” analysis is done by technical mechanisms, the decision about what constitutes the micro-zone for the use is essentially a decision about what based on why. Thus the micro-zoning strategy and application has the potential to realize substantial reductions in the ANA problem and links indirectly to addressing the issues of why.

C. The effects of micro-zoning on different insiders

As we discussed earlier, different insiders act with different access, motivation, skill, trust, and jobs to do. The limitation of effects of insiders depends on these properties as well as the way micro-zoning is used.

We will focus on 4 insider models for the purposes of this discussion; (1) the competent but imperfect **loyal sole contributor** with normal user access for their job function, just doing their job on a day-to-day basis; (2) the **disloyal executive** or manager who has decided to open a competitive company using information from the enterprise; (3) the competent but imperfect **loyal administrative functionary** who has a wide range of responsibilities and tasks using a relatively small number of applications. They attend many meetings, support a team of workers, authorizes others to perform activities, and reviews performance to support the enterprise objectives; and (4) the **disloyal systems and network administrator** who was caught running a business-in-a-business using enterprise resources but was retained for perceived operational importance, and decided to do a better job of hiding the business-in-a-business within the enterprise.

From these insider models, we will assume some notions around access, motivation, skill, and trust levels, and discuss the effects of micro-zoning on limiting their effects.

1) Loyal sole contributor

Access: Limited to projects they work on and common content available for their job.

Motivation: Benevolent, just trying to do their job.

Skill: Unskilled technical attacker, but computer savvy.

Trust: Trusted to perform their job reasonably as trained.

Effects are expected to be:

Use: Only authorized use by the insider is likely, and methods used to take advantage of their insider access during collaborative or other use of external system and resources will be limited to the period of VM use.

Movement: They will only move content as appropriate, and methods used to take advantage of their insider access during collaborative or other use of external system and resources will be limited to the period of VM use and the directly accessible content. Trojans might possibly make their way out of the VM through the shared area, but this will be a slow process if it is effective at all.

Inbound: Any inbound content will be constrained to the micro-zone during its period of use and only content required for job duties will be moved from the micro-zone by the insider's explicit acts.

Outbound: Only information made accessible by the insider for the purpose of the task will be made available to be sent out and only during the session.

Structured or unstructured data: Structured data received is likely to only be used by authorized mechanisms which affect content controls. Unstructured and structured data is likely to be inspected by the insider during use, and this is less likely to have indirect consequences. Only insider needed data will be stored past the session.

Executable content: Execution will only have scope and effect within the VM during the session. The insider will not move executable content out of the VM unless it is concealed in structured or unstructured data, as described above. Side effects to stored data within the micro-zone are possible, but likely limited to direct effects of mechanisms on that data during use. Very targeted and skilled threats would have to be at play in order to cause the insider's access to be exploited for larger scale indirect effects.

Integrity effects: Arbitrary corruption can extend to all accessible areas of the micro-zone not under additional controls and remain there for the duration of the session, except for micro-zone storage areas outside of the VM, which will retain effects. Insider inspection for use and use only by limited applications are likely to limit any such effects.

Availability effects: Executable content can directly deny services within the VM, but if the insider restarts the session, the executable is stopped until again invoked, and is not retained across sessions. It must be again executed by acts of the insider. To the extent the VM effects its enveloping operating environment, it may reduce performance of that environment, but VMs can be throttled in various ways (e.g., CPU usage, memory, disk space, network connectivity, bandwidth) within that operating environment, and again, a restart stops this until insider acts again induce the behavior. It can also potentially destroy retained writeable content within the micro-zone. For the well-intentioned insider, these effects, when noticed, will likely be countered and reported, and they will not normally persist across sessions.

Confidentiality effects: The insider must make confidential data accessible to the micro-zone in order for it to be leaked. This could only be done (1) in quantity up to the size of the micro-zone stored content, which would be limited by the insider actively uploading or enabling its use, (2) when working with trusted partners the insider deemed suitable for such access, and (3) during sessions with those partners where that information was made available. In any case, such leakage and the capability supporting it would not persist past the end of each session.

Use control effects: Only micro-zone accessible mechanisms and content can be used and use by any executable mechanism or person in the VM can potentially invoke any such use. However, the insider limits the available uses by selection of the VM and by authenticating uses where such authentication is required. Thus use is limited to sessions where the insider authorized the specific use and only continue over the remainder of the session or insider-authorized period.

Accountability effects: Acts undertaken by the insider's VM are often attributed to the logged in user identity, including acts by software within that insider's VM. Attribution to the VM and session are also typically reflected in those logs. While external subversion of the internal VM logs may be feasible, subversion of the external logs by the VM are outside of the capacity of the micro-zone, as are subversion of other aspects of the micro-zone accountability features. Also, any such subversion is limited in time to the part of the session after the subversive act.

2) Disloyal executive

Access: Wide access to information at all levels and across many areas, and the ability to cause many others to act to grant more access as desired.

Motivation: Greed motivates them to do subversive and deceptive acts to gain access and exfiltrate large amounts of confidential data from trusted and/or restricted zones.

Skill: Unskilled technical attacker aware of protections in place and able to subvert controls by power and influence.

Trust: Trusted to make and carry out executive decisions and directly in charge of many workers and managers. Feared by many workers and thus obeyed even if not trusted.

Effects when the insider is not acting with malice are expected to be similar to those of the loyal sole contributor. Thus effects are reduced even for this case. However, when acting with malice, they are expected to have effects like:

Use: Use in excess of loyal behavior is likely, and increased use is expected to be produced over time by executive action. Use of resources will be limited to periods of VM use, but that use may be extended for long periods by the insider to allow increased use, and use may be facilitated for outsiders by the insider through sharing.

Movement: They may move content they wish into the micro-zone and have it extracted by connected parties, but that content will only be available during the periods of use. To the extent some classes of VMs are not normally authorized to some areas of content, this may again be subverted by power and influence. Various controls may be used to force presence by the insider, but this will only deter, and may move the insider into other methods.

Inbound: Any inbound content will be constrained to the micro-zone during its period of use, but the insider wishing to move it into other areas will do so by explicit acts.

Outbound: Any information made accessible by the insider will be made available to be sent out, but only during sessions when it is enabled. This will require explicit action by the insider for each session to make that content available.

Structured or unstructured data: Any data received may be exploited by explicit action of the insider and may be stored past the end of the session by explicit insider acts.

Executable content: Execution of content not desired for exploitation by the insider will be limited as for the loyal insider. The insider may move executable subversive content out of the VM for explicit exploitation and execute it.

Integrity effects: Insider acts may cause corruption of content and mechanisms outside the micro-zone if and when the insider moves content from the micro-zone to other areas for exploitation purposes.

Availability effects: The large-scale exfiltration of content may cause performance effects if large enough in volume, but otherwise, the insider will not intentionally cause such effects.

Confidentiality effects: The insider likely will make confidential data accessible to the micro-zone in order for it to be leaked. To the extent this is limited in quantity, the insider may use power and influence to change these thresholds.

Use control effects: Only micro-zone accessible mechanisms and content can be used. To the extent this is limiting, the insider may use power and influence to increase the scope of use. The insider may make their credentials available for such use, or may use power and influence to create other identities that they then use for such purposes.

Accountability effects: Acts undertaken by the insider's VM are likely to be attributed to the logged in user identity. Administrative subversion may be used by the insider to deceive such attribution. Even if detectable by controls, the executive may be able to continue unhindered by subverting the control system through power and influence.

3) *Loyal administrative functionary*

Access: Access to management and detailed information associated with their areas and workers and early access to information on business unit planning efforts. Able to alter access of others within the bounds of their control.

Motivation: Benevolent, just trying to do their job.

Skill: Unskilled technical attacker, but may be computer savvy, and can get help as needed.

Trust: Trusted to perform their job reasonably as trained.

Effects are expected to be effectively the same as the effects for the loyal sole contributor, except that the scope of content and capabilities may differ depending on the uses and information made available within the micro-zone.

4) *Disloyal systems and network administrator*

Access: Ancillary access exists to large portions of enterprise content and direct access to systems and mechanisms used for in-zone protective functions. Not directly able to effect audits and limited in high-consequence actions.

Motivation: Greed and ego, trying to make more money on the side and feeling like they deserve more credit and renumeration. They often feel that managers are fools.

Skill: Skilled technical attacker with access to external expertise and tools.

Trust: Trusted to perform their job reasonably as trained and highly trusted to keep systems and networks operating within their zone.

Effects when the insider is not acting with malice are expected to be similar to those of the loyal sole contributor. Thus effects are reduced even for this case. When acting with malice, they are expected to have effects like:

Use: Use in excess of loyal behavior is likely, and increased use is expected to be produced over time by technical manipulations. Subversions undertaken to conceal use of resources may cause errors or misoperation of enterprise zoning controls, making micro-zoning protections more important because of their independence from enterprise zoning. To the extent the zoning limits exchanges of information not involving known micro-zones, this may be problematic for the insider wishing to operate a 24x7 business-in-a-business.

Movement: They may move content they wish as they wish, within the limits of zones and subzones, but this is unlikely to effect micro-zones.

Inbound: Inbound content will be dealt with as they wish, subject to the ability to gain access to it.

Outbound: Any information made accessible by the insider will be made available to be sent out.

Structured or unstructured data: Data received may be exploited by explicit action of the insider.

Executable content: They may use executable subversive content as they wish.

Integrity effects: Insider acts may cause corruption of content and mechanisms.

Availability effects: The business-in-a-business may cause performance effects if large enough in volume, but otherwise, the insider will not intentionally cause such effects.

Confidentiality effects: This insider is unlikely to produce confidentiality effects to the enterprise..

Use control effects: To the extent the insider is forced to use authorized micro-zones to pass information in and out of the enterprise, this may be problematic. The insider may create or alter micro-zones to allow their use, but if use controls limit operation of these micro-zones based on ongoing authentication or otherwise limit them, the insider may no longer be able to use enterprise resources for their business-in-a-business 24x7 operations.

Accountability effects: Acts undertaken by the insider's VM are likely to be attributed to the logged in user identity. Administrative subversion may be used by the insider to deceive such attribution. To the extent some of the movements they make are more noticeable because they are not controlled by micro-zones, this may be more readily detected.

III. Implementation and use of micro-zones

Implementation of micro-zones of the sorts described may take many forms. Some of the more common design patterns are outlined here as an example, with limited case examples used to discuss observed effects.

The basic set of choices with respect to micro-zones are to use temporary {{encrypted} remote access connections to / on-endpoint} {non-}state-retaining {terminal servers, microzones} {with controlled configurations, surveillance, recording, limited actions, {with push / pull / shared storage}} for remote {diagnosis, maintenance, supervised activities} for limited time frames.

Temporary micro-zone use is necessary in order to gain the temporal limitations on effects. Within endpoints, we generally use shared storage areas (e.g., a directory on the desktop of the user for movement to and from the micro-zone) and encrypted remote communications to micro-zone storage associated with a task or client in non-state retaining VMs with controlled configurations for supervised activities for limited time frames. For terminal server-based use (e.g., Citrix), we use encrypted remote desktop connections to non-state-retaining VMs with controlled configurations, surveillance, recording, limited actions, and push-pull storage for supervised activities for limited time frames. We tend to use a different operating system in VMs than in the enclosing operating environment out of the belief that different environments are better suited to the different tasks at hand.

Several examples of use patterns are in limited common use today. Three of them are; (1) small business endpoint risk reduction; (2) limited joint venture participation, and (3) enterprise access to key infrastructure.

A. *Small business endpoint risk reduction*

Increasingly, knowledgeable small businesses, (e.g., consulting boutiques, small engineering firms, etc.) are getting concerned about the aggregation of content and risk on user laptops. The need to have information available leads to risk aggregation on these devices, and disk encryption offers a partial solution, but fails to protect content in use. The requirement for remote desktop sessions, Web-based research and support functions, special purpose collaborative software packages, and similar things exposes them, in their view, to increased risk of machine takeover and exploitation of all available content, leading to loss of intellectual property and reputational damage. A solution implementable at low cost by the users themselves is micro-zones.

For perceived high risk activities, like Web browsing, remote desktop activities, and collaboration using software not under their control and loaded from unknown or untrusted sites, shared storage desktop directories for non-state retaining VMs for supervised activities over limited time frames are configured. They are started up as needed, run for the period of the activity, with retained or use-required content moved through the shared directory. This is believed to afford sufficient risk reduction to allow far safer operation, and user experience shows that it eliminates essentially all malware, viruses, etc. that are brought in by these loyal insiders from persisting past the end of session or having effects outside of direct effects on the micro-zone while it is active. These mechanisms are normally operated through virtual network address translation (NAT) gateways, eliminating direct attacks against the machines and leaving only Trojan horses and similar attack methods that exploit actual use. Out-of-pocket costs are nearly zero, and the time taken at start and end of session are negligible in the context of this sort of use. Administrative costs are very low, since once configured, the VMs are not typically updated, even for patches, more than quarterly. Even for small shops, VM images are built once and made available for others to reduce these costs.

The net effect is that, for practical purposes, loyal insiders no longer have to worry significantly about accidentally bringing in malware, can use the Web and other similar resources without concern about putting their entire environment at risk, and in practice, incidents involving more than micro-zones are rare.

B. *Limited joint venture participation*

For limited joint ventures (JVs), particularly where employees of several companies that are otherwise competitive are participating part-time in a collaborative mode, the mix of the need to share selectively and the need to prevent leakage or granting of access or privileges to each others' infrastructure lead to specialized collaborative environments. Without such environments, the risks of exposing internal information to competitors is perceived as being too high to bear, and a wide range of specialized solutions have been tried over time.

In these situations, many enterprises have created micro-zones as a sort of "shared space" for the joint venture. One or more of the participants host computers in zones and subzones isolated from the rest of their enterprise information infrastructure, and allow terminal server-based use (e.g., Citrix), with encrypted remote desktop connections to non-state-retaining VMs with controlled configurations, surveillance, recording, limited actions, and push-pull storage to and from the JV with shared mountable storage by the VMs in the JV zone and subzone for supervised activities for limited time frames.

The net effect is that, JV insiders are only able to access the micro-zone data associated with their tasks. They may act in benevolent or malicious ways, and to the extent they act maliciously, they may do harm to the JV or the relevant content within the JV they are allowed to access, but they are restricted from accessing content in some micro-zones, and they have no means to upload and download content other than to and from their own enterprise infrastructure. The micro-zone is surveilled and recoded so that misbehavior can be identified (in terms of information leaks or activities not strictly in keeping with the needs of the JV) and attributed (to the responsible organization). Contracts provide for the effect reduction of detected leaks, and loyal insiders can work more or less freely using the tools they always use without worrying about exposing content not authorized for the JV or used within the JV but not made available to the partners.

C. Enterprise-level access to key infrastructure

As part of widely deployed zone architectures, there is sometimes a control zone and/or audit zone, separated from functional business areas, and used for the purpose of managing the systems and networks. Access to areas within the control and audit zones are particularly sensitive because of the potential indirect aggregated effects on the enterprise of these zones and their subzones. But because of the global nature of such operations, there is a requirement for remote access, even if only from network operations centers (NOCs) or other similar areas, to support efficient 24x7 5-continent operations.

For operations that affect controls or gain access to audit information, users who work in these centers are not normally granted direct access, and the computers they use for day-to-day activities would aggregate enormous risk (e.g., the potential to bring down significant portions of the global enterprise or release financial information on a global basis before legally authorized release dates) if they were always potentially able to do anything their users had the overall capability to do (i.e., a large-scale ANA problem). A single break-in to one computer system in such an environment could be catastrophic if micro-zoning or a similar approach were not used, and this has been seen on occasion when global infrastructures have collapsed or nearly collapsed for major enterprises because of the acts (accidental or not) of trusted insiders.

A reasonably common approach over the last several years has been the use of terminal server-based use (e.g., Citrix), of encrypted remote desktop connections to state-retaining VMs with controlled configurations, surveillance, recording, limited actions, and only within-micro-zone shared storage for supervised activities over limited time frames. These micro-zones are designed to perform specific tasks, such as managing routers and switches, and have special purpose software for those purposes. Each such function is in its own micro-zone, and the micro-zones have independent storage so that they need not interact in order to support shared management from control centers across the world as their operations proceed over time. Update servers for patch management, identity management workflow servers for supporting global authorization, public key infrastructure management systems for supporting global use of encryption, and similar mechanisms are only accessible from specific micro-zones using controlled software, and these micro-zones are only accessible from global control centers.

Insiders are the only actors that are supposed to be able to gain access without physical break-ins to physically controlled data centers. By using micro-zones, insiders are limited in their activities based on available interfaces provided for periods of use of each capability they are granted access to. An accidental act has limited effect, and since access is only for use in limited interfaces, infection with malware, viruses, etc. are not feasible in the controlled systems on the terminal servers. However, while in micro-zones, the insider with a computer not as well protected may become infected, have their usage subverted, or intentionally act in a disloyal manner.

Further controls often include, without limit, separation of duties, multi-person control, submit-commit cycles, and change control. Separation of duties is used to limit the totality of functions allowed to any insider, and to separate different aspects of operational control among multiple parties. This is supported by micro-zones to the extent that they are used as an enforcement mechanism for this control. Multi-person control supports requirements that more than one person agree to a change. For example, in a workflow system, a proposed change may require one or more approvals from specific individuals from different groups before becoming effective. Micro-zoning is rarely used for this function. Submit-commit cycles are another approach to assuring independence between requests for change and invocation of those changes. For example, one micro-zone might support proposing a change and another micro-zone might support committing the change. While this can be done, it rarely is, because other mechanisms such as special purpose physical devices that display the proposed change and allow it to be approved are also available. Change control is a reasonably good place to use micro-zoning. In properly change-controlled environments, there are separate subzones an/or zones for research and development, testing, and production. But micro-zoning can be used to the same end, providing temporary access for testing and change transfer, and enforcing flow limits and read-only for changes.

IV. Conclusions

Micro-zoning is a viable technique for limiting undesired effects of insiders. While it is more effective at limiting undesired effects for loyal insiders, it also has substantial effect and value in limiting both direct and indirect effects of disloyal insiders.

References

[1] http://fuse.sourceforge.net/

[2] “Encrypted Filesystem Howto”, https://help.ubuntu.com/community/EncryptedFilesystemHowto

[3] IEEE, ``Draft Standard for Virtual Bridge Local Area Networks," P802.1Q/D1, May 16, 1997.

[4] F. Cohen, “Security Decision: Zoning your network”, 2008-11, http://all.net/Analyst/2008-11.pdf

[5] F. Cohen, “Any is not All”, http://all.net/Analyst/2011-03.pdf, 2011-03.