US-CERT:
Netflow visualization

Presented by Aaron Bossert and Jerry Derrick
10 January 2012

#Agenda

```
if($attendance > 0) {
    print "- Introduction\n";
    print "- The back-end\n";
    print "- Things to consider\n";
    print "- The front-end\n";
    print "- The Good, Bad, and Ugly\n";
    print "- Next steps\n";
    print "Exiting\n";
    exit(0);
}
else {
    print "See ya! Going to Starbucks\n";
}
```

The Back-end

- Don’t worry, pretty pictures coming next…
- Netezza™ Database Appliance
  - OLAP, not OLTP
  - Billions of records and 100’s thousands queries per day
  - Scales linearly on processing and storage
  - Open connectivity standards (ODBC, JDBC)
- Data enrichment to provide context
- Data aggregates or cubes

The Back-end (Crayola™ version)

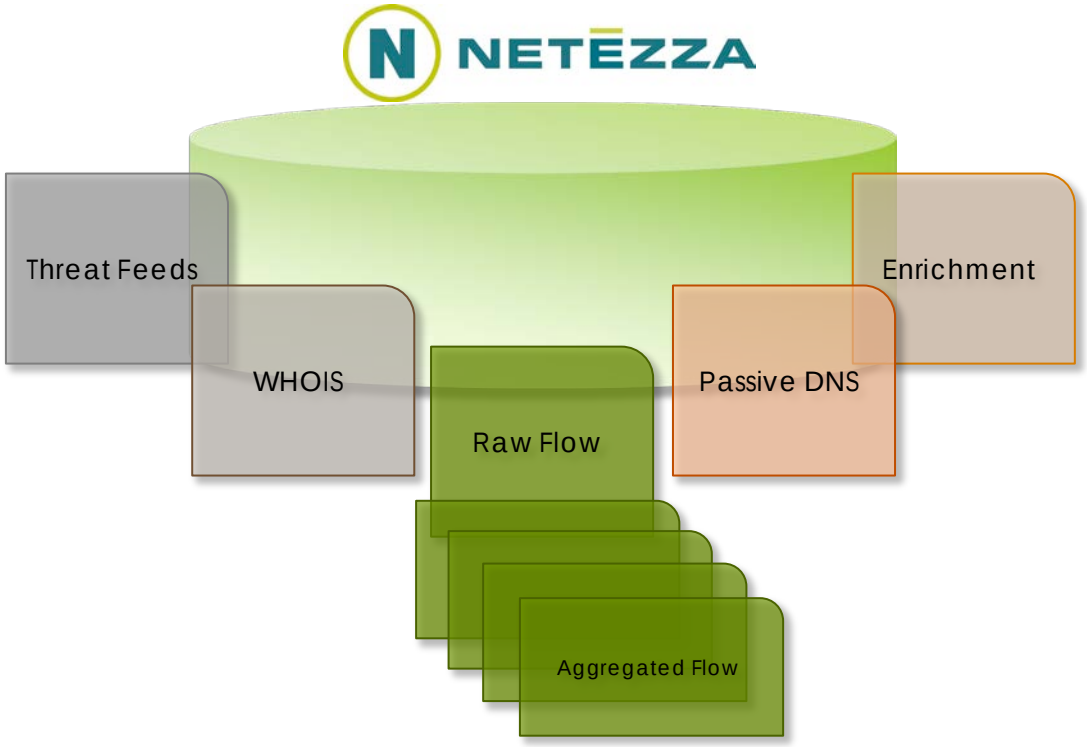

Notional Data Flow/Architecture

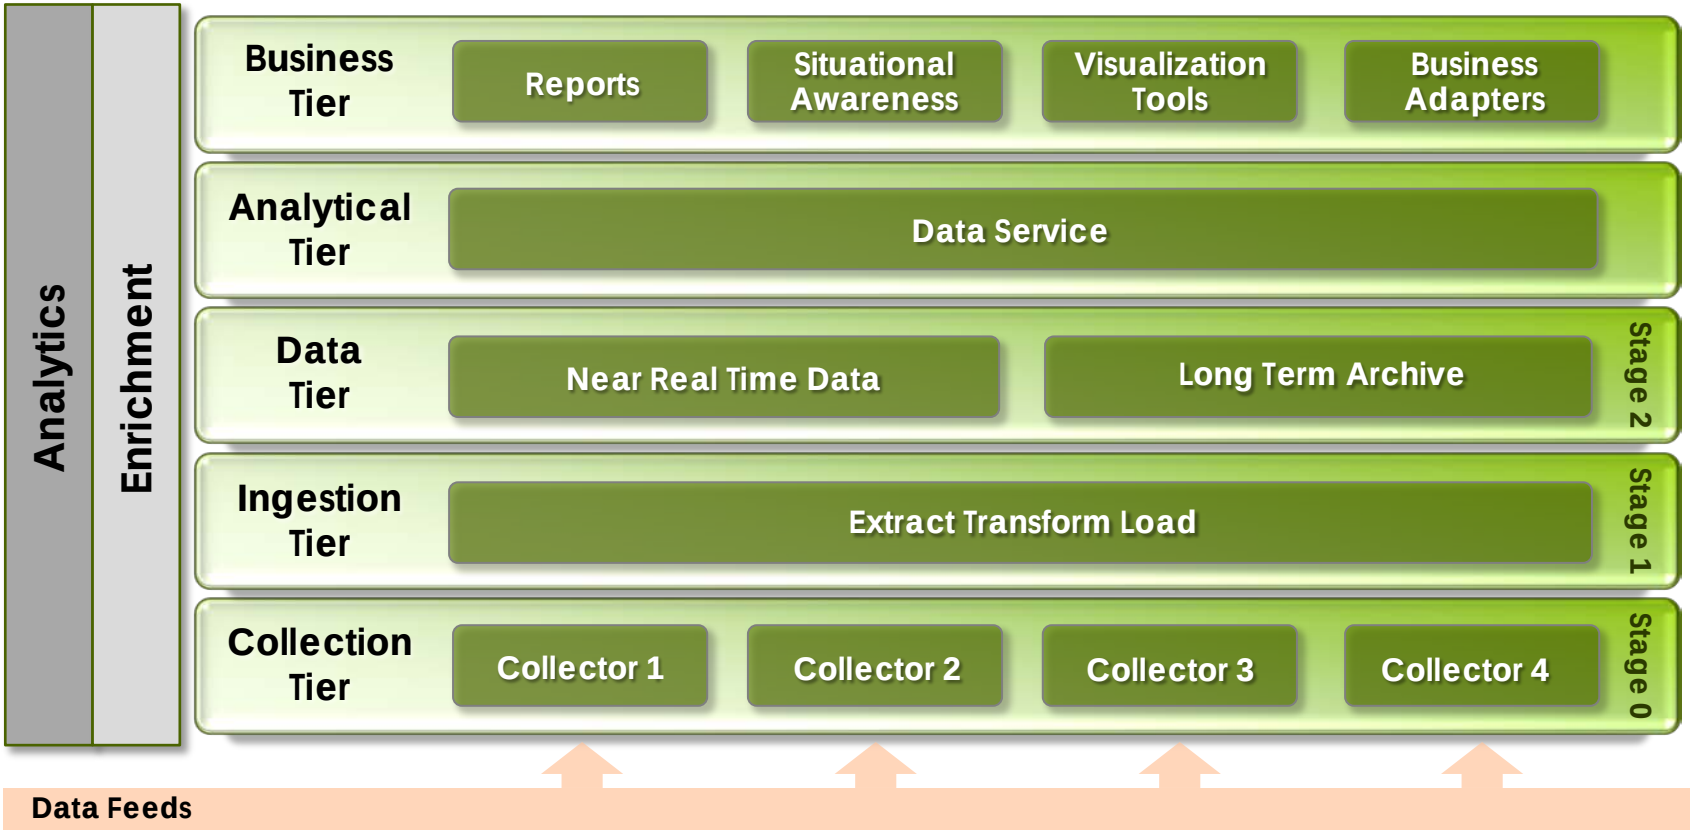

Visualization Considerations

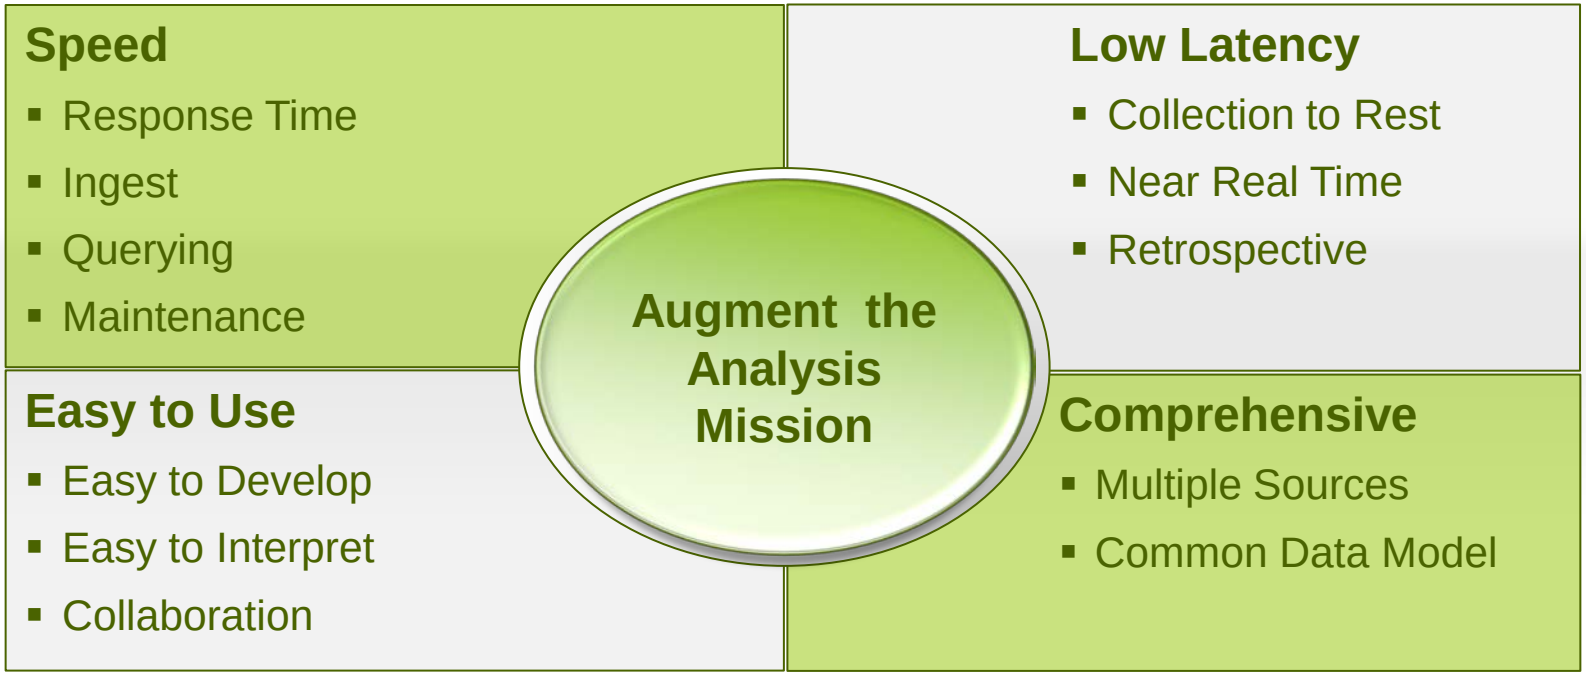

Tool Considerations

- Client or Server Based
  - Where is the data?
  - What is the “state” of your data?
  - How responsive does the system need to be?
- What does the visualization do for you?
  - Provides an alternative way of viewing the problem
  - Confirm/refute suspected activity
  - You haven’t seen what you haven’t seen

The Front-end

- Provides logical workflow
- Visual pattern recognition
  - Quick culling of benign traffic
  - Outliers stand out
  - Easier to see relationships
- Can be used to train junior analysts
- Repeatable results
- Interface can be tuned for different applications

The Front-end (Situational Awareness)

The Front-end (Link Analysis)

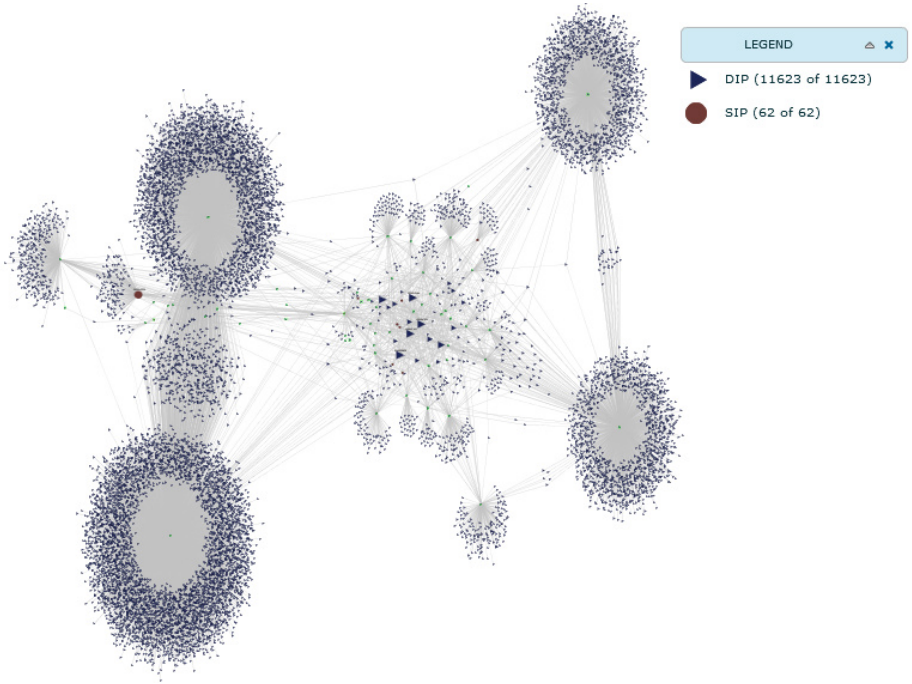

- Explore relationships
- Drill and pivot
- Identify bad and non-obvious actors
- Questions lead to more questions
- Minimize false positives
- Validate findings

The Front-end (Customized Analysis)

The Good, Bad, and Ugly

- GOOD!
  - Single pane of glass
  - More efficient workflow
  - Shorter learning curve for junior analysts
  - Knowledge transfer from senior analysts
  - Pretty
- BAD!
  - Massive data-sets are still massive (even with fusion IO cards)
  - Less flexible than command-line
  - If what you want to look at is not pre-computed, go get a coffee, or five…
- UGLY!
  - Operating on the edge of what is possible
  - No one vendor provides an ideal front-end solution
  - Custom applications take time

Next Steps

- Settle on a front-end application
- Add white-listing on the fly
- Automate pattern recognition (got your slide-ruler?)
- Tighter integration with Indicator database

Questions?