Root zone scalability model

Bart Gijsen

October 28, 2009

TNO | Kennis voor zaken

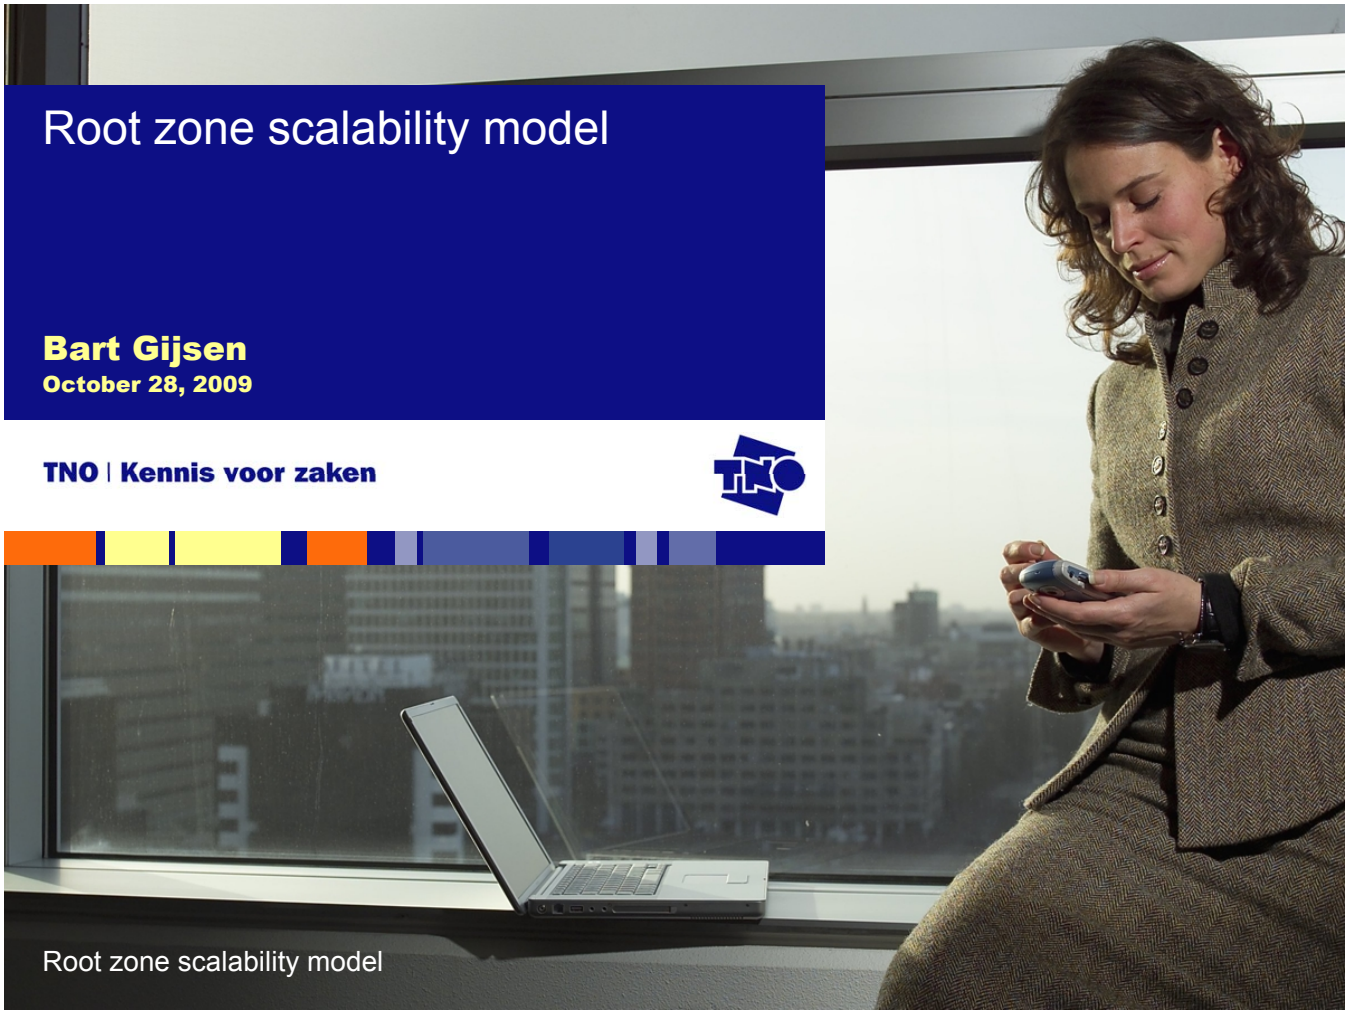

Introduction

- Development of the model by TNO *as part of* the Root Scalability Study Team
- Why quantify?

  Scalability is a quantitative topic

- What's the challenge?

  "*The challenge is to reap sound insight and understanding from simulations, while never mistaking for the simulation real world*."

  [FloydPaxson01, Simulating The Internet]

Goal of the quantitative model

- Root Scaling Study Terms of Reference
  - Primary deliverable: model of the root server system
    - showing how different parts of the system are related
    - impact of changing (combinations of) parameter values on all parts of the system
    - the model should be as quantitative as possible
    - use of the model: clarify consequences of policy decisions about the root
    - it should not try to answer: “how much is too many?”
  - Impact of growth scenarios (“Plus 1”, “Plus 2” and “Plus 4”)

- The quantitative model investigates the scalability:
  1. The parameters that dominantly influence the scalability are not a priori known => model will help to indentify them
  2. Once the scalability is understood, the model will be applied to quantify the scalability boundaries

Developing the quantitative model (1/2)

- The quantitative model is based on
  - Narratives from the Root Scaling Study Team
  - Terms of reference of ICANN
- Observed information deficiencies:
  - Some information regarding processes was not available, conflicting, or subject to change in very near future
  - Failure rates in provisioning and publication process are unknown
  - Measurement data of zone file distribution is fragmented
- Scalability questions to be answered require diverging model output metrics
  - Resource load, lead times, several types of error probabilities, and more?
  - Consequence w.r.t. model analysis techniques => use one *analytical* model per 1 or 2 metrics, or a single *simulation* based model

Developing the quantitative model (2/2)

- **Consequently, the modelling approach was chosen such that:**

- Model is easily adjustable during its development
  - Hierarchical modeling
  - Separation between workflow and resources layers
  - Use block/object oriented, event-driven simulation SW package (ExtendSim)
- Modeled processes are recognizable (enable review/feedback)
  - Simulation of workflow with graphical interface and animation
- Input parameter policy:
  - Include enough parameters to enable investigation of relevant questions,
  - While keeping the total number of input parameters as low as possible
- Model based sensitivity analysis allows to:
  - Refine the model itself and
  - Estimating the scalability ranges and numerical confidence intervals

Chosen scope of the scalability model

- Quantitative analysis of the scalability of the root-zone file provisioning and publication process

Qualitative reasoning and rough estimating within RSST pinpointed these processes as most likely bottlenecks

Overview level: workflow layer

- The root scaling model consists of the following parts:
  - Provisioning process of TLD change requests
    - receiving change requests by IANA
    - IANA – NTIA/DoC – VeriSign validation checks
  - Root-zone file publication
    - production of the zone file
    - distribution to the RSO's
- The events in the event-driven simulation model are …
  - provisioning side: TLD change requests, distinguished per type (variable rate)
  - publication side: root zone files (twice a day, variable size)

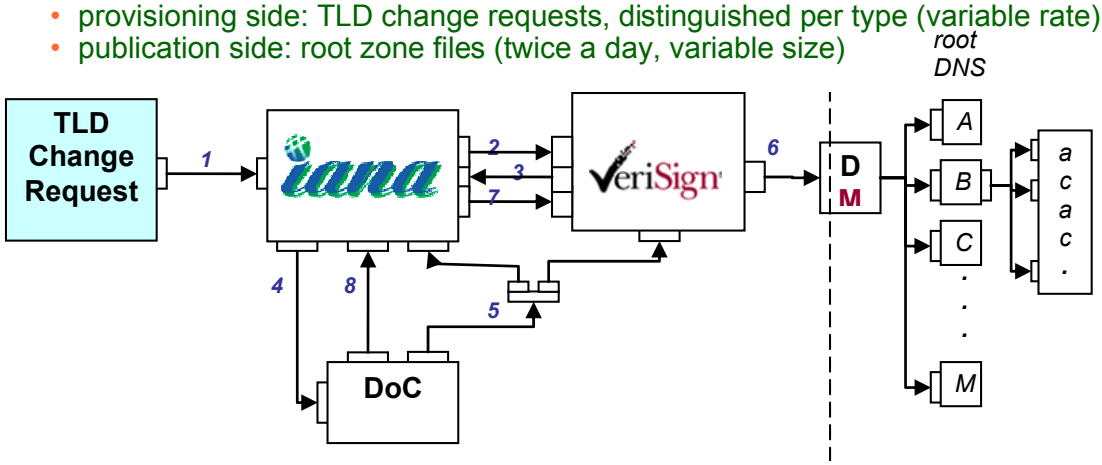

Provisioning side Publication side

ExtendSim model screenshot: top-level view

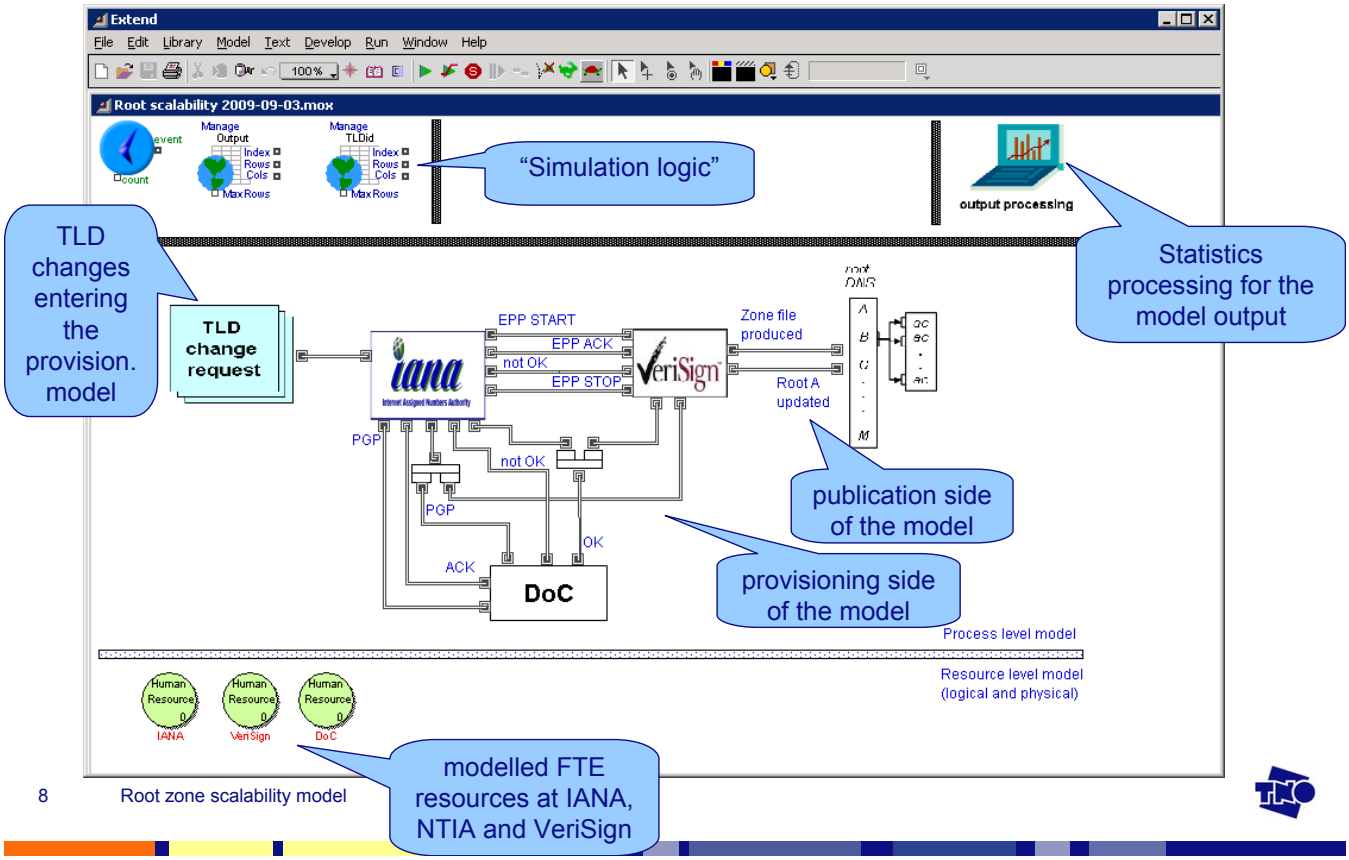

ExtendSim model screenshots
Generation of TLD requests

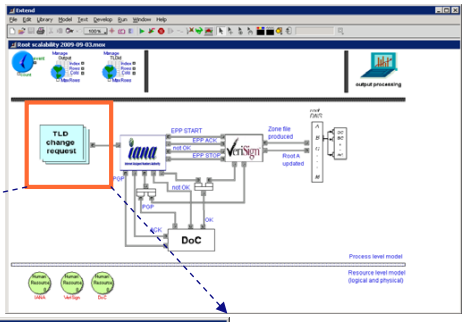

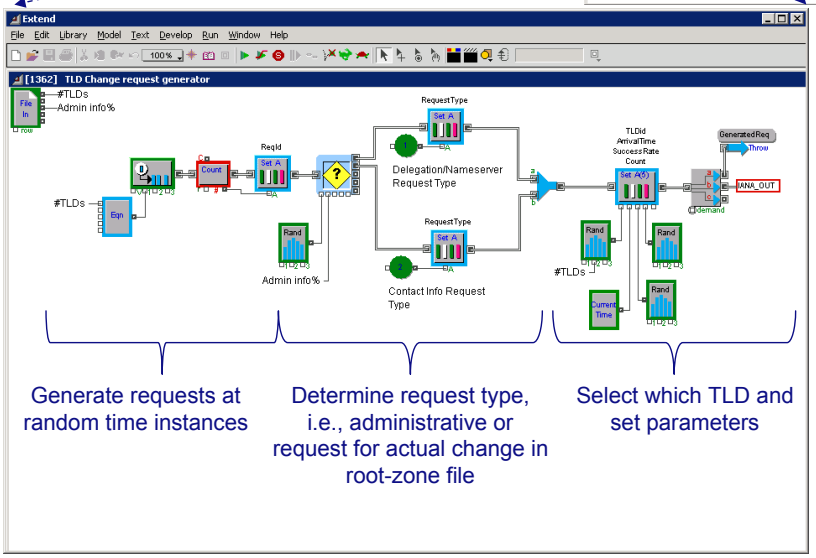

ExtendSim model screenshots

VeriSign processing of requests

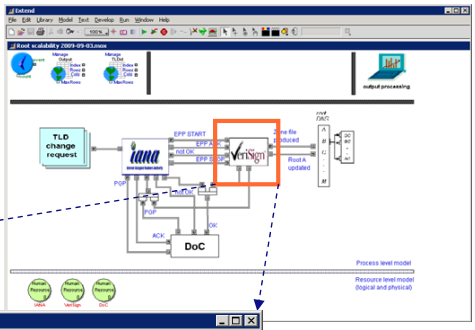

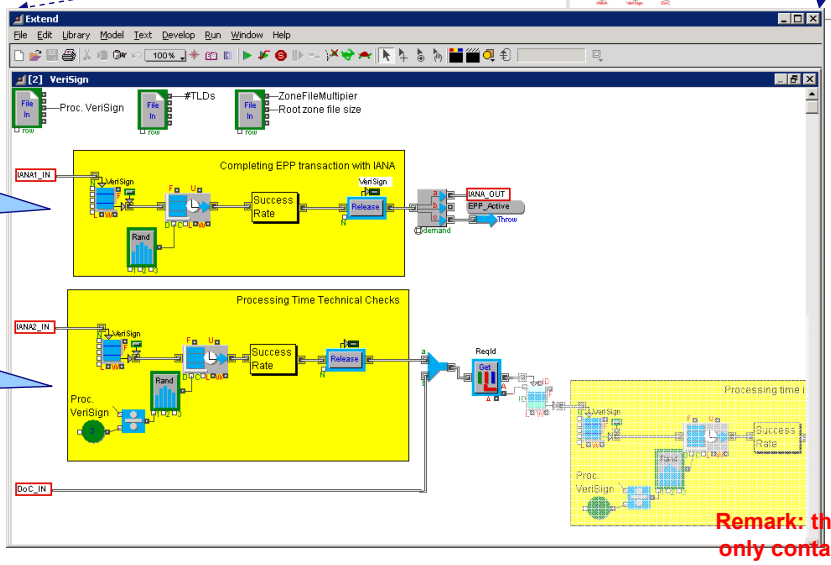

EPP transaction with DoC

Processing of technical checks

Remark: this display only contains a part of the modelled processes at VeriSign

ExtendSim model screenshots
Publication process

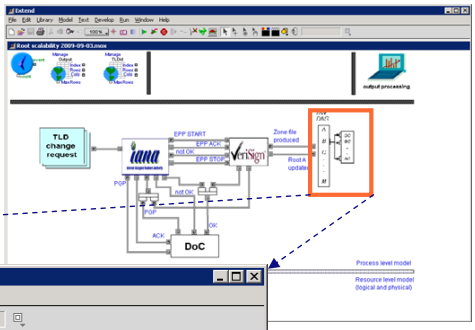

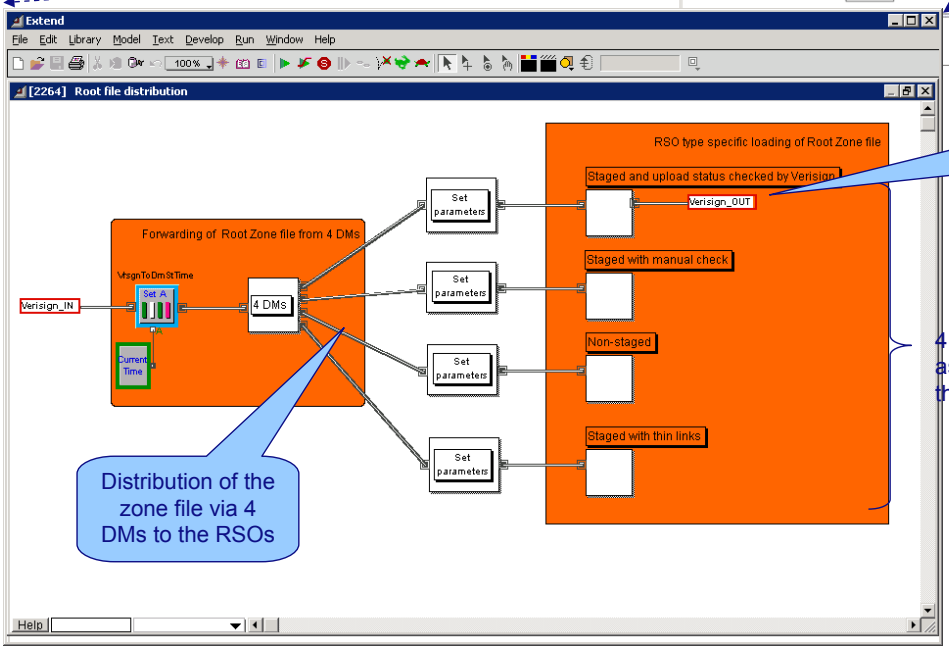

Feedback to the provisioning part of the model

4 types of assemblies of the RSOs

Distribution of the zone file via 4 DMs to the RSOs

Root-zone file publication process

- Two examples (out of four) RSO assemblies and the modelling
  - In the model we confine to the successful retrieval to a single name server

RSO assembly | Model

- **RSO with staged cluster with check from VeriSign**

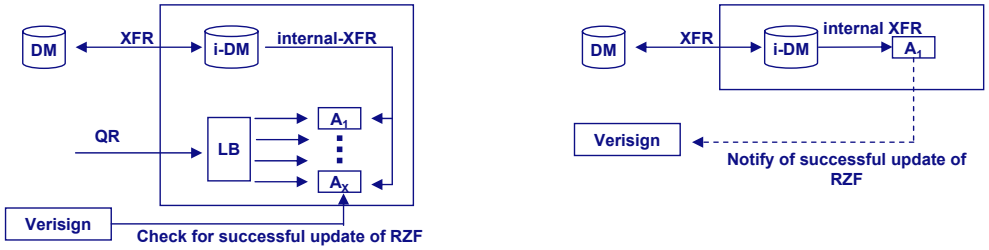

- **RSO with non-staged cluster**

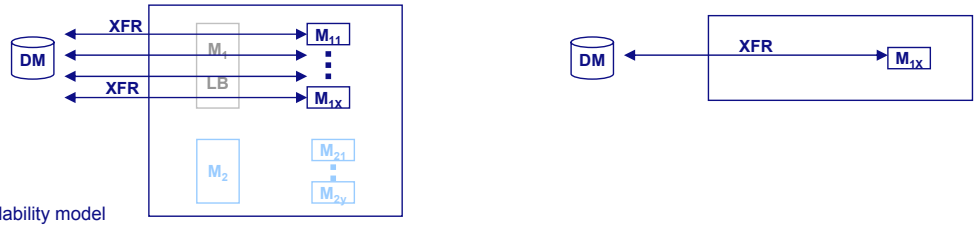

Model input and outputs

Input

- Provisioning
  - # TLDs
  - TLD change request rate
  - Fraction of “Administrative info” changes
  - Processing times at IANA, NTIA, VeriSign
  - Available FTE capacity
  - # Authorization checks per change request
  - Office hours for manual actions
- Error model in provisioning process
  - Incremental error rate per manual action
- Publication
  - Normalized root zone file size
  - File size multiplier (e.g., #TLD, DNSSEC)
  - Round-Trip Time (for DNS notify)
  - Packet-loss probability (for DNS notify)
  - DNS / SOA Number of attempts
  - DNS / SOA time-out value
  - XFR Connection goodput (Mbit/s)
  - XFR success probability

Output

- Provisioning
  - Lead time of provisioning side
  - Load on each of the manual resources
- Error rates in provisioning process
  - Cumulative error rate in provisioning process
- Publication
  - Zone file loading time in publication process

Model inputs

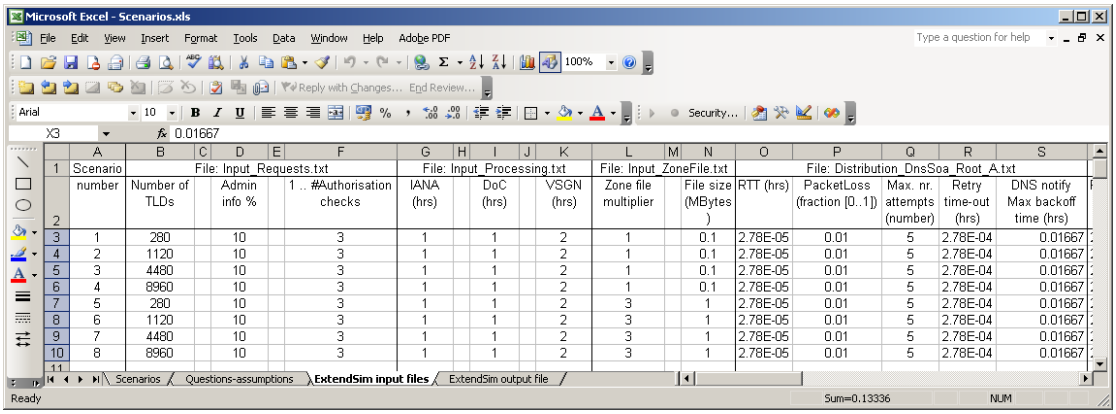

Microsoft Excel - Scenarios.xls

Scenario	File: Input_Requests.txt				File: Input_Processing.txt			File: Input_ZoneFile.txt		File: Distribution_DnsSoa_Root_A.txt				
number	Number of TLDs		Admin info %	1 .. #Authorisation checks	IANA (hrs)	DoC (hrs)	VSGN (hrs)	Zone file multiplier	File size (MBytes)	RTT (hrs)	PacketLoss (fraction [0..1])	Max. nr. attempts (number)	Retry time-out (hrs)	DNS notify Max backoff time (hrs)
1	280		10	3	1	1	2	1	0.1	2.78E-05	0.01	5	2.78E-04	0.01667
2	1120		10	3	1	1	2	1	0.1	2.78E-05	0.01	5	2.78E-04	0.01667
3	4480		10	3	1	1	2	1	0.1	2.78E-05	0.01	5	2.78E-04	0.01667
4	8960		10	3	1	1	2	1	0.1	2.78E-05	0.01	5	2.78E-04	0.01667
5	280		10	3	1	1	2	3	1	2.78E-05	0.01	5	2.78E-04	0.01667
6	1120		10	3	1	1	2	3	1	2.78E-05	0.01	5	2.78E-04	0.01667
7	4480		10	3	1	1	2	3	1	2.78E-05	0.01	5	2.78E-04	0.01667
8	8960		10	3	1	1	2	3	1	2.78E-05	0.01	5	2.78E-04	0.01667

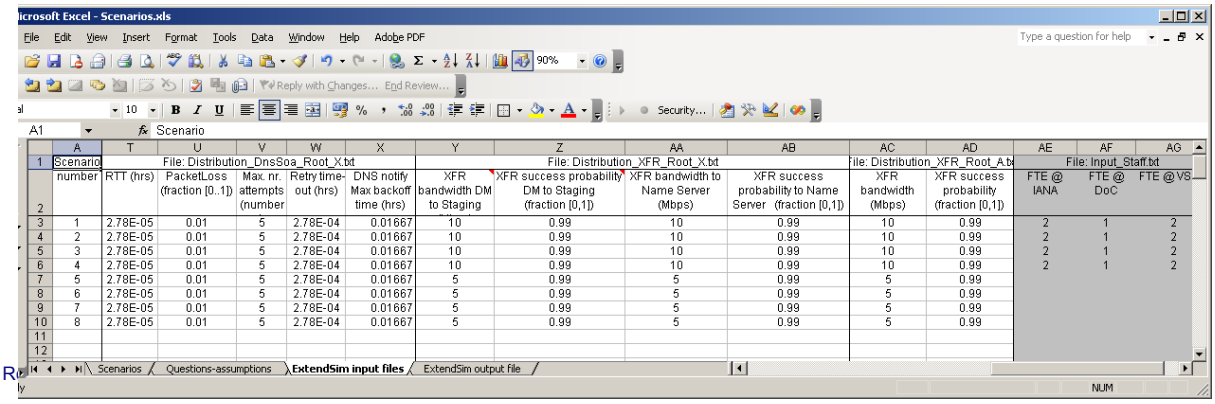

Microsoft Excel - Scenarios.xls

Scenario	File: Distribution_DnsSoa_Root_X.txt					File: Distribution_XFR_Root_X.txt				File: Distribution_XFR_Root_A.tx		File: Input_Staff.txt		
number	RTT (hrs)	PacketLoss (fraction [0..1])	Max. nr. attempts (number	Retry time-out (hrs)	DNS notify Max backoff time (hrs)	XFR bandwidth DM to Staging	XFR success probability DM to Staging (fraction [0,1])	XFR bandwidth to Name Server (Mbps)	XFR success probability to Name Server (fraction [0,1])	XFR bandwidth (Mbps)	XFR success probability (fraction [0,1])	FTE @ IANA	FTE @ DoC	FTE @VS
1	2.78E-05	0.01	5	2.78E-04	0.01667	10	0.99	10	0.99	10	0.99	2	1	2
2	2.78E-05	0.01	5	2.78E-04	0.01667	10	0.99	10	0.99	10	0.99	2	1	2
3	2.78E-05	0.01	5	2.78E-04	0.01667	10	0.99	10	0.99	10	0.99	2	1	2
4	2.78E-05	0.01	5	2.78E-04	0.01667	10	0.99	10	0.99	10	0.99	2	1	2
5	2.78E-05	0.01	5	2.78E-04	0.01667	5	0.99	5	0.99	5	0.99			
6	2.78E-05	0.01	5	2.78E-04	0.01667	5	0.99	5	0.99	5	0.99			
7	2.78E-05	0.01	5	2.78E-04	0.01667	5	0.99	5	0.99	5	0.99			
8	2.78E-05	0.01	5	2.78E-04	0.01667	5	0.99	5	0.99	5	0.99			

Model output

- Output parameters focused on:
  - load of the resources
  - provisioning process and publication lead times
  - error propagation probabilities
- Benefit of chosen simulation approach: adaption of model output metrics is very easy
- Choice to implement model in ExtendSim provides graphical interface and animation ‘as a bonus’
  - this enhances insight in the modeled processes

Example of results of the simulation model

Scenario	# TLD´s	File size		Connection quality
1	280	0.1 MB	3 MB	Good
2	1120	0.4 MB	12 MB	Good
3	4480	1.6 MB	48 MB	Good
4	8960	3.2 MB	96 MB	Good

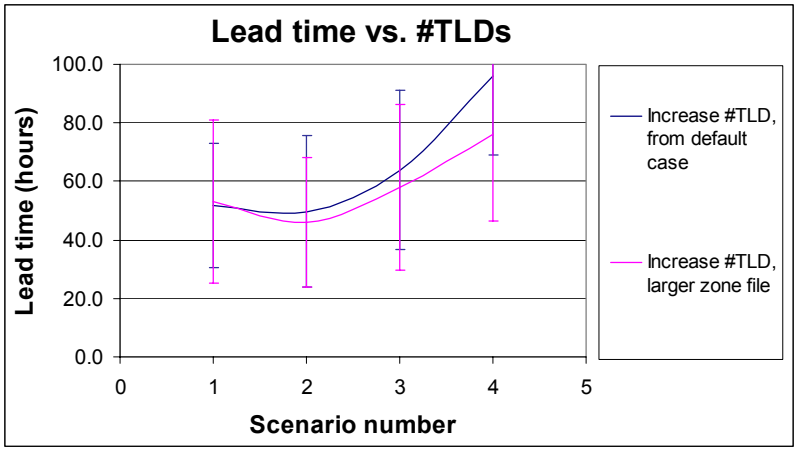

Conclusions

- Simulation model is developed and applied for scalability analysis
  - model specifies the current understanding of the TLD change provisioning and zone file publication process => “**base-line model**”
  - improving quality of model input data remains a challenge (“rubish-in = rubish-out”)
- Preliminary results from simulated cases support the conclusion in the Scaling the Root report
  - current processes can cope with addition of hunderds of TLDs
  - when adding thousands of TLDs resource capacity upgrades will become necessary

Recommended next steps

A. **Start collecting monitoring data for the root system** in order to get (a) reliable quantitative data and (b) experience with their trend patterns
   - The model input and output parameters are a starting point for the metrics to monitor; further investigation needed to find the most appropriate set

B1. **Validate and fine-tune the model**
   - Using the collected quantitative data and the more specific intended use of the model

B2. Cover the risk of quantitative numbers: *Do not pretend to be more predictive / accurate, than the quantitative facts allow you to be!* => **analyze sensitivity of the model input parameters to estimate the numerical confidence intervals**

C. **Detail the quantitative root-scaling analysis** to obtain more accurate boundaries for the scalability
   - Start simple, start with first-order-statistic: load on resources