



CSEC ITS/N2E Cyber Threat Discovery



DISCOCON 2010
A CND Perspective

*Safeguarding Canada's security through information superiority
Préserver la sécurité du Canada par la supériorité de l'information*



N2E Cyber Threat Discovery

- CSEC/ITS Discovery: Context & Sitrep
- Current Capabilities
- CND Metadata Analysis at Scale



N2E Context

- Within CSEC/ITS, CND operations concentrated in N2
 - Core: incident response team (alerts -> analysis -> mitigation advice)
 - Malware/RE & VR teams
- Discovery has always been a required activity within N2
 - IR often takes precedence
- N2E established to focus effort on discovery
 - Hunting vs. firefighting: new threats, techniques, tradecraft
 - Iterations of hypothesis based research/analysis
 - What evidence is there in the data of compromised systems?
 - What new threats or techniques can we find operating against us?
 - Development of effective anomaly detection/heuristic techniques
 - How can we better detect these new threats in the future?
 - Shape capability development priorities





N2E Sitrep

- Team formally established earlier this year
 - Growing to 7; hiring underway, strong candidates in pipeline
- Excellent access to full take data
- Analytical environment improving
- Progress on policy support
 - Use of intercepted private communications
 - Sharing mechanisms



Our Haystack

- Three individual clients
 - 10's - low 1000's of signature based alerts each day (majority false positive)
- Soon: “official” Internet gc.ca aggregation point
- Full pcaps (retention: days to months)
 - 1's - 10's TB of passively tapped network traffic each day
- Metadata (retention: months to years)
 - 10's - 100's GB of non-indexed, textual, network metadata / day (descriptive)





Toolset

- Homemade wrappers for heuristic detection tools
- Popquiz/Slipstream for metadata
- And...



```
$ find | xargs | grep | sed |  
cut | awk | sort | uniq -c |  
perl | python
```




Homemade Wrappers: PonyExpress SMTP processing

8Ball wrapper to process all email attachments


- Cuts and re-rebuilds SMTP session from full packet captures
- Extracts and logs metadata from SMTP and RFC822
- Extracts attachments and sends it to a cluster of 8Ball services
 - Wrapper able to manage other deep scan tools (e.g. AV)
 - Aggregates results from scans asynchronously
- Generates alerts for analysts to look at with full SMTP pcap
- Catches new implants or first stage delivered using attachments



Homemade Wrappers: Stripsearch

Automated binary analysis

- Recursively runs binary through the following tools:



- Correlation with file repository and reports
- Future: greater degree of automation




Metadata

- Standard Flows (with protocol guessing)
- DNS
 - Queries / Answers and beaconing
- SMTP
 - RFC 788 (SMTP), RFC 822 (2822, 5322) Internet Message Format
- HTTP
 - All server-side headers
 - User-Agent summary
 - List of POST without a GET, URLs and PE downloads



Metadata::findmask





Noteworthy Catches

WMI-based implant

- Trend Micro has a good report this type of implant
<http://us.trendmicro.com/imperial/md/content/us/trendwatch/researchanda>
- Uses WMI for persistence instead of Registry or Boot sector
 - Reside in ActiveScriptEventConsumer
 - First seen in September 2009
See RFI_20091023_280_1 and IXR_20100909_732_1
 - Caught with Pony Express
 - Stripsearch fuzzy correlation linked it with previous reports



Noteworthy Catches

Google IP addresses used as a sleep command

- From DNS metadata
 - When call back domain was queried, a Google IP was used as a sleep command. Searching for Google IPs in DNS metadata revealed new TROPUSNU domains and infected workstations.

Suspicious RDP and TOR sessions

- From Flow metadata we identified:

RDP

- Incoming RDP sessions from various locations to the RDP service (uses the same certificate)
- Still under investigation

TOR

- Was picked-up as a suspicious burst of outgoing SSL connections going to several locations from a single source



Noteworthy Catches: TOR...

- Most likely a breach of Internet usage policy
- All SSL CERTS were:
 - Self-signed
 - Certificate has a validity window of just 2 hours (e.g. 101017152421Z to 101017172421Z)
 - Issuer name appears to be a randomly generated fully qualified domain name (FQDN), unique for each destination IP (e.g. www.b2wwzduvdc5jyty3s7.net)
 - Common Name (CN) field appears to be a randomly generated FQDN unique for each session (e.g. www.xehzhl3cip.net)



Noteworthy Catches: TOR...

- All sessions for October 17 extracted from pcap using the following ngrep HEX string:

```
$ ngrep -I pcap -tq -X 0x5A170D313031303137
```

```
T 2010/10/17 15:49:00.699087 [REDACTED] -> X.X.X.43:61469 [AP]
```

```
(...
```

```
  ×...) [REDACTED] ..1010171524217..101017172421Z0%1#0!...U...www.b2wwzdu  
  vdc5jyty3s7.net0(...×...)
```

- The string is the end of validity period (Date only)
- Snort signature could be generated for looking at CERTS with validity period starting/ending same day



CND Metadata Analysis at Scale

Selected SAWUNEH Results



www.jolyon.co.uk



General Problem: Information Overload

- What is 'at scale' for CSEC CND?
- Problem: too much data
 - Acquired, retained, summarized, analyzed, presented to analysts
- Opportunities at all tiers of system. Selected SAWUNEH topics help in one or more of these areas



What is SAWUNEH?

- Literally: 'Summer Analytic Workshop Up North EH'
 - Annual data analysis workshop hosted at CSEC by crypt-math shop
 - 2010 thrust was CND
 - Reps from each of the 5 eyes + cleared researchers
- Helps fill the (sparsely inhabited) plane that often exists between OpsDev and R&D
- Special Assets: Netezza, Cray XMT



Topic 1: Correlating 'E-Mail, Web, Flow' Metadata

- E-mail based content delivery attacks against GC *very* common
- CND capability rapidly developed & deployed; heavily relies on attachment scanning
- The inevitable evasion
 - Attacks now commonly use email as inducement to URL visit, instead of direct content delivery.
- Too many benign hyperlinks delivered over email
 - Need to reduce metadata presented to front line analysts



Topic 1: Correlating 'E-Mail, Web, Flow' Metadata


Reducing Email URL metadata volume delivered to front line analysts

- Select highly suspicious subset of Email URL metadata based on correlation
- Only present 'Email URL' metadata if:
 - Email was inbound
 - Contained one or more hyperlinks
 - Hyperlink nominated as suspicious
 - Hyperlink was actually visited by a recipient.
 - Provide flow and http metadata for resulting HTTP session



Topic 1: Correlating 'E-Mail, Web, Flow' Metadata SAWUNEH approach

- Import EMAIL, URL, FLOW metadata into separate tables on the NETEZZA appliance (distributed db)
- Single SQL Query to fill requirement:





'Email + Web' attack detection SAWUNEH Style

Findings:

- 50% of the unique results were malicious
- Significantly lower false positive rates compared to URL inspection only
- Result provides anomaly tip but also context at fingertips for analyst

Next Steps:

- Better definition of 'suspicious' URL
- Automate extraction of web content from session for analysts
- Automate feeding of web payloads into content scanning system = partial mitigation of the original evasion. Higher analyst productivity



Topic 2: Malicious Attachment Prediction

Problem:

- Our current email scanning system **fully** processes **every** email it sees in a modular fashion
- Works very well today, but we can't afford to scale this approach 'as-is' to meet future requirements.
- Large proportion of time spent in *late* stages of processing.. especially deep scan

Possible Approach:

- Can we predict which emails will score in the deep scan based only on metadata extracted during earlier processing phases?
- If so, can we drop those that have low probability of deep scan success?




Malicious Attachment Prediction

- Given metadata from each processing phase, build a predictive model to selectively promote emails with high probably of non-zero scoring in 8-ball based on features extracted early in processing
- Choice of Predictive Model
 - Random Forests: decision tree based classifier (Brieman 2001)
- Take a feature set ($X_1, \dots, X_p$) and a corpus of training samples with *known* classification as input. Generate *Random Forest* of decision trees to be used to prediction classification
- New samples pass through previously trained forest which yields probability that deepscan will be nonzero



Malicious Attachment Prediction

Data Reduction vs. 'Interesting' Data Loss





Malicious Attachment Prediction

Result:

- 85% data reduction with 1-3% loss of 'interesting' emails
 - Can discard majority of emails with minimal loss in positive hits
- Prediction with model is relatively cheap (10,000s features per second)


Next steps:


- Extract those features which contribute most to prediction and incorporate appropriate filters in each stage of email processing system with aim to discard early and often
- There are likely a few simple features contributing disproportionately to predictive power: flow size, attachment size, attachment type etc



(Mini) Topic 3: PE Masquerade Detection

- Multiple threat actors observed masquerading Windows PE downloads using 'benign' filename extensions (jpg, gif, htm etc)
- We log metadata on HTTP sessions containing a PE header in content of first packet (modules available in both SLIPSTREAM and POPQUIZ)
 - Exclude entries with suffixes common to PE downloads (exe, bin, php, asp)
 - Provide analysts contextual metadata & 1 click access to payloads
- Easy to see entries that jumped out: file extensions such as .jpg, .gif highly suspicious, often malicious
- Also filter by 'uniqueness' over time to eliminate AV / OS updates
- Next Step: automate extraction of pe masquerades and push through content scanning system







Final Notes

SAWUNEH outcomes made a few things clearer to us

- Our days of plain old **grep** as primary search tool are nearing an end
- Relational DB's provide us with a few key benefits:
 - Simple indexing of existing metadata (sub second searches)
 - ability to correlate easily across 'primary' metadata outputs
 - significantly faster hypothesis testing
- Depending on scale, DB over head may be too costly
 - In such cases, still very useful tool for discovery work on finite snapshots
 - Case be used to test correlation hypothesis on finite data sets before spending significantly more cycles implementing a streaming / on-line version.
- Simple correlation techniques can provide high yield metadata to analysts
 - Islands of primary metadata are becoming unwieldy
 - Post processing of primary metadata (correlation, newness, uniqueness, reduction techniques etc) becoming a requirement to mitigate information overload.



Thanks

[REDACTED]
[REDACTED]@cse-cst.gc.ca

[REDACTED]
[REDACTED]@cse-cst.gc.ca

[REDACTED]
[REDACTED]@cse-cst.gc.ca



Communications Security
Establishment Canada

Centre de la sécurité
des télécommunications Canada

TOP SECRET

CSEC ITS/N2E Cyber Threat Discovery



DISCOCON 2010
A CND Perspective

Safeguarding Canada's security through information superiority
Préserver la sécurité du Canada par la supériorité de l'information

Canada



N2E Cyber Threat Discovery

- CSEC/ITS Discovery: Context & Sitrep
- Current Capabilities
- CND Metadata Analysis at Scale



N2E Context

- Within CSEC/ITS, CND operations concentrated in N2
 - Core: incident response team (alerts -> analysis -> mitigation advice)
 - Malware/RE & VR teams
- Discovery has always been a required activity within N2
 - IR often takes precedence
- N2E established to focus effort on discovery
 - Hunting vs. firefighting: new threats, techniques, tradecraft
 - Iterations of hypothesis based research/analysis
 - What evidence is there in the data of compromised systems?
 - What new threats or techniques can we find operating against us?
 - Development of effective anomaly detection/heuristic techniques
 - How can we better detect these new threats in the future?
 - Shape capability development priorities



N2E Sitrep

- Team formally established earlier this year
 - Growing to 7; hiring underway, strong candidates in pipeline
- Excellent access to full take data
- Analytical environment improving
- Progress on policy support
 - Use of intercepted private communications
 - Sharing mechanisms



Our Haystack

- Three individual clients
 - 10's - low 1000's of signature based alerts each day (majority false positive)
- Soon: "official" Internet gc.ca aggregation point
- Full pcaps (retention: days to months)
 - 1's - 10's TB of passively tapped network traffic each day
- Metadata (retention: months to years)
 - 10's - 100's GB of non-indexed, textual, network metadata / day (descript e)



Notes on gc.ca:

- System capacity at 400TB per month
- Expected to start at 150TB per month (68 departments) but expected to grow to capacity the next 2 years (>100 departments)



Toolset

- Homemade wrappers for heuristic detection tools
- Popquiz/Slipstream for metadata
- And...



```
$ find | xargs | grep | sed |  
cut | awk | sort | uniq -c |  
perl | python
```

We have a bunch of heuristic tools which are good at detecting unknown malware. They are mainly wrappers around tools like 8Ball and metadata produced out of raw network traffic.

Detection from wrappers handled through alert management system but metadata analyzed manually.

Using good old unix text filters!!!



Homemade Wrappers: PonyExpress SMTP processing

8Ball wrapper to process all email attachments

- Cuts and re-rebuilds SMTP session from full packet captures
- Extracts and logs metadata from SMTP and RFC822
- Extracts attachments and sends it to a cluster of 8Ball services
 - Wrapper able to manage other deep scan tools (e.g. AV)
 - Aggregates results from scans asynchronously
- Generates alerts for analysts to look at with full SMTP pcap
- Catches new implants or first stage delivered using attachments

- Currently processes 400,000 emails per day
 - These are all emails that were previously filtered by Anti-SPAM, Anti-Virus softwares
 - Out of these, 400 gets promoted to the alert system
 - 1% of these alerts are worthy of reporting to client
- With gc.ca coming
 - PonyExpress will be in front of Anti-Virus, Anti-SPAM due to location of collection points.
 - Looking at deploying appliances in front of PonyExpress



Homemade Wrappers: Stripsearch

Automated binary analysis

- Recursively runs binary through the following tools:



- Correlation with file repository and reports
- Future: greater degree of automation

- Modules are run in parallel and some sequentially: Filters are run first to exclude known good
 - Each module can be selectively disabled or enabled
- Recursive means when 8Ball finds an embedded file, it extracts it and runs it through stripsearch independently
- Currently processes our malware repository (~1500 samples) in 1h30 minutes.
- Capable of processing folders and subfolders



Metadata

- Standard Flows (with protocol guessing)
- DNS
 - Queries / Answers and beaconing
- SMTP
 - RFC 788 (SMTP), RFC 822 (2822, 5322) Internet Message Format
- HTTP
 - All server-side headers
 - User-Agent summary
 - List of POST without a GET, URLs and PE downloads

SMTP comes from PonyExpress

All others from popquiz

HTTP server headers and User-Agent summary are based on the first payload packet so when headers are using more than one packet data gets truncated. Sorting yields strange side effects like showing nice staircase-like output

User-Agent: b

User-Agent: bla

User-Agent: blab


User-Agent: blabl

User-Agent: blabla

Found some SQL injection attempts in Server String: 'DROP TABLE servertypes; --



Metadata::findmask



When we have more than 7 – 10 unique known strings in the same flow its usually worth investigating. Anything less is most likely a false positive.

Looks for XORed strings such as:

- Windows API calls
- Part of Registry keys



Noteworthy Catches

WMI-based implant

- Trend Micro has a good report this type of implant
<http://us.trendmicro.com/imperial/md/content/us/trendwatch/researchanda>
- Uses WMI for persistence instead of Registry or Boot sector
 - Reside in ActiveScriptEventConsumer
 - First seen in September 2009
See RFI_20091023_280_1 and IXR_20100909_732_1
 - Caught with Pony Express
 - Stripsearch fuzzy correlation linked it with previous reports

Since first version, we saw increase in sophistication

At first WMI event created through a script run by
CSCRIPT.EXE

Now, is able to create the WMI event from the binary



Noteworthy Catches

Google IP addresses used as a sleep command

- From DNS metadata
 - When call back domain was queried, a Google IP was used as a sleep command. Searching for Google IPs in DNS metadata revealed new TROPUSNU domains and infected workstations.

Suspicious RDP and TOR sessions

- From Flow metadata we identified:
 - RDP
 - Incoming RDP sessions from various locations to the RDP service (uses the same certificate)
 - Still under investigation
 - TOR
 - Was picked-up as a suspicious burst of outgoing SSL connections going to several locations from a single source



Noteworthy Catches: TOR...

- Most likely a breach of Internet usage policy
- All SSL CERTS were:
 - Self-signed
 - Certificate has a validity window of just 2 hours (e.g. 101017152421Z to 101017172421Z)
 - Issuer name appears to be a randomly generated fully qualified domain name (FQDN), unique for each destination IP (e.g. www.b2wwzduvdc5jty3s7.net)
 - Common Name (CN) field appears to be a randomly generated FQDN unique for each session (e.g. www.xehzh13cip.net)



Noteworthy Catches: TOR...

- All sessions for October 17 extracted from pcap using the following ngrep HEX string:

```
$ ngrep -I pcap -tq -X 0x5A170D313031303137
```

```
T 2010/10/17 15:49:00.699087 [REDACTED] -> X.X.X.43:61469 [AP]
```

```
(...  
>[REDACTED]...1010171524217...101017172421Z0%1#0!...U...www.b2wwzdu  
vdc5jyty3s7.net0(...>[REDACTED])
```

- The string is the end of validity period (Date only)
- Snort signature could be generated for looking at CERTS with validity period starting/ending same day

The easiest signature would be to have snort look at starting end ending in the same month. However, the signature will have to be changed every month.

A search for October only looking at CERTs ending in October (1010) resulted in a lot of false positives but easily managed by greping start date of October. Only one false positive left.



CND Metadata Analysis at Scale

Selected SAWUNEH Results



www.sajon.co.uk

Safeguarding Canada's security through information superiority
Préserver la sécurité du Canada par la supériorité de l'information

Canada



General Problem: Information Overload

- What is 'at scale' for CSEC CND?
- Problem: too much data
 - Acquired, retained, summarized, analyzed, presented to analysts
- Opportunities at all tiers of system. Selected SAWUNEH topics help in one or more of these areas

Today:

1-10's TB pcap retained each day

10's-100's GB of metadata retained each day

Future:

metadata continues to increase linearly with new access points

attempt to hold pcap steady at current rates through smarter retention policies

Tiers:



What is SAWUNEH?

- Literally: 'Summer Analytic Workshop Up North EH'
 - Annual data analysis workshop hosted at CSEC by crypt-math shop
 - 2010 thrust was CND
 - Reps from each of the 5 eyes + cleared researchers
- Helps fill the (sparsely inhabited) plane that often exists between OpsDev and R&D
- Special Assets: Netezza, Cray XMT

Many researchers but focus is 'applied'

Netezza: distributed data appliance made available to facilitate workshop

OpsDev vs R&D rather than Ops vs R&D



Topic 1: Correlating 'E-Mail, Web, Flow' Metadata

- E-mail based content delivery attacks against GC *very* common
- CND capability rapidly developed & deployed; heavily relies on attachment scanning
- The inevitable evasion
 - Attacks now commonly use email as inducement to URL visit, instead of direct content delivery.
- Too many benign hyperlinks delivered over email
 - Need to reduce metadata presented to front line analysts

PonyExpress:

Yields many detections with low false positive rates

Commercial sector has responded in similar fashion

Email as inducement

Delivery of HTTP URL over E-mail, exploit is delivered via HTTP



Topic 1: Correlating 'E-Mail, Web, Flow' Metadata

Reducing Email URL metadata volume delivered to front line analysts

- Select highly suspicious subset of Email URL metadata based on correlation
- Only present 'Email URL' metadata if:
 - Email was inbound
 - Contained one or more hyperlinks
 - Hyperlink nominated as suspicious
 - Hyperlink was actually visited by a recipient.
 - Provide flow and http metadata for resulting HTTP session

Initial definition of 'suspicious' was naive: suffix dictionary

Have actual visit helps by:


more likely to be successful attack

means analyst likely has access to related web traffic to assist with triage



Topic 1: Correlating 'E-Mail, Web, Flow' Metadata SAWUNEH approach

- Import EMAIL, URL, FLOW metadata into separate tables on the NETEZZA appliance (distributed db)
- Single SQL Query to fill requirement:





'Email + Web' attack detection SAWUNEH Style

Findings:

- 50% of the unique results were malicious
- Significantly lower false positive rates compared to URL inspection only
- Result provides anomaly tip but also context at fingertips for analyst

Next Steps:

- Better definition of 'suspicious' URL
- Automate extraction of web content from session for analysts
- Automate feeding of web payloads into content scanning system = partial mitigation of the original evasion. Higher analyst productivity

Possible 'suspicious' definitions: URLs tagged with 'ObfJavaScript, FlashContent, LedToDownload' etc
Correlated on URL 'trees' not isolated URLs (otherwise redirects, remote hrefs etc defeat scanning)..
But somewhat expensive (not sure if this is solved in feasible way ?)

primary metadata sources too large for discovery... need value-add reductions
reduced 'working sets' of metadata = higher yield for analyst time.

They evade by moving payload delivery to different session. We follow them and (attempt) to close evasion.



Topic 2: Malicious Attachment Prediction

Problem:

- Our current email scanning system **fully** processes **every** email it sees in a modular fashion
- Works very well today, but we can't afford to scale this approach 'as-is' to meet future requirements.
- Large proportion of time spent in *late* stages of processing.. especially deep scan

Possible Approach:

- Can we predict which emails will score in the deep scan based only on metadata extracted during earlier processing phases?
- If so, can we drop those that have low probability of deep scan success?

Fully process = sessionize, smtp, rfc822, mime, deep scan(8Ball) of all attach...etc

Optimization axiom: find your hotspot



Malicious Attachment Prediction


- Given metadata from each processing phase, build a predictive model to selectively promote emails with high probability of non-zero scoring in 8-ball based on features extracted early in processing
- Choice of Predictive Model
 - Random Forests: decision tree based classifier (Breiman 2001)
- Take a feature set $(X_1, \dots, X_p)$ and a corpus of training samples with *known* classification as input. Generate *Random Forest* of decision trees to be used to prediction classification
- New samples pass through previously trained forest which yields probability that deepscan will be nonzero

Phases: (ip, tcp, smtp, rfc822, mime etc)

Classifier:

Based on probability of scoring non-zero in deep scan

Example features include:




In our case: email metadata types + > 10000
previously scanned samples with full metadata
including overall scan score used to build the Forest

Note: Only build/train forest once. Subsequently you
are just passing metadata through and getting an
 $P(\text{email scores non-zero})$



Malicious Attachment Prediction Data Reduction vs. 'Interesting' Data Loss



Safeguarding Canada's security through information superiority
Préserver la sécurité du Canada par la supériorité de l'information

Canada



Malicious Attachment Prediction

Result:

- 85% data reduction with 1-3% loss of 'interesting' emails
 - Can discard majority of emails with minimal loss in positive hits
- Prediction with model is relatively cheap (10,000s features per second)

Next steps:

- Extract those features which contribute most to prediction and incorporate appropriate filters in each stage of email processing system with aim to discard early and often
- There are likely a few simple features contributing disproportionately to predictive power: flow size, attachment size, attachment type etc

Could be used either permanently, or simply as a learning tool to find which features you might want to filter on.

Emphasis on feature predictors available early in processing (flow, smtp, 822 header rather than MIME content etc)



(Mini) Topic 3: PE Masquerade Detection


- Multiple threat actors observed masquerading Windows PE downloads using 'benign' filename extensions (jpg, gif, htm etc)
- We log metadata on HTTP sessions containing a PE header in content of first packet (modules available in both SLIPSTREAM and POPQUIZ)
 - Exclude entries with suffixes common to PE downloads (exe, bin, php, asp)
 - Provide analysts contextual metadata & 1 click access to payloads
- Easy to see entries that jumped out: file extensions such as .jpg, .gif highly suspicious, often malicious
- Also filter by 'uniqueness' over time to eliminate AV / OS updates
- Next Step: automate extraction of pe masquerades and push through content scanning system



Communications Security
Establishment Canada

Centre de la sécurité
des télécommunications Canada

TOP SECRET



Safeguarding Canada's security through information superiority
Préserver la sécurité du Canada par la supériorité de l'information

Canada


27



Communications Security
Establishment Canada

Centre de la sécurité
des télécommunications Canada

TOP SECRET



Safeguarding Canada's security through information superiority
Préserver la sécurité du Canada par la supériorité de l'information

Canada

28



Final Notes

SAWUNEH outcomes made a few things clearer to us

- Our days of plain old **grep** as primary search tool are nearing an end
- Relational DB's provide us with a few key benefits:
 - Simple indexing of existing metadata (sub second searches)
 - ability to correlate easily across 'primary' metadata outputs
 - significantly faster hypothesis testing
- Depending on scale, DB over head may be too costly
 - In such cases, still very useful tool for discovery work on finite snapshots
 - Case be used to test correlation hypothesis on finite data sets before spending significantly more cycles implementing a streaming / on-line version.
- Simple correlation techniques can provide high yield metadata to analysts
 - Islands of primary metadata are becoming unwieldy
 - Post processing of primary metadata (correlation, newness, uniqueness, reduction techniques etc) becoming a requirement to mitigate information overload.



Thanks

[REDACTED]
[REDACTED]@cse-cst.gc.ca

[REDACTED]
[REDACTED]@cse-cst.gc.ca

[REDACTED]
[REDACTED]@cse-cst.gc.ca