

Lock&Track Online AIRS: Availability, Integrity, Reliability & Security

— A Layered Approach —

Dean Woodward and Lorin Ricker

Introduction

The issues of computer system *availability*, *integrity*, *reliability* and *security* are of justifiable concern to all users, managers and technical support staff, especially when that computer system performs a mission-critical, 24×365 function. This is particularly true for the customers of the **Lock&TrackSM Corrections Information System** and the **Lock&ForceSM Sheriff's Information System**, both as on-site systems or premise-based installations, and in the form of the **Lock&Track Online** service, each of which provides mission-critical offender management functions for law enforcement and corrections organizations nationwide.

This paper addresses various aspects of **Lock&Track** “AIRS” (availability, integrity, reliability and security) attributes in order to make a convincing case for the product’s fitness-to-duty in its intended application niche, jail and corrections management. Each reference to and assertion about **Lock&Track** and/or **Lock&Track Online** throughout this paper can just as easily be applied to **Lock&Force**.

The format of this discussion is that of a FAQ (frequently asked questions), written for corrections and law enforcement managers and their technical support staff. The discussion focuses on the top-level overview — specific technical details about security issues are omitted for obvious reasons.¹

AIRS FAQ

1. Why should we consider Lock&Track to be reliable and secure?

There are several reasons why you can have confidence in **Lock&Track**:

- **Lock&Track (L&T)** has been engineered with the goals of AIRS built-in from the beginning; security and reliability are not after-thoughts, but have been designed in. The **L&T** applications and database are protected by a *layered approach*, much like the layers of an onion — no single component guarantees security or reliability, but these layers ensure that the product meets and exceeds all reasonable expectations for AIRS.
- **L&T** is based on *foundation technologies* (hardware platform, operating system and database management system) which themselves are “very high achievers” with regard to AIRS attributes. Several of these contributing technologies are discussed in more detail later in this paper.
- The **Lock&Track Online (L&TO) server systems** are physically housed at a secure “co-location” site, supported by uninterruptible electrical power, high-speed Internet network feeds, state-of-the-art environmental controls (including fire protection), physical security, and 24×365 system and network operations monitoring. On-premise **L&T** installations are similarly designed to include UPS, high-speed Internet and LAN network support, and other aspects of mission-critical information system support as appropriate to that site or facility.

¹ We don't want to invite trouble.

- The **L&T technical staff** takes AIRS issues very seriously, and has a high level of competence and experience in these areas. Their expertise is put into practice in numerous ways, including 24×365 system and network monitoring, policies and procedures, durable and well-tested operational scripts and procedures, including systems and database backups, intrusion monitoring, and much more.
- Furthermore, we *practice what we preach*: Our technical staff actually hold routine drills and practice scenarios to ensure that we're competent at intrusion detection, data verification, disaster recovery and database restorations.

2. What exactly do you mean by “AIRS”? Don't these words all mean pretty much the same thing?

While the qualities of computer system *availability*, *integrity*, *reliability* and *security* are all interrelated, they don't mean the same thing. And, to confuse things, casual use of these terms blurs the important distinction. Here's what we mean by each term as applied to a computer and application system:

- **Availability** — Refers to the uptime of the system: how much time the system is able to be used as intended to benefit its user community. Ideally, a mission-critical system is available to its users 100% of the time, 24 hours a day, 365 days a year, with no time off (downtime) for system malfunctions, power outage, network communications disruptions, maintenance activities, or even routine activities such as system and database backups.
- **Integrity** — Refers to the overall accuracy, completeness and soundness of the data and applications hosted on the system. Data integrity usually refers to freedom from data errors from whatever causes. Application and system integrity refer to the absence of bugs and immunity from external disruptive influences such as viruses and other malware.
- **Reliability** — Refers to both the dependability of the computer system hardware and the stability of the entire computer system as it performs its functions, including all aspects of operational procedures which contribute to that stability.
- **Security** — Refers to the ability of the system to withstand malicious or mischievous actions, either by outsiders or legitimate users, which are intended to disrupt, damage or shake confidence in the system, its functions and/or data.

From the perspective of the **L&T** management and technical staff, AIRS is a multi-faceted, challenging set of operational goals — from *your* perspective, as a prospective or current **L&T** or **L&TOnline** customer, there is just one overriding question:

Is Lock&Track at least as reliable and secure as my own in-house JMS run by my own DP staff?

The remainder of this paper is intended to convince you that **L&T** will *meet or exceed* your organization's requirements for availability, integrity, reliability and security.

3. What is a “co-location” and why is it a good thing?

The **Lock&Track Online** servers are housed at a professionally managed “co-location” (co-lo) facility near downtown Denver. This facility has multiple high-speed, high-capacity network feeds to the Internet; these independent network lines are physically separate from each other, and connect to different ISPs and telcos.² Note that, since reputable co-lo businesses exist in all major U.S. cities, and are expanding into many smaller metropolitan areas as well, it is conceivable that a traditional, on-premise **L&T** installation could be moved to a co-lo site to enjoy the same operational benefits as discussed here.

² Computerese for “Internet Service Providers” and “telephone companies.”

The co-lo's onsite NOC (Network Operations Center) is staffed 24x365 with experts who manage networks; they are capable of diagnosing and responding to all kinds of network-related problems, including any Denial of Service (DoS) attacks, network congestion, physical link problems, *etc.*, that their customers may report to them. The **L&T** technical staff have direct access to the “tier two” (second level) co-lo technical support team³ who are ready to actually help resolve any issues we may face.

The Denver co-lo facility is located near Invesco Field (the new stadium that replaces the city's famous Mile High Stadium) and the Pepsi Center (where the Denver Nuggets and Colorado Avalanche play), which puts it in the near-unique position of having *separate power feeds from two different power companies and grids*, in addition to the high capacity onsite *UPS (uninterruptible power systems) and backup electrical generator* that are standard for this type of facility.

State-of-the-art environmental controls provide *constant temperature and humidity* which is perfect for computer systems. *Fire detection and suppression systems* are likewise state-of-the-art.

The servers themselves are *physically locked into rack storage*. Only three people (members of our **L&T** technical staff) are allowed unescorted physical access to the machines — and they must present photo-IDs and log into/out of the building.

In short, we have elected to site **Lock&Track Online** servers at a reliable co-lo because this affords us, and you, the very best environmental services, physical security and operational stability for this mission-critical application. Our current co-lo vendor provides superior services, and contributes at least an order of magnitude improvement in AIRS for **L&TO**.

4. What kind of firewall technology does Lock&Track Online use?

L&TO servers are protected behind a Symantec (formerly Axent) **Raptor** firewall (for product information, see <http://enterprisesecurity.symantec.com/products/>). Raptor is a “stateful, proxying” firewall; that is, the firewall software inspects the contents of each packet it passes to ensure that the data inside is consistent with the protocol or application it purports to be. Network traffic which fails inspection is discarded and logged as a potential attack. **L&TO** technical staff routinely reviews and monitors firewall logs to detect any “barbarians at the gate,” our early-warning alert for possible service intrusions and breakins.

No user workstations are “inside” the firewall, only the **L&TO** servers themselves; any attacks against the machine are by nature from “outside” the firewall and subject to its inspection and rules base.

The **L&TO** technical support team also accesses the servers through the firewall and are subject to the same access rules as “regular” users. When, on the odd occasion they need a higher level of access, they have the option of connecting to the servers through a VPN (encrypted virtual private network connection), protecting their “privileged” sessions and data.

5. What is the hardware platform used by Lock&Track?

Lock&Track runs exclusively on HP/Compaq **AlphaServer** systems — currently, the **L&TO** production server is a Model DS20, with dual 533 MHz CPUs, 4 GB (gigabytes) of physical memory (RAM), high speed 10/100 Mb Ethernet, and approximately 50 GB of RAID-5 disk storage. The dual CPUs allow computing to continue in the event of a failure of either CPU; RAM is ECC, and can detect and correct any single-bit errors, while all of physical memory is controlled by highly sophisticated memory management algorithms; RAID-5 disk permits the hot-swap of any single failed disk drive without loss of data or down-time. At other customer on-premise JMS sites, **Lock&Track** runs on various other models

³ We bypass the tier one team who ask questions like “Is your PC plugged in?” — in a true emergency, we expect to get to the core of the matter very quickly!

of AlphaServer systems,⁴ and, when the situation and workload warrants, we will be able to scale up the **L&T** service as needed by moving it to a larger, higher capacity AlphaServer system and/or by adding VMScluster nodes to the existing configuration.

L&T development activity is carried out on a second AlphaServer system, a Model 1200 which is configured nearly identically to the DS20 production server. In the insignificantly rare instance of a catastrophic failure of the DS20, the development system can be quickly reconfigured for production use within no more than 4 elapsed hours.

Vendor hardware maintenance contracts are carried for all **L&T** servers, with guaranteed 4 hour response times and escalating levels of service to full repair. On-premise **L&T** servers can be similarly protected with hardware maintenance contracts with guaranteed response times and service levels.

6. What is the operating system used by Lock&Track?

All **L&T** AlphaServers are running the most current and stable release of **OpenVMS** (also historically known as **VMS**). VMS was selected as the OS of choice for **L&T** because of the 24×365 nature of correctional management operations and VMS's *reputation for reliability*. Many hospitals, CAD/911, industrial and manufacturing processes,⁵ and financial centers around the country, as well as in Europe and Australia, rely on VMS-based systems for their truly mission-critical support functions, where a system being down or unavailable is not acceptable.

Furthermore, VMS enjoys a commendable reputation as a *highly secure* operating system. VMS was the first commercial operating system to be evaluated and certified at the C2 security level (discretionary access controls) by the U.S. Dept. of Defense (*c.f.*, DOD "Orange Book"). VMS provides secure, password mediated logins, fully isolated processes, multi-threading, and access controls on *all* objects such as files, databases, tables, physical and virtual devices, and memory spaces. System management and operations resources, tools, procedures and techniques are second-to-none.

There are *no default passwords* on any user accounts, and especially not on system (privileged) accounts. User password management can be fully controlled as to password length, content⁶ and expiration. User accounts are completely managed and limited as to object access, privilege and rights identifiers, login type and time-of-day, *etc.* Any aspect of system operation and/or user access can be audited and monitored. Finally, no unprivileged process (or user) can have any effect on any other process; each is completely isolated from the others.

This past summer, a VMS system was given the acid test: it was taken to a system cracker's convention in Las Vegas and subjected to every known hacker attack for two and a half days. In the end, the contest judges declared the VMS system "*virtually unhackable*," as it was the only machine to survive the event unscathed. See <http://www.openvms.compaq.com/openvmstimes/openvmstimes.pdf> for details on this event.

⁴ Production servers at our premise-based **L&T** installations range from DS20 class systems to larger VMSclustered ES40 systems. **L&T** will run on *any* contemporary VMS-based AlphaServer system, from the smallest to the largest models, and we are committed to port **L&T** to future Itanium™-based server systems as OpenVMS itself is ported to that architecture.

⁵ You may be surprised to discover that Intel has been using VMS for years to support its semiconductor fab lines, including those that manufacture the Pentium CPU chips.

⁶ For example, passwords are pre-screen for words that appear in the dictionary, for re-use of the username as a password, and for other obvious "weak password" attributes. The user's attempt to choose such a password is gently rejected, and the user is encouraged to determine a more secure password for his or her account.

7. But isn't OpenVMS a "dead operating system"? Haven't we heard that VMS is at end-of-life, and should be avoided? And isn't the Alpha technology dead, too?

If VMS could speak, it would quote Mark Twain's rejoinder that "Rumors of my untimely demise have been greatly exaggerated!" Contrary to numerous articles written by misinformed technical pundits that "VMS is obsolete/dead..." and "VMS and Alpha are at end-of-life, overdue for retirement...", here are the facts:

- HP's OpenVMS is supported by a dedicated team of experienced professionals engaged in an ongoing and aggressive product development plan. In addition to maintaining VMS as the most robust and secure operating system in the market, this team's current efforts include:
 - § Certification of VMS as Defense Information Initiative Common Operating Environment (DII COE) compliant, which carries with it a *20-year commitment to the U.S. Government* to support VMS.
 - § Itanium/VMS product port (see below).
 - § Significant technology advancements in Galaxy, networking, VMScluster, security and storage management/device support. VMS development is frequently in the vanguard of such efforts, and usually defines the direction that other industry-standard and commodity computing products will ultimately take.
- On June 25, 2001, Compaq (now part of "the new HP") management held a news conference at which they announced that Alpha, long the fastest processor on the planet, would come to its design conclusion with the late 2004 release of the Alpha 21364 EV79 part. Compaq would subsequently base *all of its enterprise server systems on forthcoming members of Intel's IPF (Itanium Processor Family), the IA-64 microprocessor.*
- HP has recently *completed the port of OpenVMS to the IPF architecture, and has transferred key Alpha processor and compiler technologies, tools and engineering resources to Intel*, all as part of a non-exclusive technology agreement between HP and Intel. This is the most concrete and firm rebuttal of claims that "HP is not committed to VMS" that you'll find. Itanium-based OpenVMS systems and servers are anticipated to be generally available in early calendar 2006.
- In early 2005, LockWorks LLC successfully completed the port of *all* of its applications products, including **Lock&Track**, **Lock&Track Online** and **Lock&Force**, together with all underlying support software such as RAPT, to the Itanium/OpenVMS environment. There were *no significant problems or challenges* encountered during this port, and LockWorks is among the earliest application vendors to have completed this product port. LockWorks products and services will be available on Itanium/OpenVMS platforms just as soon as these are ready for release by HP.
- AlphaServer computer systems and all Alpha-based products will be supplied and supported by HP/Compaq for at least 20 years, well into the 2020's. HP (and its predecessors Compaq and DEC) has an excellent track record in supporting current and aging systems products — the VAX product line, technology from the 1980s, continues to be supported in hardware and software to this day, and new versions of OpenVMS are continually released to support VAX systems in synchrony with Alphas.

For corroboration of these assertions, refer to <http://www.hp.com/products/servers/> and <http://www.openvms.hp.com/> (both of which are vendor web-sites), and <http://www.openvms.org/> (which is an independent web-site).

8. What database technology is used by Lock&Track?

The **Lock&Track Offender Database** is based on the **Oracle/Rdb Database Management System**. Oracle/Rdb was selected based on the following criteria:

- It is designed for *high-speed, real-time data access* and *superior database performance in highly demanding applications*.
- It is a fully *relational database*, and uses *industry-standard SQL* for all data access and operations.
- All *data access is transactional*, with commit or rollback completion, to ensure *multi-user concurrency and data integrity*.
- All *data operations are journaled* in an After-Image Journal (AIJ) file, which enables complete database recovery up to the point of last-committed transaction in the rare case of a database or disk failure.
- All *data tables and views are subject to VMS access controls* to fully mediate any user's access to data. Each individual **L&T** user is granted a specific data access profile which is appropriate to his/her job function and "need to know" — users can see *only the data for which they are trained and authorized*, and all other data is secured against their use.

Furthermore, because of the combination of VMS and Oracle/Rdb data access controls, *no unauthorized, non-L&T user can ever gain access to any part of the Lock&Track Offender Database*.

- It includes *sophisticated database management tools and utilities* to support 24×365 database access and operations, including online database verification, backups, and numerous data management activities.
- Oracle has committed to follow Compaq's port of VMS to the Itanium/IPF architecture and *will port Oracle/Rdb* as soon as Compaq can provide VMS/IPF development capabilities.

In summary, Oracle/Rdb is designed to guarantee application data integrity, availability, security and reliability. **L&TO** exploits nearly every trick and resource inherent in both VMS and Oracle/Rdb to ensure that this remains true.

9. What about user access and security?

Lock&Track users are all given *unprivileged* accounts; that is, the account that the user logs into has only the minimum set of access rights needed to run the application. In fact, the VMS security model lets us set users up so that they can't even see, let alone access or damage, the **Lock&Track Offender Database** except through the **L&T** applications. In short, users are permitted access to just the limited set of data elements and operations required to perform a specific job function, nothing more.

Each user of **L&T** is given his or her own account to use. Account sharing is specifically *not* allowed, and is considered a violation of the **L&T** use agreement. Since there is no cost basis to the customer on the number of user accounts set up, there is no benefit to account sharing; to the contrary, giving each user their own account encourages users' ownership and accountability towards the data, and this in turn helps to accrete overall data integrity.

Data operations are audited in various ways within **L&T**, which allows technical and corrections management to identify and assist specific users who need additional application training or remedial attention regarding the correct use of the data — this also serves to immediately identify any user who attempts to perform malicious activities against the data or system.

Unfortunately, users are often the *weak link* in computer system security. To mitigate this, **L&T** provides all users with a self-paced study guide entitled "**Lock&Track Security and Passwords (Everything you**

need to know...)”, and we encourage the management at each customer site to incorporate this study guide into every **L&T** user’s training. We also encourage all users to take ownership of the data, its quality and integrity, and we encourage management to direct and empower their staff accordingly — after all, it’s *your data*.

10. System and data backups? How often is the L&T database backed up?

All **Lock&Track** installations, both on-premise and **L&TO** itself, are managed according to the same rigorous backup schedule and procedure:

- Backup tapes are rotated weekly, and are stored in an off-site facility.
- When a fresh backup tape is rotated into the tape drive, it will be initialized and a full disk backup (*all* disk volumes) is performed as scheduled that night under automatic system control.
- Each subsequent night during the week, an incremental disk backup is performed.
- Prior to each nightly disk backup, a **Lock&Track Offender Database** full online backup to disk save-set, plus a backup and reset of the After-Image Journal (AIJ), is performed — this save-set of the database and the AIJ-copy are both included in each backup-to-tape, ensuring full database recovery at any point.
- Note that all database and disk backup operations are performed *online*, with *no interruption of services* to the user community. **L&T** is *never* taken down for backup operations.
- Also note that these backups are all performed automatically (in regularly scheduled batch jobs), and all resulting operations are monitored and checked each morning by **L&T** operations staff.
- Finally, the **L&T** technical staff routinely practices all aspects of disk and database recovery and restoration procedures, so this is never “just a drill” to us — nor would we panic in the face of an actual disaster and recovery situation. System and database recovery procedures are a skill that we maintain, not a “policy” which is documented “just in case.”

11. How do Lock&Track users connect to and use the applications and data?

An authorized user accesses **L&T** (either the **L&TO** service or an on-premise **L&T** system) from a standard Windows PC (workstation) upon which the **Lock&Track** “client” application (*LTclient*) has been installed. The user activates *LTclient* from an icon and is presented with a **Lock&Track** sign-on screen which collects his/her username, password, the name of the **L&T** application server (host) to connect, and the **Lock&Track** desktop to activate.

When the user clicks the Start button to initiate the application login process, a Telnet session is initiated between *LTclient* and the **L&T** application host server. For **Lock&Track Online**, this Telnet session is initially directed to the public IP-address for **L&TO**, namely lto.locktrack.com, which is first intercepted, examined and blessed by the Raptor firewall as acceptable incoming Telnet traffic. The firewall’s built-in rule set automatically directs this Telnet session to the **L&TO** production server, and specifically prevents Telnet traffic from accessing any other computer resource behind the firewall (*i.e.*, browsing and probing is prohibited). Similar network and firewall services can be applied to an on-premise **L&T** system as well.

The **L&TO** production server responds to this Telnet initial conversation by executing a standard VMS login challenge and authentication (and see questions 11 and 12 below). If a valid username/password combination has been provided by the user, a server-side application process is created and connected to the Telnet session, which then carries on a **Lock&Track** application conversation until the user request a logout action.

Note that there are several points where the application session is monitored and controlled (layers of security and protection), and each of these is routinely monitored by our technical staff members.

12. Isn't Telnet insecure?

We are frequently asked about our decision to use Telnet as the base protocol for access to **Lock&Track Online** over the Internet. Telnet is very useful because it is a routable protocol, simple to implement, imposes reasonable load on available network bandwidth, and is available nearly everywhere (*i.e.*, on all Windows PC/client workstations). **L&T** implements a private application protocol on top of Telnet, embedding a two-way, in-band command and data conversation between each PC/client (user) and the application server.

It is, theoretically, possible for a hacker to “hijack” a Telnet session. In practice, it would be extremely difficult to pull together all the necessary pieces to confuse both our firewall and server enough to take over an established session, together with physical and surreptitious access to the network and/or system under attack. In particular, the slightest error in synchronizing the client-server conversation is sufficient to terminate a user session, making a hijacked **L&T** session prohibitively difficult to sustain.

13. Doesn't Telnet send insecure passwords?

It is true that native Telnet sends all traffic, including usernames and passwords, in “clear text.” **L&TO**, however, does not. The **LTclient** encrypts the password using a “one-time pad cipher” function; the application server has a “hook” built into its login function that intercepts the password and decrypts it to complete the login function. Clear text passwords are never sent between the **LTclient** and the **L&TO** application server, and the password encryption changes from session to session.

14. What about other Internet services provided by Lock&Track?

Both FTP and HTTP services are supported by a **L&T** production server, and each provides valuable auxiliary features and functions to enhance the application service. HTTP provides read-only data and web-page access, and various CGI (*etc.*) scripts are carefully designed accordingly.

FTP access is similarly limited to out-bound support services only, and is typically used by **L&T** customers for **LTclient** client upgrades (via web-page mediated downloads). Anonymous FTP is not enabled on the production server.

The Raptor firewall includes appropriate rules to mediate and detect inappropriate use of all of the various and commonly used TCP/IP protocols.

15. Are unauthorized system breakins possible or detectible?

Presuming a hostile or unauthorized agent could defeat the Raptor firewall and other outer layers of **L&T** service protection, that perpetrator would ultimately have to penetrate the VMS operating system login and intrusion detection safeties. To do so, he would have to either: i) Compromise a legitimate application user to gain access to that user's username and password;⁷ or ii) Make repeated attempts to guess usernames and passwords to gain login access.⁸

⁷ This is typically done by what is called “social engineering,” *i.e.*, a hacker contacts a legitimate user while masquerading as an “official representative” of the **L&T** team or your own organization, and attempting to convince or coerce your honest user to divulge his username and password to the perpetrator. **Lock&Track** includes user training and documentation which is specifically designed to teach and alert our user community to this possibility and how to avoid falling victim to it.

⁸ Sometimes called “war-dialing.”

Both user training and vigilance are sufficient to thwart the first situation — this is also a primary reason why shared user accounts (*i.e.*, “group logons”) are *never permitted* in **L&T**.⁹

To guard and thwart the second kind of attack, VMS provides an always-on intrusion (or breakin) detection function which continually monitors all attempts to login to the system. If a legitimate user flubs the entry of his password, he is allowed to retry the login procedure — and most users can routinely login to the system within one, sometimes two, attempts.

However, if an outsider (sometimes referred to as a “hacker”) attempts login by guessing at usernames and passwords, the following occurs:

1. First of all, the odds of a perpetrator correctly guessing a correct username and a strong password on the first, or even second, try is highly unlikely — on the first unsuccessful attempt to login (including innocent typing errors by a legitimate user), VMS goes into “intrusion alert” mode.
2. An intrusion becomes a “suspect” at the third unsuccessful login attempt, and the system posts alert messages to system managers and operators, including a constantly active intrusion analysis batch job in **L&TO**.
3. By the fifth unsuccessful login attempt, VMS is convinced that this “suspect” is a bona-fide intruder, so it posts additional system alerts and enters “evasion mode.”
4. In evasion mode, VMS continues to solicit the perpetrator’s efforts with additional username/password prompting. However, this is now just a ruse to keep the perp “on the hook” for possible identification and apprehension.

From this point on, even if the perpetrator gets lucky and correctly guesses a valid username/password combination, VMS *will not login the user process*, but simply keeps prompting for login input. Logging in from this particular terminal/source is now disabled pending system management intervention — this disabled state persists until either a random time interval has elapsed *or* a system manager specifically cancels that status for that login source.

5. Note that: a) Other legitimate user logins and processing are not affected in any way by this activity for a breakin perpetrator. b) Unauthorized login activity is *always* immediately detected and quickly defeated, without the perpetrator being tipped off that his activities are being monitored. c) This process forgives the occasional typing mistake by a legitimate user. d) All such intrusion detection and evasion activity is audited and monitored automatically by the system, with full visibility of each occurrence by the system management staff.

Other “classic” system breakin attacks and hacks which have been historically successful against various Unix, Linux and Windows/NT systems (including, recently, a vulnerability found and exploited in SSH v1.0) are well documented to be computationally infeasible¹⁰ or practically impossible against a VMS system, including such hacker favorites as: i) Theft and decoding of the system’s “password file”; ii) Buffer overflow attacks; iii) Trojan horses; iv) UID manipulation to gain “root” access; v) Misappropriation and misuse of privileged commands or utilities; vi) Viruses and worms — for a variety of technical reasons, VMS is immune to these...and many more.

In conclusion, VMS is “virtually unhackable”¹¹ on its own, especially under experienced and watchful systems management; the additional layers of extra-VMS security defense provided for the **Lock&Track Online** service are simply the icing on the cake.

⁹ A *shared* secret password isn’t (...secret).

¹⁰ As in: there aren’t enough years left in the lifetime of the planet Earth to compute all the possibilities.

¹¹ Again, see the report on VMS’s performance at the DEF CON 9 hackers convention as described in the Appendix to this paper.

16. What about computer viruses? Is Lock&Track susceptible to these?

As noted above, VMS and VMS-based applications such as **L&T** are immune to viruses. Although *in theory* it is conceivable that an executable image under VMS might be infectable with a virus, in practice this would be extremely difficult to do: unlike a Windows, Linux or Unix executable program, which is a simple stream of executable bytes and data, a VMS executable is encoded as a highly structured file, which makes grafting on a malicious chunk of virus code nearly impossible — doing so would render the executable file invalid and corrupt to the VMS image activator, and it just wouldn't run.

Furthermore, VMS executable image files are subject to the same access controls as are all other protected objects within the operating system, and any agent attempting to attach a virus to a program file would also have to defeat access control mechanisms, which is impossible without privileged access.

There are *no Microsoft Windows-based computer systems* within the set that comprise the **Lock&Track Online** service itself;¹² this fact certainly mitigates against virus infection within the bounds of the **L&TO** service environment.

Of course, *your* users will be using Windows-based PCs as workstations to access **L&T**, and these are certainly subject to PC virus attacks via the usual channels (infected floppies and FTP file transmissions, e-mail macro viruses, *etc.*). It remains *your organization's responsibility* to safeguard your PCs and other computer systems against the various kinds of malware which are prevalent in commerce, government and industry today.

17. What about the customer's end of the connection? How can my organization connect securely to Lock&Track Online?

Customer sites have several options in how they would like to connect to **L&TO**. Most customers without a designated site security manager have rather simple rules blocking most inbound traffic; for these sites, connecting to **L&TO** is not an issue.

More advanced security sites may adjust their firewalls and/or proxy servers to allow outbound Telnet (port 23) to our **L&TO** server site.¹³ Those sites that have VPN (virtual private network) capability, or desire to implement it, can be accommodated by the **L&TO** Raptor firewall's VPN capabilities. In particular, a VPN connection provides end-to-end encryption of the entire Telnet session, including the application conversation itself, providing an extra layer of impenetrable data privacy over the Internet.

Note that in order to provide Raptor-compatible VPN capabilities,¹⁴ the customer may be required to purchase an additional VPN appliance for outbound connectivity.

18. Can other corrections users at other customer sites access my organization's data?

The answer here is both *yes* and *no*... er, it depends. There are two issues here:

- **Yes** — One of the *compelling benefits* of the **Lock&Track Online** approach to jail management systems is that, indeed, the *Offender Database* is a common resource shared among *all*

¹² We *do* use a couple of Windows/NT-based systems for non-critical functions, including an Exchange e-mail server. And, lest you think that there's an anti-Microsoft bias within our **L&TO** support team, we also do use NT-based PCs as our routine development and support workstations — these, like yours, are located *outside* of our secure firewall perimeter.

¹³ For those sites not comfortable with this, the application has a configuration option to run on another port — 2001 by default — but it can in fact run on any port not taken up by another standard service that L&TO uses (these include HTTP port 80, and FTP ports 20 and 21, for example) with a safe and simple registry tweak.

¹⁴ A VPN is a VPN, is... not always a VPN. Certain brands or models of VPN devices have limited or "hardwired" settings that are mutually incompatible with what the L&TO Raptor firewall is capable of; unfortunately, this is yet another example of where a computer standard isn't. L&TO engineering and data security staff are ready and willing to work out these issues with our customers' technical staff as necessary to satisfy your organization's level of comfort regarding secure connectivity.

of the subscribing governmental organizations who use it. A common complaint (or wish) among many corrections and sheriff's departments is "If only we could share data about these offenders with our colleagues and other law enforcement partners..." — well, **L&TO** enables an appropriate level of data sharing as a matter of course.

This means that: a) Users do not have to cope with *duplicate data entry* — if an offender's data has been entered in one jurisdiction, then that same data is available for use and updating by the next facility which incarcerates that same offender. b) The offender's data set accretes constructively over time, forming a complete and accurate criminal history for him — and this history is *comprehensive*, no longer restricted only to his activities within your jurisdiction. c) Furthermore, you have access to his entire behavioral and disciplinary history, including any misbehavior at other institutions, which often enhances the security and safety of your own staff and other inmates.

The data elements shared among all L&TO users are those which are specific to the offender himself — his physical characteristics and descriptors, his name/aliases, DOB, identifiers such as SSN and DL, his criminal history (all of it), his behavioral and general disciplinary history, keep-separate information, court appearances, charges, sentences, and time-served, *etc.*

This shared data *enhances* your ability to manage this inmate and his incarceration at your facility. You *want* this level of data sharing.

- **No** — There is a set of **Lock&Track Online** data which is considered private and of concern to *your organization alone*, and this data is *not shared* with other **L&TO** customer organizations.¹⁵ Private data consists of: a) Information concerning your physical management and housing of the inmate, including current and past housing (cell) assignments; b) inmate's property inventory while in your custody; c) specific details of disciplinary actions, including sanctions, which you may take against the inmate; d) inmate trust accounting, including commissary purchases; e) risk assessments and other evaluations; f) medical records (which tend to be held highly confidential anyway).

The data elements which are not shared are those which are specific to your management of inmates in your custody — these records are typically not germane to other facilities, and certainly should not be viewed or updated by users in other facilities. They've got their own private set, and you've got yours. You *do not want* to share this kind of data.

So, with **Lock&Track Online**, you have the best of both worlds: controlled, yet effortless, data sharing with your corrections colleagues and law enforcement partners, *and* a private, secure database concerning your own inmate management activities "in your scope of control."

And, with an on-premise **Lock&Track** installation, the degree of data sharing that you enable for your colleagues and partners is entirely up to you. The base technologies, OpenVMS, Oracle/Rdb, RAPT and **L&T**, permit you to manage your data resource as openly, or close it entirely, as suits your organization's requirements and operating style.

¹⁵ Unless you specifically print out relevant reports and send them to your colleagues — this is in *your* control, not ours.

Conclusions

Hopefully, this FAQ discussion has been persuasive, and that you're convinced, or at least ready to consider, that both **Lock&Track** and **Lock&Track Online** are indeed secure and reliable. We've made several points:

- The desirable qualities of availability, integrity, reliability and security (AIRS) are explicit goals of **Lock&Track**, and the means to achieve these goals are inherent in the product and the service.
- We've used an onion-like approach to **L&T AIRS** which gives us a multi-layered set of defenses and tools to protect your corrections/JMS data.
- **L&T** is engineered for mission-critical, 24×365 duty, and is fit for service in your demanding operational environment. It is based on hardware, operating system, database and application foundations which are *proven* in other demanding industrial and governmental environments similar to yours. We will continue to evaluate and employ new state-of-the-art technologies as they develop to enhance all aspects of AIRS.
- The **L&T** technical staff can and will take care of your corrections/JMS data at least as well as your own technical staff would — and we'll strive to overachieve in this responsibility.
- The **Lock&Track** community of users is a key element in achieving AIRS goals.

The primary business goal of the **Lock&Track** team is long-term customer satisfaction — we want your organization to be part of the **L&T** community of users for years to come. AIRS principles are key to achieving that level of customer satisfaction, and our team is dedicated to excellence in pursuit of these goals.

— Appendix —

DEF CON 9 Hacker's Convention declares OpenVMS “cool and unhackable”

Note: The following text is taken unedited from the web-site URL provided below, and is included here simply for convenience.

DEF CON, a military term that refers to escalating military conflict conditions, is also the name of a computer hackers' group that meets every year in Las Vegas. At the DEF CON 9 convention, hackers from around the world got together to swap ideas, test and hone their hacker skills, and learn new techniques by playing a game called Capture the Flag.

To many professionals in the computer business, taking an OpenVMS system to a place where 4,300 hackers can try to break in for two and a half days is analogous to walking into a back street bar and flashing money around before stepping into the alley for some night air! Three members of the Dallas/Ft. Worth Compaq User Group (the DFWCUG) decided to take an OpenVMS system to DEF CON 9 and play Capture the Flag. The contender was a Compaq AlphaStation™ 4/233 system with 512 megabytes of memory, OpenVMS v7.2-1H1 operating system, TCP/IP v5.0a, Apache, and Point Secure security software. All software was standard and installed out-of-the-box. Also loaded onto the system were a few added services such as WEBserver pages, interactive Telnet accounts for any hackers who logged into the OpenVMS system to hack from the inside, and a public “Games” account for hackers who got tired of hacking!

For two and a half days, the hackers bombarded the server with different TCP/IP attacks and some internal attacks- but none of them was able to break the security or hack into the OpenVMS server. Throughout the event, Point Secure Software's System Detective product recorded every attack and every keystroke, and gave the system an extra layer of protection from the hackers. On the last day of the event, during the last half hour of the Capture the Flag contest, the judges put a note on the scoreboard that they thought the “OpenVMS system was virtually unhackable.” Immediately, all hacking attempts against the AlphaStation system ceased.

For the last half hour of the contest, the OpenVMS system coasted across the finish line with not one of the hackers bothering to waste their time on the OpenVMS server! At the DEF CON 9 wrap-up session, the judges declared the OpenVMS server “cool” because its services were continuously available and never hacked during the contest. The rest of the hacker teams also gave the server “props” (kudos) as well because they were not able to “root” the system or break in. Steve Smiley of the DFWCUG delivered a security white paper on what was learned about the hacker's attacks at CETS 2001 in Anaheim, California this month.

For more information, see the DFWCUG Quadwords newsletter at <http://www.dfwcug.org/>. Also visit <http://www.defcon.org/>, <http://www.pointsecure.com/> and <http://www.cets2001.com/>.