Chapter 7: From Chaos Is Born...

Returning to the battle the next day after breakfast, coffee call with the brew crew, morning email check, and normal cycling/showering, Lionel restarts his efforts. He turns-on the audio/video equipment. He then restarts the 'ole' IBM ThinkCentre, LinuxCNC simulator in its Internet isolation state.

Maintaining the machine Internet isolation, Lionel begins a brief manual scan of the toerr, deyfine, and then gonza.

"The size of toerr's home directory is only 239 MB. There is nothing in the Download directory. Its Desktop only includes links to Arduino and Sloeber. The snap directory only has Firefox. From these observations it can concluded that this account was reserved for recreational tinkering with Arduino components. Sloeber is an enhanced eclipse specifically designed for working with Arduino projects. The Arduino technology is centered on hobbyist and low-end prototyping. The likelihood of something being derived from the toerr home directory is very small.

The deyfine home directory has a size of 1.4GB. It's Download directory has Firefox and LinuxAirCombat sub-directories. The Desktop directory has links to the 'gl-117_Flight_Sim, mbot, and a gamepad controller. As previously noted the deyfine has an unused Thunderbird client installed.

"There is nothing remarkable about either the toerr or deyfine home directories.

"The size of the gonza home folder is significantly larger at 122.3GB. It overshadows the other two home directories. But the majority of the size is due to the Picture directory. It contains 30.9k images with a total size of 71GB. The Download directory is also large with 32.9k files and a size of 24.1 GB. Scanning across the gonza home directory, it is observed that numerous software development applications are installed. These include java, C++, eclipse, Arduino, MySQL, Android Studio, and freePascal. The snap directory contains twenty-three application. Of which 'ant', 'zoom-client', eclipse, and 'guvcview' appear to be additions that occurred in late September and early October 2020. That would coincide with the Spinnaker project effort.

"Keeping with the plan from last night of conducting a forensic analysis of sys' level directories 'bin' ,'media', 'opt','usr', and 'var' . We will now move to the restored version of Frank's computer. We will scroll through each of the directories as planned last night. If something is found that might be of interest we will halt and take a deeper look to our satisfaction. Then we will move forward.

"Starting with the 'bin' directory, let's begin.

"The 'bin' directory appears to be an example of a typical Ubuntu installation. It is unremarkable.

"The 'media' directory is unremarkable

"The 'opt' directory contains 76k files with a total size of 7.3GB. The applications contained in that directory are typical for a developer engage in code development. They include MySQL and eclipse as already having been mentioned from the look at the gonza home directory.

"The 'usr' directory contains 545 files with a combined size of 12.9 GB. Its 'bin' folder contains 3K files with a size of .9GB. Its other major supporting directory , 'lib' , contains 87.9K files with an aggregate size of 7.0GB. Within those two sub-directories of 'usr' we find java, fpc, gcc, python, javascript, lazarus, mysql, mysqlworkbench, odbc, ssl, svn, git, and graphviz. These are tools that an adventuresome developer might employ. There is 'a lot of meat here'.

“Moving to the ‘var’ directory we find it to contain 23.7K files with a total size of 9.3BG. Let’s look first at the /var folder and then we will look at the /var/log/syslog grouping of files. If there are any other “interesting” folders or files we will then look at them briefly.

“The /var folder appears to contain a normal collection of sub-folders and log files. Of the log file groupings only the system appears to be of interest. The ‘/var/tmp’ folder has a typically large number of “systemd-private-GUID-somedescriptor-text”. These are locked folders but they are typical output folders for Ubuntu services that run in background. If time permits we will take a look at specific date ranges. But this will be very consuming with low return on our time.

“Sincw Frank was engaged in software and Web development, the ‘/var/metrics’, ’/var/svn’ , and ‘/var/www’ folders may be telling.

“First, let’s look at the ‘/var/metrics’ . This folder is typically used by clients to collect developer performance metrics. A client sometimes require that the developer permit to embed an ‘agent’ that collects specific development data. The ‘agent” runs in background assembling the data as it relates to the project. At a some point or time-trigger the collected data is uploaded to a host for analysis in terms of the project’s hours and developer parametrics.

“A review of Frank’s ‘/var/metrics’ fails to show any old metrics files. But there are daily new files being collected and archived since we do have the ‘ole’ IBM ThinkCentre, LinuxCNC simulator isolated from the Internet. Now that we know we should look for an ‘agent’ working in background. We must also look at the metrics files that have been collect during the last few times we have had the simulator powered. They may be innocent or they may be ‘spying’ surreptitiously on the developers.

“Let’s move next to the ’/var/svn’. This is not a default folder but it is very commonly used for local project configuration. As such seeing it on Frank’s machine is expected. Withing the ’/var/svn’ is the ’/var/svn/repos’ folder. Franks’ contains ’/var/svn/repos’ folder contains ten project sub-folders. Switching back to Frank’s home folder for a moment, we find that he has a shell script for invoking the svn, ‘start-vn.sh’. We will give that a kick later to see what each of the project sub-folders might add to our understanding. Four of these sub-folders have modification dates that suggest they might have been used during Frank’s last project.

“If Frank is as lazy a software engineer as I was, he will have the ‘svn’ configuration application from its web version. That would have allowed him to enter a couple of commands and then with a couple of clicks he could navigate to a web browser ‘read-only’ presentation of this local repository.

“Running the ‘./start-svn.sh’ and ‘./start-web.sh scripts from the user ‘gonza’ account produces a local web page *http://localhost/svn/* showing root of the repositories, ‘repos’. With a triage priority we will take a look later as shown in Table 2.

Table 2: Collection of Repositories under the ‘repos’ svn

aapMbot/	==> need follow-up ==> Scan only
asr4android/	==> need follow-up ==> Scan only
asr4arduino/	==> need follow-up ==> Scan only
eclipse_devwrk/	==> need follow-up ==> Number 1
heeks/	==> not relative to project
instructables/	==> need follow-up ==> Scan only
openbox_del/	==> not relative to project
openbox_laz/	==> not relative to project
freecad/	==> not relative to project
sloeber/	==> need follow-up ==> Number 2
test/	==> need follow-up ==> Scan only

“Let’s take a short lunch break and resume with the ‘/var/www’ . Lionel powers-off the audio/video equipment and heads to the kitchen for lunch.

Returning from lunch he restarts the audio/video equipment and commence where he had stopped.

“Having exhausted the ‘svn’ lead it is time to move to the last ‘/var/ sub-folder of interest, ‘/var/www’ . This is the place where he would have placed any web application or web service that he may have developed as part of this last project. Within the ‘/var/www’ there are six folders. The ‘/var/www/wrkjar’ seems to be the only one with any promise of enhancing our understanding of his project work. It will need to be investigate as time permits.

“Let’s make a list of action items to follow-up and set their relative order of priority:
1> Scan syslog group of files of ‘/var/log’ with ‘grep’ using the 014_keyfile pattern file. Follow-up with manual scan using ‘genay’ on any which seem to be of interest. The Pattern file 013_keyfile will be updated to 014_keyfile using pattern tokens to be added in a moment.
2> Attempt to determine the ‘agent’ that is running and its bandwidth of surveillance, the destination of the collect project data.
3> Scan the web presentation of the ’/var/svn’ focusing on those repositories which were identified as needing follow-up.
4> Attempt to run a local Apache web host session to view the ‘content’ of the ‘/var/www/wrkjar’ folder.

“To prevent destroying any data, we are copying all of the contents from ‘/var/log/syslog*’ to a work folder. Now with all eight of the syslog* files copied we need to extract the contents from six of those eight files that were archive files so we can ‘grep’ them properly. The Ubuntu Archive Manager allows us to extract and list the files. I am doing the extract in sequential order.”

Waiting while the archive files are extracted Lionel thinks aloud, “It is important to keep in mind that the ‘eight’ syslog* files that were on the micro-SD card from Frank’s computer compose a snapshot of the what was happening on the computer. Specifically we are looking at the Jan 31, Feb 1, and Feb 2, 2020.

“Now with all eight files ready we will modify the dogrep4home.sh using 014_keyfile. We are adding six new patterns to our scan filter list. Those six new terms are ntpdate, ‘Jan 31’, ‘Feb 1’, ‘Feb 2’ ,‘mounted’ and ‘unmounted’ . The terms ntpdate is important because it will indicate the hosting company with which Frank’s machine performed date-time checks while associated with the Spinnaker project. The mounted, and unmounted wit inform us of when the SD-card was installed and uninstalled during Frank’s Clonezilla image creation.

“After a couple of trial passes we suspect that we shall see the results for that period of time that shows us whether the ‘Spinnaker’ project was developed using external media, local environment, or a cloud like GITHUB. “

As the grep scan of the syslog* files Lionel observes, “We seem to have caught a bit of luck. The first pass showed that ‘syslog’ shows the clear date demarcation between ‘Jan 31’, ‘Feb 1’. There does not seem to be any doubt that Frank stop using his computer and made a Clonezilla backup on February 1, 2020. After that he was in the ‘Spinnaker” project closeout meeting. Whatever happened from that time forward there is no record.

“Returning to the ‘grep’ results of ‘syslog’ we know that a final backup was probably performed using the micro-SD card labeled as SD_200201. As previously documented the micro-SD card contained a Clonezilla archive of Frank’s computer.”

Abstract from ‘syslog’ showing two log entries related to micro-SD card labeled: SD-200201 related to SD_200201:

```
Feb   1 15:53:46 mypond udisksd[6346]: Mounted /dev/sdd1 at
      /media/gonza/SD_200201 on behalf of uid 1000
Fed   1 18:24:413 mypond udisksd[6346]: Dismounted /dev/sdd1 at
      /media/gonza/SD_200201 on behalf of uid 1000
```

“We need to note two items before moving forward.

“ First, the last dated entry in the ‘syslog’ file before the entries related to our recent efforts was February 1, 2020,” continues Lionel. “That suggest that this archive is very nearly a complete collection of Frank’s work on the “Spinnaker” project.

“The ‘syslog’ also documents communication with a Sunnyvale, CA hosting company, VilleiNet via ntpdate transactions. A search of the Internet failed to verify any current entity by that moniker.

Abstract from ‘syslog’ showing a log entry related to network time date check with an host that no longer exists:

```
Jan 30 11:26:37 mypond ntpdate[1709]: step time server 96.45.34.9 offset
      25.865755 sec
```

“Moving next to the ‘/var/metric’ folder we find only current entries related to our efforts. That suggest that there is still an agent at work that is collecting “project data”. A search of the Internet using the query ‘how to determine what daemons are running on Ubuntu’ produced a useful OpenSource tool, rcconf’ to resolve which agents or daemons are running during an Ubuntu session on Frank’s computer. When invoked ‘rcconf’ produced a list of agents that were working in background.

That list contained only one entry of particular interest:

Abstract from rcconf showing list of active agents:

[*] indicates activated at boot
[ ] indicates deactivated prior to boot
….

```
      [*] mysql          Start and stop the mysql databas
      [*] nmbd           start Samba NetBIOS nameserver
>>>   [*] ntimesh        Jib Timesheet          <<<<<< Item of Interest
      [*] ondemand       Set the CPU Frequency Scaling go
      [*] plymouth       Stop plymouth during boot and st….
….
```

“Let’s deactivate the ‘ntimesh’ agent and reboot the computer to verify that the keystroke and mouse inputs are no longer being captured and logged to the ‘var/metrics’ folder.

After going through a power-off/power-on reboot, Lionel invokes the ‘geany’ editor and begins to type ‘Four score and seven years ago our fathers brought forth, upon this continent, a new nation’.

“Let’s see what today’s ‘/var/metrics/jib_yymmdd’ file has in it. Opening, the today’s file, ‘/var/metrics/jib_220110’ reveals a text file shows a collection of keystrokes, mouse clicks, and windows in which they occurred. Manually, scanning the file in the ‘geany’ editor fails to locate the Lincoln Gettysburg Address segment.

“Now we will re-run the ‘rcconf’ tool but enable the nitimesh agent this time, reboot, and type another phrase to verify that the agent is active and logging to ‘/var/metrics/jib_220110’. TO maintain our anonymity we will continue to keep the computer isolated from the Internet. “

Making a final pass through ‘rcconf’ tool Lionel sets the enable for the nitimesh agent. He then reboots the computer. Once he has login as gonza he activates a ‘geany’ editor session and types “i am responsible for my own debts and obligations only.’ Still in the ‘geany’ session he navigates to the ‘/var/metrics folder and opens jib_220110’. From the ‘geany’ session he scan for and confirms the presence of his last statement about debts and obligations” is now shown as a set of collected keystrokes. That action confirms that the computer has returned to the original status with the ‘ntimesh’ actively collecting keystrokes.

“We seem to have wrapped-up the agent issue with the exception of confirming where the ‘jib_yymmdd’ files were being sent. But the last thing we want to do is ‘touch’ the agent’s host before we know who exactly they are.

“We could ‘ping’ the IP address of the former agent’s host site, 96.45.34.9 from my laptop to see if it is still active. But the bundle of data that composes a ‘ping packet’ would contain our IP address. That we do not want to do. I will put that on my list of things to do the next time I am at the County Library Branch.

“Instead let’s run a search of the IP address. Invoking *https://whatismyipaddress.com/ip-lookup* and entering the IP Address of 96.45.34.9 we learn that it is a high security enterprise firm specializing in cybersecurity. We will call this a wall or obstacle which we will avoid for the moment.

“This appears to be a good point to break for supper. Tonight it is barbecue burgers.

Over supper conversation turns to the Stacy-Frank project.

“Well geek detective what’s the crime scene theory now?” quips Sandi.

“I do not really know. Except for knowing that they are both deceased now, I would still holding-out for the elopement to the islands for a few months of honeymooning.

“But that fire at State and then discovering an embedded agent recording keystroke/mouse clicks with likely transmission of surveillance files a Sunnyvale IP is beginning to hoist warning flags. I am keeping my ‘setup’ isolated from the Internet just in case there is something bogus boiling.

Following supper Lionel returns for a deep dive into the var/svn repository folder.

“So let’s move forward to the ’/var/svn’ folder. Frank wisely setup a instance of the svn configuration management application. He used it in three different modes.

“The first is the standard command line that allowed him to interact with large amounts of project artifacts. The command line version allowed him to push and pull his work, create new repositories for the project archives as he need them, and make global corrections and updates.

“The second mode was the html version of svn. This allowed him to navigate and rapidly view any specific project artifact in his collection.

“The final mode is a user friendly graphical mode or application called RabbitVCS. Frank used the RabbitVCS for the day to day configuration. With it he maintained archive control of his project artifacts and of the many changes to them. The RabbitVCS application hides the complexity of keeping project artifacts and their changes.

“For anyone like us coming behind Frank it provides the opportunity to glimpse into what he was doing in the project. RabbitVCS first provides a log file specific to Frank that maintains a record of events. Then RabbitVCS appends icons to each folder and file in Frank’s project work space. These icons present the developed with visual cue about the configuration status of folders and files.“

“Bottom-line if Frank ‘touched’ the artifact file and if he ever added that artifact at least once to the ‘Spinnaker’ project repository then we ‘see’ that artifact and all of the changes. There is a high probability that Frank did just that for every artifact he developed for ‘Spinnaker’. Plus the caveat is that if his artifacts needed data or supporting files from other ‘Spinnaker’ participants: imhotep, patel, judge, gator, and turtle Frank’s repository would also contain current version of those artifacts.

“The **BAD NEWS** is that we do not have Frank’s project workspace. If the ‘Spinnaker’ project was managed in a typical manner, the Catalina Software Engineering provided at least the project storage media. The closure email ,‘ 01/21/2021 from zbozz’, discussed inventorying and returning hardware and intellectual property. As such Catalina must have Frank’s workspace as it was on the storage media.

“**BUT** we have the local svn copy that Frank create in ‘/var/svn’ with all eleven repositories that he ever created. If we recall correctly we know there are eight which we identified previously that probably need follow-up. Triaging that eight there appear to be two which need priority: eclipse_devwrk/ and sloeber/. These two both use the ‘eclipse’ development environment that supports complex code projects and automated, daily ‘project build’. The other five are basically toy builders compared to these two.

“ We know we can display the contents of the last version of each project artifact that Fran sent to the Repositories using the local web page, *http://localhost/svn/*. Viewing the contents may provide us with our best chance of learning what the ‘Spinnaker’ was and what Frank’s contribution might have been.

“Opening a web session with *http://localhost/svn/* and selecting the first of the two most likely, repositories, ‘eclipse_devwrk/’. We find that only the trunk/ sub-folder is populated. This is not unusual since the ‘trunk/’ is where the main development resides. The other two are used for short term experiments, sometimes called spikes. Since the ‘Spinnaker’ project was near its final release and closure this would be an expected observation.

“Under ‘eclipse_devwrk/trunk’ eight projects are shown. Of these eight projects three have names which suggest a possible match or might relevant to ‘Spinnaker’. Those three are ‘e4Spn’, ‘html_index_rana’, and ‘uc2ump’. Without becoming bogged down we will spend a few minutes in each of of these three projects to see if we can gain any insight that might be helpful.

“First ‘e4Spn’ appears to be a Java language project using the MySQL database with a number of useful Apache web tools settings for controlling access, interfaces, and displays. Let’s see if we can quickly drill down to anything specific to Spinnaker or Catalina. Doing a quick ‘find in files’ using the ‘geany’ editor ‘grep’ for Spinnaker or Catalina produced 743 hits.

“CONCLUSION: This is the mother-lode which will need to be mined. But let’s also look at the other two.

“Looking next at the ‘eclipse_devwrk/trunk/html_index_rana’ we find a collection of web pages that appear to be Frank’s location where he places artifacts for his personal web page. Scanning there the collection we quickly see several branches to sub-web pages that deal with various interest he has. This repository is of little to no value in helping to understand the ‘Spinnaker’ project.

“Pushing late into the night now we will quickly look at ‘eclipse_devwrk/trunk/usc2ump’. This is a collection of web pages that documents a construction project in which Frank was once engaged. Once again this project will provide no information about ‘Spinnaker’”

“Based upon what we discovered tonight, all my attention tomorrow will be draw to ‘eclipse_devwrk/trunk/e4Spn’ svn project,” concludes Lionel as he shuts-down the audio/video equipment, turns-off the lights , and finally heads for the door.

“So Geek Detective is it time to fill the sandbags?” greets Sandi as Lionel enters the family room.

“It just might be, it might just be,” replies Lionel. “I really do not know what is going on. But I think I will be ‘packing’ the Springfield KDE in the AM when we go for coffee.”

“Come-on it can not be that grim.”

“First, State did have that fire. I think that one of Prof Yakubu’s students may have tripped the keystroke monitoring agent and alerted someone related to the ‘Spinnaker’ project that there was a new set of eyes looking over old information.

“Second Frank had an archive repository project which threw me 743 hits when I searched the project for ‘Spinnaker’ and ‘Catalina’.

Does that sound suspect enough for you?

“We carried while that nut-job Mudd was running around near Claude’s. I am afraid that it is time to be cautious again.”

“Anyway I am done for the night. Goodnight, Love” closes Lionel as he marches towards the bedroom.